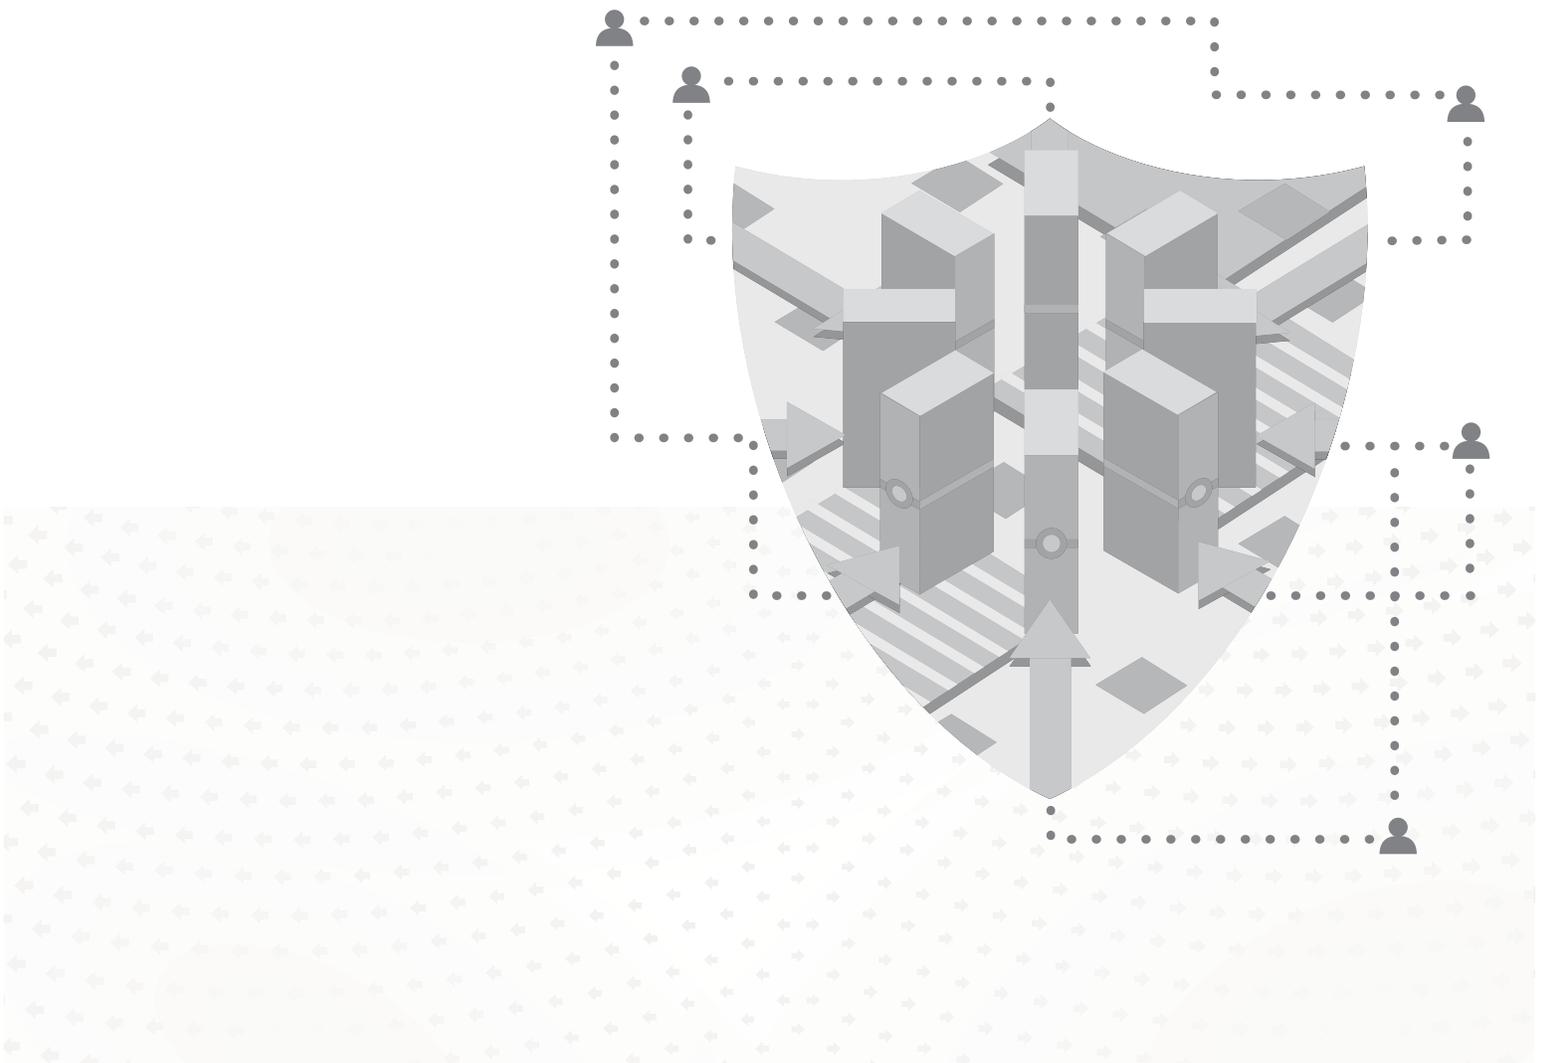


Sourcefire 3D System

User Guide

Version 5.2



Legal Notices

Cisco, the Cisco logo, Sourcefire, the Sourcefire logo, Snort, the Snort and Pig logo, and certain other trademarks and logos are trademarks or registered trademarks of Cisco and/or its affiliates in the United States 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.

The legal notices, disclaimers, terms of use, and other information contained herein (the "terms") apply only to the information discussed in this documentation (the "Documentation") and your use of it. These terms do not apply to or govern the use of websites controlled by Cisco or its subsidiaries (collectively, "Cisco") or any Sourcefire-provided or Cisco-provided products. Sourcefire and Cisco products are available for purchase and subject to a separate license agreement and/or terms of use containing very different terms and conditions.

The copyright in the Documentation is owned by Cisco and is protected by copyright and other intellectual property laws of the United States and other countries. You may use, print out, save on a retrieval system, and otherwise copy and distribute the Documentation solely for non-commercial use, provided that you (i) do not modify the Documentation in any way and (ii) always include Cisco's copyright, trademark, and other proprietary notices, as well as a link to, or print out of, the full contents of this page and its terms.

No part of the Documentation may be used in a compilation or otherwise incorporated into another work or with or into any other documentation or user manuals, or be used to create derivative works, without the express prior written permission of Cisco. Cisco reserves the right to change the terms at any time, and your continued use of the Documentation shall be deemed an acceptance of those terms.

© 2004 - 2014 Cisco and/or its affiliates. All rights reserved.

Disclaimers

THE DOCUMENTATION AND ANY INFORMATION AVAILABLE FROM IT MAY INCLUDE INACCURACIES OR TYPOGRAPHICAL ERRORS. CISCO MAY CHANGE THE DOCUMENTATION FROM TIME TO TIME. CISCO MAKES NO REPRESENTATIONS OR WARRANTIES ABOUT THE ACCURACY OR SUITABILITY OF ANY CISCO-CONTROLLED WEBSITE, THE DOCUMENTATION AND/OR ANY PRODUCT INFORMATION. CISCO-CONTROLLED WEBSITES, THE DOCUMENTATION AND ALL PRODUCT INFORMATION ARE PROVIDED "AS IS" AND CISCO DISCLAIMS ANY AND ALL EXPRESS AND IMPLIED WARRANTIES, INCLUDING BUT NOT LIMITED TO WARRANTIES OF TITLE AND THE IMPLIED WARRANTIES OF MERCHANTABILITY AND/OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL CISCO BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING BUT NOT LIMITED TO PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOSS OF DATA, LOSS OF PROFITS, AND/OR BUSINESS INTERRUPTIONS), ARISING OUT OF OR IN ANY WAY RELATED TO CISCO-CONTROLLED WEBSITES OR THE DOCUMENTATION, NO MATTER HOW CAUSED AND/OR WHETHER BASED ON CONTRACT, STRICT LIABILITY, NEGLIGENCE OR OTHER TORTUOUS ACTIVITY, OR ANY OTHER THEORY OF LIABILITY, EVEN IF CISCO IS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES/JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATIONS MAY NOT APPLY TO YOU.

2014-Mar-25 10:22

Table of Contents

Chapter 1:	Introduction to the Sourcefire 3D System	37
	Sourcefire 3D System Appliances	38
	Defense Centers.....	39
	Managed Devices	39
	Understanding Appliance Series, Models, and Capabilities.....	40
	Sourcefire 3D System Components	44
	Redundancy and Resource Sharing	44
	Network Traffic Management	45
	FireSIGHT.....	46
	Access Control.....	46
	Intrusion Detection and Prevention	47
	File Tracking, Control, and Malware Protection.....	48
	Application Programming Interfaces.....	49
	Security, Internet Access, and Communication Ports.....	50
	Internet Access Requirements	50
	Open Communication Ports Requirements	52
	Documentation Resources	54
	Documentation Conventions	54
	License Conventions.....	55
	Supported Device and Defense Center Conventions	56
	Access Conventions	56
	IP Address Conventions.....	57
	Logging into the Appliance	58
	Logging into the Appliance to Set Up an Account.....	60
	Logging Out of the Appliance	62
	Using the Context Menu	63

Chapter 2:	Using Dashboards	66
	Understanding Dashboard Widgets.....	70
	Understanding Widget Availability	71
	Understanding Widget Preferences	73
	Understanding the Predefined Widgets	74
	Understanding the Appliance Information Widget.....	75
	Understanding the Appliance Status Widget.....	76
	Understanding the Correlation Events Widget	77
	Understanding the Current Interface Status Widget	77
	Understanding the Current Sessions Widget	78
	Understanding the Custom Analysis Widget.....	79
	Understanding the Disk Usage Widget	96
	Understanding the Interface Traffic Widget.....	97
	Understanding the Intrusion Events Widget.....	97
	Understanding the Network Compliance Widget.....	99
	Understanding the Product Licensing Widget.....	101
	Understanding the Product Updates Widget.....	101
	Understanding the RSS Feed Widget.....	103
	Understanding the System Load Widget.....	104
	Understanding the System Time Widget	104
	Understanding the White List Events Widget	105
	Working with Dashboards	106
	Creating a Custom Dashboard.....	106
	Viewing Dashboards	109
	Modifying Dashboards.....	111
	Deleting a Dashboard	116
Chapter 3:	Using the Context Explorer.....	117
	Understanding the Context Explorer	119
	Understanding the Traffic and Intrusion Event Counts Time Graph	119
	Understanding the Network Information Section	120
	Understanding the Application Information Section	126
	Understanding the Intrusion Information Section.....	131
	Understanding the Files Information Section	135
	Understanding the Geolocation Information Section.....	140
	Understanding the URL Information Section.....	143
	Refreshing the Context Explorer.....	146
	Setting the Context Explorer Time Range	147
	Minimizing and Maximizing Context Explorer Sections	147
	Drilling Down on Context Explorer Data	148
	Working with Filters in the Context Explorer	150
	Adding and Applying Filters	151
	Creating Filters with the Context Menu	155
	Bookmarking Filters	156

Chapter 4:	Using Objects and Security Zones.....	158
	Using the Object Manager.....	159
	Grouping Objects.....	159
	Browsing, Sorting, and Filtering Objects	161
	Working with Network Objects	161
	Working with Security Intelligence Lists and Feeds	162
	Working with the Global Whitelist and Blacklist	166
	Working with the Sourcefire Intelligence Feed.....	167
	Working with Custom Security Intelligence Feeds.....	168
	Manually Updating Security Intelligence Feeds.....	170
	Working with Custom Security Intelligence Lists.....	170
	Working with Port Objects.....	173
	Working with VLAN Tag Objects	174
	Working with URL Objects	175
	Working with Application Filters	176
	Working with the Global Malware Whitelist	180
	Adding a File to the Whitelist from the Event View.....	181
	Adding a File to the Whitelist by Uploading the File	182
	Adding a File to the Whitelist Using the SHA-256 Value	184
	Modifying Files on the Global Malware List	185
	Working with Security Zones.....	187
Chapter 5:	Managing Devices.....	190
	Management Concepts	191
	What Can Be Managed by a Defense Center?	191
	Beyond Policies and Events	192
	Using Redundant Defense Centers	193
	Working in NAT Environments.....	193
	Configuring High Availability	194
	Using High Availability.....	195
	Guidelines for Implementing High Availability	199
	Setting Up High Availability	200
	Monitoring and Changing High Availability Status	202
	Disabling High Availability and Unregistering Devices	204
	Pausing Communication Between Paired Defense Centers.....	205
	Restarting Communication Between Paired Defense Centers.....	205
	Working with Devices.....	206
	Understanding the Device Management Page.....	206
	Adding Devices to the Defense Center	208
	Applying Changes to Devices	211
	Using the Device Management Revision Comparison Report.....	212
	Deleting Devices.....	212

Table of Contents

Configuring Remote Management	213
Editing Remote Management.....	216
Changing the Management Port.....	217
Managing Device Groups	217
Adding Device Groups	218
Editing Device Groups	218
Deleting Device Groups.....	219
Clustering Devices.....	220
Establishing Device Clusters	223
Editing Device Clusters.....	224
Configuring Individual Devices in a Cluster.....	225
Configuring Individual Device Stacks in a Cluster	226
Configuring Interfaces on a Clustered Device	227
Switching the Active Peer in a Cluster	228
Placing a Clustered Device into Maintenance Mode	229
Replacing a Device in a Clustered Stack.....	229
Establishing Clustered State Sharing.....	230
Troubleshooting Clustered State Sharing.....	233
Separating Clustered Devices.....	236
Managing Stacked Devices.....	237
Establishing Device Stacks	239
Editing Device Stacks	241
Configuring Individual Devices in a Stack	242
Configuring Interfaces on a Stacked Device	243
Separating Stacked Devices	243
Editing Device Configuration	244
Editing Assigned Device Names.....	244
Enabling and Disabling Device Licenses.....	246
Editing Device System Settings.....	247
Viewing the Health of a Device.....	248
Editing Device Management Settings	249
Understanding Advanced Device Settings.....	250
Editing Advanced Device Settings	251
Configuring Fast-Path Rules.....	253
Configuring Interfaces	258
Configuring the Management Interface.....	260
Configuring HA Link Interfaces	262
Configuring the Interface MTU	263
Disabling Interfaces	264
Preventing Duplicate Connection Logging.....	265
Chapter 6: Setting Up an IPS Device	266
Understanding Passive IPS Deployments.....	266
Configuring Passive Interfaces	267

	Understanding Inline IPS Deployments	269
	Configuring Inline Interfaces	269
	Configuring Inline Sets.....	271
	Viewing Inline Sets	271
	Adding Inline Sets.....	271
	Configuring Advanced Inline Set Options	275
	Deleting Inline Sets.....	279
Chapter 7:	Setting Up Virtual Switches	281
	Configuring Switched Interfaces.....	282
	Configuring Physical Switched Interfaces	283
	Adding Logical Switched Interfaces.....	285
	Deleting Logical Switched Interfaces.....	287
	Configuring Virtual Switches	288
	Viewing Virtual Switches.....	288
	Adding Virtual Switches	289
	Configuring Advanced Virtual Switch Settings	291
	Deleting Virtual Switches	294
Chapter 8:	Setting Up Virtual Routers.....	295
	Configuring Routed Interfaces	296
	Configuring Physical Routed Interfaces	296
	Adding Logical Routed Interfaces	300
	Deleting Logical Routed Interfaces.....	304
	Configuring SFRP.....	304
	Configuring Virtual Routers	306
	Viewing Virtual Routers.....	307
	Adding Virtual Routers	307
	Setting up DHCP Relay	310
	Setting up Static Routes	313
	Setting up Dynamic Routing	316
	Setting up RIP Configuration.....	316
	Setting Up OSPF Configuration	323
	Setting up Virtual Router Filters	335
	Adding Virtual Router Authentication Profiles	339
	Viewing Virtual Router Statistics.....	340
	Deleting Virtual Routers	341
	Configuring NAT on a Virtual Router	341
	Understanding the NAT Entry List.....	343
	Adding NAT Entries.....	344
	Deleting NAT Entries	350

Chapter 9:	Setting Up Hybrid Interfaces	352
	Adding Logical Hybrid Interfaces	352
	Deleting Logical Hybrid Interfaces	356
Chapter 10:	Using Gateway VPN	358
	Understanding IPSec	359
	Understanding IKE	359
	Understanding VPN Deployments	360
	Understanding Point-to-Point VPN Deployments	360
	Understanding Star VPN Deployments	360
	Understanding Mesh VPN Deployments	361
	Managing VPN Deployments	362
	Configuring VPN Deployments	363
	Configuring Advanced VPN Deployment Settings	374
	Applying a VPN Deployment	377
	Viewing VPN Deployment Status	377
	Viewing VPN Statistics and Logs	378
	Using the VPN Deployment Comparison View	381
Chapter 11:	Using NAT Policies	383
	Planning and Implementing a NAT Policy	384
	Configuring NAT Policies	385
	Managing NAT Policy Targets	387
	Organizing Rules in a NAT Policy	389
	Working with NAT Rule Warnings and Errors	390
	Managing NAT Policies	391
	Creating a NAT Policy	392
	Editing a NAT Policy	394
	Copying a NAT Policy	395
	Viewing a NAT Policy Report	396
	Comparing Two NAT Policies	397
	Applying a NAT Policy	401
	Creating and Editing NAT Rules	404
	Understanding NAT Rule Types	406
	Understanding NAT Rule Conditions and Condition Mechanics	409
	Understanding NAT Rule Conditions	410
	Adding Conditions to NAT Rules	411
	Searching NAT Rule Condition Lists	413
	Adding Literal Conditions to NAT Rules	414
	Using Objects in NAT Rule Conditions	415

	Working with Different Types of Conditions in NAT Rules.....	415
	Adding Zone Conditions to NAT Rules.....	415
	Adding Source Network Conditions to Dynamic NAT Rules.....	418
	Adding Destination Network Conditions to NAT Rules.....	419
	Adding Port Conditions to NAT Rules	421
Chapter 12:	Using Access Control Policies	424
	Configuring Policies	426
	Setting the Default Action.....	428
	Logging Connections for the Default Action.....	431
	Using Custom User Roles with Access Control Policies	433
	Managing Policy Targets	434
	Adding an HTTP Response Page	436
	Filtering Traffic Based on Security Intelligence Data.....	438
	Configuring Advanced Access Control Policy Settings	447
	Organizing Rules in a Policy	449
	Working with Rule Categories	451
	Searching for Rules.....	453
	Filtering Rules by Device	454
	Working with Warnings and Errors	455
	Managing Access Control Policies.....	456
	Creating an Access Control Policy	458
	Editing an Access Control Policy.....	459
	Copying an Access Control Policy.....	460
	Viewing an Access Control Policy Report	461
	Comparing Two Access Control Policies.....	462
	Applying an Access Control Policy.....	466
Chapter 13:	Understanding and Writing Access Control Rules	472
	Creating and Editing Access Control Rules	474
	Understanding Rule Actions	478
	Understanding Rule Conditions and Condition Mechanics	482
	Understanding Rule Conditions	483
	Adding Rule Conditions.....	485
	Searching Condition Lists	489
	Adding Literal Conditions.....	490
	Using Objects in Conditions	491

Table of Contents

Working with Different Types of Conditions	491
Adding Zone Conditions.....	492
Adding Network Conditions.....	493
Adding VLAN Tag Conditions.....	495
Adding User Conditions.....	497
Working with Application Conditions.....	498
Adding Port Conditions.....	504
Adding URL Conditions.....	506
Performing File and Intrusion Inspection on Allowed Traffic.....	511
Logging Connection, File, and Malware Information.....	515
Adding Comments to a Rule.....	521
Chapter 14:	
 Configuring External Alerting.....	524
Working with Alert Responses	526
Creating an Email Alert Response	527
Creating an SNMP Alert Response.....	528
Creating a Syslog Alert Response.....	530
Modifying an Alert Response	534
Deleting an Alert Response	534
Enabling and Disabling Alert Responses	534
Configuring Impact Flag Alerting	535
Configuring Discovery Event Alerting	536
Configuring Advanced Malware Protection Alerting.....	537
Chapter 15:	
 Working With Connection Data.....	539
Understanding Connection Data.....	540
Understanding Connection Summaries	542
Connection Data Fields.....	543
Information Available in Connection Events.....	550
Uses for Connection Data in the Sourcefire 3D System.....	554
Viewing Connection Data Graphs and Tables	555
Working with Connection Graphs	555
Changing the Graph Type.....	557
Selecting Datasets.....	560
Viewing Information About Aggregated Connection Data	563
Manipulating a Connection Graph on a Workflow Page.....	563
Drilling Down Through Connection Data Graphs	564
Recentering and Zooming on Line Graphs	565
Selecting Data to Graph.....	565
Detaching Connection Graphs	569
Exporting Connection Data.....	569

Table of Contents

	Working with Connection Data Tables	570
	Working with Connection Events Associated with Monitor Rules	571
	Viewing Files Detected in a Connection	572
	Viewing Intrusion Events Associated with a Connection	574
	Searching for Connection Data	575
	Viewing the Connection Summary Page	578
Chapter 16:	Introduction to Sourcefire Intrusion Prevention	581
	Understanding How Traffic Is Analyzed	583
	Capturing and Decoding Packets	584
	Processing Packets	585
	Generating Events	586
	Analyzing Intrusion Event Data	587
	Using Intrusion Event Responses	588
	Understanding Intrusion Prevention Deployments	588
	The Benefits of Custom Intrusion Policies	591
Chapter 17:	Working with Intrusion Events.....	593
	Viewing Intrusion Event Statistics	595
	Host Statistics.....	596
	Event Overview	597
	Event Statistics	598
	Viewing Intrusion Event Performance.....	598
	Generating Intrusion Event Performance Statistics Graphs.....	599
	Viewing Intrusion Event Graphs.....	601
	Viewing Intrusion Events	602
	Understanding Intrusion Events	604
	Viewing Connection Data Associated with Intrusion Events	610
	Reviewing Intrusion Events	611
	Understanding Workflow Pages for Intrusion Events	612
	Using Drill-Down and Table View Pages	616
	Using the Packet View	622
	Viewing Event Information.....	625
	Viewing Frame Information.....	634
	Viewing Data Link Layer Information	635
	Viewing Network Layer Information	636
	Viewing Transport Layer Information	638
	Viewing Packet Byte Information.....	641

	Using Impact Flags to Evaluate Events	641
	Searching for Intrusion Events	644
	Using the Clipboard	652
	Generating Clipboard Reports.....	652
	Deleting Events from the Clipboard.....	654
Chapter 18:	Handling Incidents.....	656
	Incident Handling Basics.....	657
	Definition of an Incident.....	657
	Common Incident Handling Processes.....	657
	Incident Types in the Sourcefire 3D System.....	660
	Creating an Incident.....	661
	Editing an Incident	663
	Generating Incident Reports.....	664
	Creating Custom Incident Types.....	665
Chapter 19:	Configuring Intrusion Policies	667
	Planning and Implementing an Intrusion Policy	668
	Managing Intrusion Policies	670
	Creating an Intrusion Policy	672
	Editing an Intrusion Policy.....	675
	Using the Navigation Panel.....	678
	Committing Intrusion Policy Changes	679
	Reapplying an Intrusion Policy	680
	Viewing an Intrusion Policy Report	682
	Comparing Two Intrusion Policies	685
	Setting Drop Behavior in an Inline Deployment	689
	Understanding the Base Policy	691
	Using Default Intrusion Policies	692
	Using a Custom Base Policy	693
	Allowing Rule Updates to Modify the Base Policy.....	694
	Selecting the Base Policy.....	695
	Accepting Rule Setting Changes from a Custom Base Policy	696
	Managing Variables.....	698
	Understanding Existing Variables.....	699
	Modifying Variables.....	701
	Creating New Variables.....	703
	Deleting Unused Variables.....	709
	Understanding Custom Variables.....	710
	Defining IP Addresses and Ports for Your Network	711
	Defining IP Addresses in Variables and Rules.....	713
	Defining Ports in Variables and Rules	715

Chapter 20:	Managing Rules in an Intrusion Policy	717
	Understanding Intrusion Prevention Rule Types	718
	Viewing Rules in an Intrusion Policy	719
	Sorting the Rule Display.....	722
	Viewing Rule Details	723
	Filtering Rules in an Intrusion Policy	729
	Understanding Rule Filtering in an Intrusion Policy.....	730
	Setting a Rule Filter in an Intrusion Policy	741
	Setting Rule States	743
	Filtering Intrusion Event Notification Per Policy	746
	Configuring Event Thresholding.....	747
	Configuring Suppression Per Intrusion Policy	753
	Adding Dynamic Rule States	756
	Understanding Dynamic Rule States	757
	Setting a Dynamic Rule State	758
	Adding Alerts	761
	Adding SNMP Alerts	761
	Adding Rule Comments.....	762
	Managing FireSIGHT Rule State Recommendations	764
	Understanding Basic Rule State Recommendations	765
	Understanding Advanced Rule State Recommendations	766
	Using FireSIGHT Recommendations	768
Chapter 21:	Using Advanced Settings in an Intrusion Policy	773
	Modifying Advanced Settings	774
	Understanding Preprocessors	780
	Meeting Traffic Challenges with Preprocessors	781
	Understanding Preprocessor Execution Order	782
	Reading Preprocessor Events	784
	Automatically Enabling Advanced Settings	787
	Understanding Troubleshooting Options	790
Chapter 22:	Using Layers in an Intrusion Policy	792
	Understanding Intrusion Policy Layers.....	792
	Sharing Layers	794
	Using Rules in Layers	795
	Removing Multi-Layer Rule Settings.....	797
	Using the FireSIGHT Recommendations Layer	799
	Using Layers with Advanced Settings	800
	Configuring User Layers	803

Chapter 23: Using Application Layer Preprocessors 809

- Decoding DCE/RPC Traffic 810
 - Selecting Global DCE/RPC Options 811
 - Understanding Target-Based DCE/RPC Server Policies 813
 - Understanding DCE/RPC Transports..... 814
 - Selecting DCE/RPC Target-Based Policy Options 818
 - Configuring the DCE/RPC Preprocessor 821
- Detecting Exploits in DNS Name Server Responses..... 827
 - Understanding DNS Preprocessor Resource Record Inspection 827
 - Detecting Overflow Attempts in RData Text Fields 829
 - Detecting Obsolete DNS Resource Record Types..... 829
 - Detecting Experimental DNS Resource Record Types 830
 - Configuring the DNS Preprocessor..... 830
- Decoding FTP and Telnet Traffic..... 832
 - Understanding Global FTP and Telnet Options 832
 - Configuring Global FTP/Telnet Options 833
 - Understanding Telnet Options 835
 - Configuring Telnet Options 836
 - Understanding Server-Level FTP Options 838
 - Configuring Server-Level FTP Options 842
 - Understanding Client-Level FTP Options 846
 - Configuring Client-Level FTP Options 847
- Decoding HTTP Traffic 850
 - Selecting Global HTTP Normalization Options 851
 - Configuring Global HTTP Configuration Options 853
 - Selecting Server-Level HTTP Normalization Options 854
 - Selecting Server-Level HTTP Normalization Encoding Options..... 862
 - Configuring HTTP Server Options..... 865
 - Enabling Additional HTTP Inspect Preprocessor Rules 868
- Using the Sun RPC Preprocessor 869
 - Configuring the Sun RPC Preprocessor 870
- Decoding the Session Initiation Protocol 872
 - Selecting SIP Preprocessor Options 873
 - Configuring the SIP Preprocessor..... 875
 - Enabling Additional SIP Preprocessor Rules 876
- Configuring the GTP Command Channel..... 878
- Decoding IMAP Traffic 880
 - Selecting IMAP Preprocessor Options 881
 - Configuring the IMAP Preprocessor 882
 - Enabling Additional IMAP Preprocessor Rules 884
- Decoding POP Traffic 884
 - Selecting POP Preprocessor Options 885
 - Configuring the POP Preprocessor..... 887
 - Enabling Additional POP Preprocessor Rules 889

Decoding SMTP Traffic	889
Understanding SMTP Decoding	890
Configuring SMTP Decoding	895
Enabling SMTP Maximum Decoding Memory Alerting	899
Detecting Exploits Using the SSH Preprocessor	899
Selecting SSH Preprocessor Options	901
Configuring the SSH Preprocessor	903
Using the SSL Preprocessor	905
Understanding SSL Preprocessing	905
Enabling SSL Preprocessor Rules	907
Configuring the SSL Preprocessor	907
Working with SCADA Preprocessors	909
Configuring the Modbus Preprocessor	909
Configuring the DNP3 Preprocessor	911

Chapter 24: Using Transport & Network Layer Preprocessors 915

Verifying Checksums	915
Ignoring VLAN Headers	917
Normalizing Inline Traffic	918
Understanding Protocol Normalization	919
Configuring Inline Normalization	922
Defragmenting IP Packets	927
Understanding IP Fragmentation Exploits	928
Target-Based Defragmentation Policies	929
Selecting Defragmentation Options	930
Configuring IP Defragmentation	932
Understanding Packet Decoding	934
Configuring Packet Decoding	938
Using TCP Stream Preprocessing	940
Understanding State-Related TCP Exploits	941
Initiating Active Responses with Drop Rules	941
Selecting TCP Global Options	943
Understanding Target-Based TCP Policies	943
Selecting TCP Policy Options	945
Reassembling TCP Streams	949
Configuring TCP Stream Preprocessing	952
Using UDP Stream Preprocessing	956
Configuring UDP Stream Preprocessing	957

Chapter 25:	Detecting Specific Threats	959
	Detecting Back Orifice	959
	Detecting Portscans.....	961
	Configuring Portscan Detection	965
	Understanding Portscan Events.....	968
	Preventing Rate-Based Attacks	971
	Understanding Rate-Based Attack Prevention	971
	Rate-Based Attack Prevention and Other Filters	974
	Configuring Rate-Based Attack Prevention	981
	Detecting Sensitive Data	983
	Deploying Sensitive Data Detection	985
	Selecting Global Sensitive Data Detection Options	985
	Selecting Individual Data Type Options.....	987
	Using Predefined Data Types.....	988
	Configuring Sensitive Data Detection	989
	Selecting Application Protocols to Monitor.....	993
	Special Case: Detecting Sensitive Data in FTP Traffic	995
	Using Custom Data Types	996
Chapter 26:	Using Adaptive Profiles	1003
	Understanding Adaptive Profiles	1004
	Using Adaptive Profiles with Preprocessors.....	1004
	Adaptive Profiles and FireSIGHT Recommended Rules.....	1005
	Configuring Adaptive Profiles.....	1006
Chapter 27:	Using Global Rule Thresholding	1009
	Understanding Thresholding.....	1009
	Understanding Thresholding Options	1010
	Configuring Global Thresholds	1012
	Disabling the Global Threshold	1014
Chapter 28:	Using Performance Settings in an Intrusion Policy	1015
	Event Queue Configuration	1016
	Understanding Packet Latency Thresholding.....	1017
	Setting Packet Latency Thresholding Options	1019
	Configuring Packet Latency Thresholding.....	1020
	Understanding Rule Latency Thresholding	1021
	Setting Rule Latency Thresholding Options.....	1024
	Configuring Rule Latency Thresholding	1025

	Performance Statistics Configuration	1026
	Constraining Regular Expressions	1028
	Rule Processing Configuration.....	1030
Chapter 29:	Configuring External Responses to Intrusion Events.....	1033
	Using SNMP Responses	1034
	Configuring SNMP Responses	1037
	Using Syslog Responses	1038
	Configuring Syslog Responses	1040
	Understanding Email Alerting	1042
	Configuring Email Alerting	1044
Chapter 30:	Understanding and Writing Intrusion Rules.....	1046
	Understanding Rule Anatomy	1047
	Understanding Rule Headers.....	1049
	Specifying Rule Actions	1050
	Specifying Protocols	1051
	Specifying Source and Destination IP Addresses.....	1051
	Specifying Source and Destination Ports.....	1053
	Specifying Direction.....	1055

Table of Contents

Understanding Keywords and Arguments in Rules	1055
Defining Intrusion Event Details	1057
Searching for Content Matches	1063
Constraining Content Matches	1065
Replacing Content in Inline Deployments.....	1078
Using Byte_Jump and Byte_Test.....	1079
Searching for Content Using PCRE.....	1086
Adding Metadata to a Rule	1096
Inspecting IP Header Values	1101
Inspecting ICMP Header Values	1104
Inspecting TCP Header Values and Stream Size.....	1106
Enabling and Disabling TCP Stream Reassembly	1112
Extracting SSL Information from a Session	1113
Inspecting Application Layer Protocol Values	1116
Inspecting Packet Characteristics	1153
Reading Packet Data into Keyword Arguments	1156
Initiating Active Responses with Rule Keywords.....	1160
Filtering Events	1165
Evaluating Post-Attack Traffic	1166
Detecting Attacks That Span Multiple Packets	1168
Generating Events on the HTTP Encoding Type and Location	1175
Pointing to a Specific Payload Type.....	1177
Pointing to the Beginning of the Packet Payload.....	1178
Decoding and Inspecting Base64 Data.....	1179
Constructing a Rule	1181
Writing New Rules.....	1182
Modifying Existing Rules	1184
Adding Comments to Rules.....	1187
Deleting Custom Rules.....	1188
Searching for Rules.....	1189
Filtering Rules on the Rule Editor Page	1192
Using Keywords in a Rule Filter	1192
Using Character Strings in a Rule Filter	1194
Combining Keywords and Character Strings in a Rule Filter.....	1195
Filtering Rules	1195
Chapter 31: Working with Malware Protection and File Control.....	1197
Understanding Malware Protection and File Control	1199
Configuring Malware Protection and File Control	1201
Logging Events Based on Malware Protection and File Control	1202
Integrating FireAMP with the Sourcefire 3D System	1202
Network-Based AMP vs Endpoint-Based FireAMP	1204

Table of Contents

Understanding and Creating File Policies	1206
Creating a File Policy	1213
Working with File Rules	1214
Comparing Two File Policies	1217
Working with Sourcefire Cloud Connections for FireAMP.....	1219
Creating a Sourcefire Cloud Connection	1220
Deleting or Disabling a Sourcefire Cloud Connection	1221
Working with File Events	1222
Viewing File Events.....	1223
Understanding the File Events Table	1224
Searching for File Events	1227
Working with Malware Events.....	1229
Viewing Malware Events	1232
Understanding the Malware Events Table	1233
Searching for Malware Events	1239
Working with Network File Trajectory.....	1241
Reviewing Network File Trajectory	1242
Analyzing Network File Trajectory.....	1244
Chapter 32: Introduction to Network Discovery	1250
Understanding Discovery Data Collection	1251
Understanding Host Data Collection	1252
Understanding User Data Collection	1252
Understanding Application Detection	1263
Importing Third-Party Discovery Data	1269
Uses for Discovery Data	1270
Understanding NetFlow.....	1271
Differences Between NetFlow and FireSIGHT Data	1272
Preparing to Analyze NetFlow Data	1274
Creating a Network Discovery Policy	1275
Working with Discovery Rules	1276
Restricting User Logging	1285
Configuring Advanced Network Discovery Options.....	1287
Applying the Network Discovery Policy	1297
Obtaining User Data from LDAP Servers	1298
Creating LDAP Connections with the Defense Center	1298
Enabling and Disabling User Awareness LDAP Connections	1306
Performing an On-Demand User Data Retrieval for Access Control..	1307
Configuring Defense Center-User Agent Connections	1307

Chapter 33:	Using the Network Map.....	1314
	Understanding the Network Map	1315
	Working with the Hosts Network Map	1316
	Working with the Network Devices Network Map.....	1318
	Working with the Mobile Devices Network Map.....	1319
	Working with the Applications Network Map	1321
	Working with the Vulnerabilities Network Map	1323
	Working with the Host Attributes Network Map	1325
	Working with Custom Network Topologies	1326
	Creating Custom Topologies.....	1327
	Managing Custom Topologies	1332
Chapter 34:	Using Host Profiles	1334
	Viewing Host Profiles.....	1337
	Working with Basic Host Information in the Host Profile	1338
	Working with IP Addresses in the Host Profile.....	1341
	Working with Operating Systems in the Host Profile	1341
	Viewing Operating System Identities.....	1343
	Editing an Operating System	1344
	Resolving Operating System Identity Conflicts	1345
	Working with Servers in the Host Profile.....	1347
	Server Detail	1349
	Editing Server Identities.....	1352
	Resolving Server Identity Conflicts.....	1353
	Working with Applications in the Host Profile	1354
	Viewing Applications in the Host Profile	1354
	Deleting Applications from the Host Profile	1356
	Working with VLAN Tags in the Host Profile	1357
	Working with User History in the Host Profile.....	1357
	Working with Host Attributes in the Host Profile.....	1358
	Assigning Host Attribute Values	1358
	Working with Host Protocols in the Host Profile	1359
	Working with White List Violations in the Host Profile	1360
	Creating a White List Host Profile from a Host Profile	1361
	Working with Malware Detections in the Host Profile	1362

Table of Contents

Working with Vulnerabilities in the Host Profile	1363
Viewing Vulnerability Details	1365
Setting the Vulnerability Impact Qualification	1367
Downloading Patches for Vulnerabilities	1368
Setting Vulnerabilities for Individual Hosts	1368
Working with the Predefined Host Attributes	1369
Working with User-Defined Host Attributes	1370
Creating User-Defined Host Attributes	1372
Editing a User-Defined Host Attribute	1374
Deleting a User-Defined Host Attribute	1375
Working with Scan Results in a Host Profile	1375
Scanning a Host from the Host Profile	1376

Chapter 35: Working with Discovery Events..... 1377

Viewing Discovery Event Statistics.....	1378
Statistics Summary.....	1379
Event Breakdown.....	1381
Protocol Breakdown.....	1382
Application Protocol Breakdown	1382
OS Breakdown.....	1383
Viewing Discovery Performance Graphs.....	1384
Understanding Discovery Event Workflows	1386
Working with Discovery and Host Input Events	1388
Understanding Discovery Event Types	1389
Understanding Host Input Event Types	1394
Viewing Discovery and Host Input Events.....	1396
Understanding the Discovery Events Table	1397
Searching for Discovery Events	1399
Working with Hosts	1401
Viewing Hosts.....	1402
Understanding the Hosts Table	1403
Creating a Traffic Profile for Selected Hosts	1407
Creating a Compliance White List Based on Selected Hosts	1408
Searching for Hosts	1408
Working with Host Attributes	1412
Viewing Host Attributes	1412
Understanding the Host Attributes Table.....	1413
Setting Host Attributes for Selected Hosts.....	1415
Searching for Host Attributes.....	1416
Working with Servers	1418
Viewing Servers.....	1418
Understanding the Servers Table.....	1419
Searching for Servers.....	1422

Table of Contents

Working with Applications	1424
Viewing Applications	1425
Understanding the Applications Table	1425
Searching for Applications	1427
Working with Application Details	1429
Viewing Application Details	1429
Understanding the Application Detail Table	1430
Searching for Application Details	1432
Working with Sourcefire Vulnerabilities	1434
Viewing Sourcefire Vulnerabilities	1435
Understanding the Sourcefire Vulnerabilities Table	1436
Deactivating Sourcefire Vulnerabilities	1438
Searching for Sourcefire Vulnerabilities	1439
Working with Third-Party Vulnerabilities	1441
Viewing Third-Party Vulnerabilities	1441
Understanding the Third-Party Vulnerabilities Table	1442
Searching for Third-Party Vulnerabilities	1444
Working with Users	1445
Viewing Users	1447
Understanding the Users Table	1447
Understanding User Details and Host History	1449
Searching for Users	1451
Working with User Activity	1453
Viewing User Activity Events	1454
Understanding the User Activity Table	1455
Searching for User Activity	1456
Chapter 36: Configuring Correlation Policies and Rules.....	1459
Creating Rules for Correlation Policies	1461
Providing Basic Rule Information	1464
Specifying Correlation Rule Trigger Criteria	1464
Adding a Host Profile Qualification	1481
Adding a Connection Tracker	1486
Adding a User Qualification	1496
Adding Snooze and Inactive Periods	1498
Understanding Rule Building Mechanics	1499
Managing Rules for Correlation Policies	1508
Modifying a Rule	1508
Deleting a Rule	1509
Creating a Rule Group	1509

Table of Contents

Grouping Correlation Responses	1510
Creating a Response Group	1511
Modifying a Response Group	1512
Deleting a Response Group	1512
Activating and Deactivating Response Groups	1512
Creating Correlation Policies	1513
Providing Basic Policy Information	1515
Adding Rules and White Lists to a Correlation Policy	1515
Setting Rule and White List Priorities	1516
Adding Responses to Rules and White Lists	1517
Managing Correlation Policies	1519
Activating and Deactivating Correlation Policies	1520
Editing a Correlation Policy	1520
Deleting a Correlation Policy	1520
Working with Correlation Events	1521
Viewing Correlation Events	1521
Understanding the Correlation Events Table	1524
Searching for Correlation Events	1526

Chapter 37: Using the Sourcefire 3D System as a Compliance Tool .. 1530

Understanding Compliance White Lists	1532
Understanding White List Targets	1533
Understanding White List Host Profiles	1534
Understanding White List Evaluations	1538
Understanding White List Violations	1539
Creating Compliance White Lists	1541
Surveying Your Network	1543
Providing Basic White List Information	1545
Configuring Compliance White List Targets	1545
Configuring Compliance White List Host Profiles	1548
Managing Compliance White Lists	1563
Modifying a Compliance White List	1563
Deleting a Compliance White List	1564
Working with Shared Host Profiles	1564
Creating Shared Host Profiles	1565
Modifying a Shared Host Profile	1567
Deleting a Shared Host Profile	1571
Resetting Built-In Host Profiles to Their Factory Defaults	1571
Working with White List Events	1572
Viewing White List Events	1573
Understanding the White List Events Table	1575
Searching for Compliance White List Events	1576

	Working with White List Violations	1579
	Viewing White List Violations	1579
	Understanding the White List Violations Table	1581
	Searching for White List Violations	1582
Chapter 38:	Creating Traffic Profiles	1585
	Providing Basic Profile Information	1587
	Specifying Traffic Profile Conditions	1588
	Syntax for Traffic Profile Conditions	1589
	Adding a Host Profile Qualification	1590
	Syntax for Host Profile Qualifications	1591
	Setting Profile Options	1593
	Saving a Traffic Profile	1595
	Activating and Deactivating Traffic Profiles	1595
	Editing a Traffic Profile	1596
	Understanding Condition-Building Mechanics	1596
	Building a Single Condition	1597
	Adding and Linking Conditions	1599
	Using Multiple Values in a Condition	1602
	Viewing Traffic Profiles	1603
Chapter 39:	Configuring Remediations	1606
	Creating Remediations	1607
	Configuring Remediations for Cisco IOS Routers	1608
	Configuring Remediations for Cisco PIX Firewalls	1617
	Configuring Nmap Remediations	1623
	Configuring Set Attribute Remediations	1629
	Working with Remediation Status Events	1633
	Viewing Remediation Status Events	1633
	Working with Remediation Status Events	1636
	Understanding the Remediation Status Table	1636
	Searching for Remediation Status Events	1638
Chapter 40:	Enhancing Network Discovery	1641
	Assessing Your Detection Strategy	1642
	Are Your Managed Devices Correctly Placed?	1642
	Do Unidentified Operating Systems Have a Unique TCP Stack?	1643
	Can the Sourcefire 3D System Identify All Applications?	1644
	Have You Applied Patches that Fix Vulnerabilities?	1644
	Do You Want to Track Third-Party Vulnerabilities?	1644

Table of Contents

Enhancing Your Network Map	1644
Understanding Passive Detection.....	1645
Understanding Active Detection.....	1645
Understanding Current Identities.....	1646
Understanding Identity Conflicts	1648
Using Custom Fingerprinting	1649
Fingerprinting Clients.....	1650
Fingerprinting Servers.....	1655
Managing Fingerprints	1660
Activating Fingerprints	1660
Deactivating Fingerprints	1661
Deleting Fingerprints	1661
Editing Fingerprints.....	1662
Working with Application Detectors	1663
Creating a User-Defined Application Protocol Detector	1666
Managing Detectors	1673
Importing Host Input Data	1680
Enabling the Use of Third-Party Data.....	1681
Managing Third-Party Product Mappings.....	1682
Mapping Third-Party Vulnerabilities.....	1687
Managing Custom Product Mappings	1688
Chapter 41: Configuring Active Scanning	1692
Understanding Nmap Scans	1693
Understanding Nmap Remediations.....	1693
Creating an Nmap Scanning Strategy	1697
Sample Nmap Scanning Profiles.....	1699
Setting up Nmap Scans	1702
Creating an Nmap Scan Instance.....	1702
Creating an Nmap Scan Target	1704
Creating an Nmap Remediation	1705
Managing Nmap Scanning.....	1710
Managing Nmap Scan Instances	1710
Managing Nmap Remediations	1712
Running an On-Demand Nmap Scan	1713
Managing Scan Targets	1714
Editing a Scan Target	1715
Deleting a Scan Target.....	1716

Table of Contents

Working with Active Scan Results	1716
Viewing Scan Results	1716
Understanding the Scan Results Table	1718
Analyzing Scan Results	1719
Monitoring Scans	1719
Importing Scan Results	1720
Searching for Scan Results	1721
Chapter 42: Working with Reports.....	1724
Generating Reports.....	1725
Creating a Report Template from an Event View	1725
Creating a Report Template by Importing a Dashboard or Workflow	1727
Generating Reports from a Report Template	1729
Using Report Generation Options.....	1732
Managing Reports	1733
Understanding Report Templates	1733
Using Report Templates	1736
Creating Report Templates from Existing Templates.....	1737
Creating New Report Templates.....	1740
Editing the Sections of a Report Template	1743
Working with Searches in Report Template Sections	1749
Using Input Parameters	1750
Editing Document Attributes in a Report Template	1756
Customizing a Cover Page	1759
Managing Logos	1760
Using Report Generation Options.....	1763
Generating Reports Using the Scheduler	1763
Distributing Reports by Email at Generation Time.....	1763
Using Remote Storage for Reports.....	1765
Managing Report Templates and Report Files.....	1766
Exporting and Importing Report Templates	1766
Deleting Report Templates	1768
Downloading Reports	1768
Deleting Reports.....	1769
Chapter 43: Searching for Events.....	1770
Performing and Saving Searches	1770
Performing a Search.....	1771
Loading a Saved Search	1774
Deleting a Saved Search	1774
Using Wildcards and Symbols in Searches.....	1775
Using Objects and Application Filters in Searches.....	1775

	Specifying Time Constraints in Searches.....	1775
	Specifying IP Addresses in Searches.....	1776
	Specifying Ports in Searches.....	1777
Chapter 44:	Using Custom Tables	1779
	Understanding Custom Tables.....	1780
	Understanding Possible Table Combinations.....	1780
	Creating a Custom Table.....	1783
	Modifying a Custom Table	1786
	Deleting a Custom Table.....	1786
	Viewing a Workflow Based on a Custom Table	1787
	Searching Custom Tables	1788
 Chapter 45:	 Understanding and Using Workflows.....	 1791
	Components of a Workflow.....	1792
	Comparing Predefined and Custom Workflows	1794
	Comparing Workflows for Predefined and Custom Tables	1794
	Predefined Intrusion Event Workflows	1794
	Predefined Malware Workflows	1797
	Predefined File Workflows.....	1798
	Predefined Connection Data Workflows.....	1798
	Predefined Host Workflows.....	1800
	Predefined Applications Workflows	1801
	Predefined Application Details Workflows.....	1802
	Predefined Servers Workflows	1803
	Predefined Host Attributes Workflows	1803
	Predefined Discovery Events Workflows.....	1804
	Predefined User Workflows.....	1804
	Predefined Vulnerabilities Workflows	1804
	Predefined Third-Party Vulnerabilities Workflows	1805
	Predefined Correlation and White List Workflows.....	1805
	Predefined System Workflows	1806
	Saved Custom Workflows	1807

Table of Contents

Using Workflows	1808
Selecting Workflows	1809
Understanding the Workflow Toolbar	1812
Using Workflow Pages.....	1812
Setting Event Time Constraints	1820
Constraining Events	1830
Using Compound Constraints.....	1832
Sorting Table View Pages and Changing Their Layout	1834
Sorting Drill-Down Workflow Pages.....	1835
Selecting Rows on a Workflow Page	1835
Navigating to Other Pages in the Workflow.....	1836
Navigating Between Workflows	1836
Using Bookmarks.....	1838
Using Custom Workflows.....	1840
Creating Custom Workflows.....	1841
Creating Custom Connection Data Workflows	1843
Viewing Custom Workflows	1846
Editing Custom Workflows	1847
Deleting Custom Workflows.....	1847
Chapter 46: Managing Users.....	1849
Understanding Sourcefire User Authentication	1849
Understanding Internal Authentication	1851
Understanding External Authentication	1851
Understanding User Privileges	1852
Managing Authentication Objects	1854
Understanding LDAP Authentication	1854
Preparing to Create an LDAP Authentication Object	1859
Quick Start to LDAP Authentication.....	1860
Creating Advanced LDAP Authentication Objects	1866
LDAP Authentication Object Examples	1880
Editing LDAP Authentication Objects	1884
Understanding RADIUS Authentication	1885
Creating RADIUS Authentication Objects.....	1885
RADIUS Authentication Object Examples	1894
Editing RADIUS Authentication Objects	1896
Deleting Authentication Objects.....	1897

Table of Contents

Managing User Accounts	1897
Viewing User Accounts.....	1898
Adding New User Accounts.....	1899
Managing Command Line Access	1901
Managing Externally Authenticated User Accounts.....	1902
Managing User Login Settings.....	1903
Configuring User Roles	1905
Managing Custom User Roles	1908
Modifying User Privileges and Options	1912
Understanding Restricted User Access Properties.....	1912
Modifying User Passwords.....	1913
Deleting User Accounts.....	1914
User Account Privileges.....	1914
Managing User Role Escalation	1925
Configuring the Escalation Target Role	1926
Configuring a Custom User Role for Escalation.....	1926
Escalating Your User Role	1928
Chapter 47: Scheduling Tasks	1929
Configuring a Recurring Task	1930
Automating Backup Jobs	1932
Automating Certificate Revocation List Downloads.....	1934
Automating Nmap Scans	1936
Preparing Your System for an Nmap Scan	1936
Scheduling an Nmap Scan	1936
Automating Applying an Intrusion Policy.....	1938
Automating Reports.....	1940
Automating Geolocation Database Updates.....	1942
Automating FireSIGHT Recommendations.....	1943
Automating Software Updates	1945
Automating Software Downloads.....	1946
Automating Software Pushes	1948
Automating Software Installs.....	1949
Automating Vulnerability Database Updates	1951
Automating VDB Update Downloads.....	1952
Automating VDB Update Installs	1953
Automating URL Filtering Updates	1955
Viewing Tasks	1957
Using the Calendar	1957
Using the Task List.....	1958

	Editing Scheduled Tasks	1959
	Deleting Scheduled Tasks	1959
	Deleting a Recurring Task	1960
	Deleting a One-Time Task	1960
Chapter 48:	Managing System Policies.....	1961
	Creating a System Policy	1962
	Editing a System Policy.....	1963
	Applying a System Policy	1965
	Comparing System Policies	1965
	Deleting System Policies	1969
	Configuring a System Policy	1969
	Configuring Access Control Policy Preferences	1970
	Configuring the Access List for Your Appliance	1971
	Configuring Audit Log Settings	1973
	Configuring Authentication Profiles	1975
	Configuring Dashboard Settings	1977
	Configuring Database Event Limits	1978
	Configuring DNS Cache Properties	1981
	Configuring a Mail Relay Host and Notification Address	1983
	Configuring Intrusion Policy Preferences	1985
	Specifying a Different Language	1986
	Adding a Custom Login Banner	1987
	Configuring SNMP Polling.....	1988
	Synchronizing Time	1991
	Configuring User Interface Settings.....	1994
	Mapping Vulnerabilities for Servers	1996
Chapter 49:	Configuring Appliance Settings.....	1998
	Viewing and Modifying the Appliance Information	1999
	Using Custom HTTPS Certificates	2002
	Viewing the Current HTTPS Server Certificate	2002
	Generating a Server Certificate Request	2003
	Uploading Server Certificates	2004
	Configuring User Certificates.....	2006
	Enabling Access to the Database	2007
	Configuring Network Settings.....	2009
	Editing Management Interface Configurations	2012
	Shutting Down and Restarting the System.....	2014

Table of Contents

Setting the Time Manually	2015
Managing Remote Storage	2017
Using Local Storage	2018
Using NFS for Remote Storage	2019
Using SSH for Remote Storage	2020
Using SMB for Remote Storage	2022
Understanding Change Reconciliation	2024
Managing Remote Console Access	2025
Configuring Remote Console Settings on the Appliance	2026
Enabling Lights-Out Management User Access	2028
Using a Serial Over LAN Connection	2029
Using Lights-Out Management	2030
Enabling Sourcefire Cloud Communications	2032
Chapter 50: Licensing the Sourcefire 3D System	2036
Understanding Licensing	2036
License Types and Restrictions	2037
Licensing High Availability Pairs	2045
Licensing Stacked and Clustered Devices	2045
Licensing Series 2 Appliances	2045
Understanding FireSIGHT Host and User License Limits	2045
Viewing Your Licenses	2048
Adding a License to the Defense Center	2050
Deleting a License	2052
Changing a Device's Licensed Capabilities	2052
Chapter 51: Updating System Software	2054
Understanding Update Types	2055
Performing Software Updates	2056
Planning for the Update	2056
Understanding the Update Process	2058
Updating a Defense Center	2062
Updating Managed Devices	2064
Monitoring the Status of Major Updates	2066
Uninstalling Software Updates	2068
Updating the Vulnerability Database	2070
Importing Rule Updates and Local Rule Files	2072
Using One-Time Rule Updates	2074
Using Recurring Rule Updates	2077
Importing Local Rule Files	2080
Viewing the Rule Update Log	2082
Updating the Geolocation Database	2092

Chapter 52:	Monitoring the System.....	2095
	Viewing Host Statistics.....	2096
	Monitoring System Status and Disk Space Usage	2099
	Viewing System Process Status	2100
	Understanding Running Processes.....	2103
	Understanding System Daemons.....	2103
	Understanding Executables and System Utilities	2105
Chapter 53:	Using Health Monitoring.....	2109
	Understanding Health Monitoring	2110
	Understanding Health Policies	2111
	Understanding Health Modules	2112
	Understanding Health Monitoring Configuration	2115
	Configuring Health Policies.....	2116
	Understanding the Default Health Policy	2117
	Creating Health Policies	2118
	Applying Health Policies.....	2145
	Editing Health Policies	2147
	Comparing Health Policies	2150
	Deleting Health Policies	2154
	Using the Health Monitor Blacklist	2155
	Blacklisting Health Policies or Appliances	2156
	Blacklisting a Health Policy Module	2158
	Configuring Health Monitor Alerts.....	2159
	Creating Health Monitor Alerts	2159
	Interpreting Health Monitor Alerts.....	2161
	Editing Health Monitor Alerts	2162
	Deleting Health Monitor Alerts	2163
	Using the Health Monitor	2163
	Interpreting Health Monitor Status	2165
	Using Appliance Health Monitors	2166
	Viewing Alerts by Status.....	2167
	Running All Modules for an Appliance	2167
	Running a Specific Health Module.....	2168
	Generating Health Module Alert Graphs.....	2170
	Using the Health Monitor to Troubleshoot.....	2171
	Working with Health Events	2174
	Understanding Health Event Views	2174
	Viewing Health Events.....	2175
	Understanding the Health Events Table.....	2183
	Searching for Health Events.....	2184

Chapter 54:	Auditing the System	2187
	Managing Audit Records	2187
	Viewing Audit Records.....	2188
	Suppressing Audit Records.....	2191
	Understanding the Audit Log Table.....	2196
	Using the Audit Log to Examine Changes	2196
	Searching Audit Records.....	2197
	Viewing the System Log.....	2200
	Filtering System Log Messages	2201
Chapter 55:	Using Backup and Restore.....	2204
	Creating Backup Files	2205
	Creating Backup Profiles.....	2208
	Backing up Your Managed Devices with a Defense Center.....	2209
	Uploading Backups from a Local Host	2210
	Restoring the Appliance from a Backup File	2211
Chapter 56:	Specifying User Preferences.....	2215
	Changing Your Password	2216
	Changing an Expired Password	2216
	Specifying Your Home Page.....	2217
	Configuring Event View Settings	2218
	Event Preferences	2218
	Default Time Windows	2219
	Default Workflows	2221
	Setting Your Default Time Zone	2222
	Specifying Your Default Dashboard.....	2223
Appendix A:	Importing and Exporting Objects.....	2225
	Exporting Objects	2226
	Importing Objects	2231
Appendix B:	Purging Discovery Data from the Database.....	2236
Appendix C:	Viewing the Status of Long-Running Tasks.....	2238
	Viewing the Task Queue	2238
	Managing the Task Queue	2240

Appendix D: **Command Line Reference** 2242

- Basic CLI Commands 2243
 - configure password 2244
 - end 2244
 - exit 2244
 - help 2245
 - history 2245
 - logout 2245
 - ? (question mark) 2246
 - ?? (double question marks) 2246

Table of Contents

Show Commands	2246
access-control-config	2248
alarms	2248
arp-tables	2249
audit-log	2249
clustering	2249
cpu	2250
database	2251
device-settings	2252
disk	2252
dns	2252
expert.....	2253
bypass.....	2253
fan-status	2253
fastpath-rules	2254
gui	2254
hostname.....	2254
hyperthreading.....	2255
inline-sets.....	2255
interfaces	2255
lcd	2256
link-state	2256
log-ips-connection.....	2256
managers	2257
memory	2257
model.....	2257
mpls-depth.....	2258
NAT	2258
network.....	2260
network-modules.....	2260
ntp.....	2260
perfstats.....	2261
portstats.....	2261
power-supply-status	2261
process-tree	2262
processes	2262
routing-table.....	2262
serial-number	2263
stacking.....	2263
summary.....	2263
time.....	2264
traffic-statistics	2264
user.....	2264
users	2265
version	2266
virtual-routers.....	2266
virtual-switches	2267
VPN.....	2267

Table of Contents

Configuration Commands	2268
clustering	2269
bypass	2269
gui	2269
lcd	2270
log-ips-connections	2270
manager	2270
mpls-depth	2271
network	2271
password	2275
stacking disable	2275
user	2275
System Commands	2279
access-control	2279
disable-http-user-cert	2280
file	2280
generate-troubleshoot	2282
ldapsearch	2282
lockdown-sensor	2283
nat rollback	2283
reboot	2283
restart	2284
shutdown	2284
Appendix E: Third-Party Products	2285
Appendix F: End User License Agreement	2286
Glossary	2299
Index	2330

CHAPTER 1

INTRODUCTION TO THE SOURCEFIRE 3D SYSTEM

The Sourcefire 3D® System combines the security of an industry-leading network intrusion protection system with the power to control access to your network based on detected applications, users, and URLs. You can also use Sourcefire appliances to serve in a switched, routed, or hybrid (switched and routed) environment; to perform network address translation (NAT); and to build secure virtual private network (VPN) tunnels among the virtual routers on Sourcefire managed devices, or from managed devices to remote devices or other third-party VPN endpoints.

The Sourcefire Defense Center® provides a centralized management console and database repository for the Sourcefire 3D System. Managed devices installed on network segments monitor traffic for analysis.

Devices in a passive deployment monitor traffic flowing across a network, for example, using a switch SPAN, virtual switch, or mirror port. Passive sensing interfaces receive all traffic unconditionally and no traffic received on these interfaces is retransmitted.

Devices in an inline deployment allow you to protect your network from attacks that might affect the availability, integrity, or confidentiality of hosts on the network. Inline interfaces receive all traffic unconditionally, and traffic received on these interfaces is retransmitted unless explicitly dropped by some configuration in your deployment. Inline devices can be deployed as a simple intrusion prevention system. You can also configure inline devices to perform access control as well as manage network traffic in other ways.

In addition to physical appliances, Sourcefire packages 64-bit virtual Defense Centers® and virtual devices for the VMware ESX/ESXi hosting environment.

This guide provides information about the features and functionality of the Sourcefire 3D System. The explanatory text, graphics, and procedures in each chapter provide detailed information to help you navigate the user interface, maximize the performance of your system, and troubleshooting complications.

The topics that follow introduce you to the Sourcefire 3D System, describe its key components, explain how to log in and out of Sourcefire appliances, contain some basic information about using the system's web interface, and help you understand how to use this guide:

- [Sourcefire 3D System Appliances](#) on page 38
- [Sourcefire 3D System Components](#) on page 44
- [Security, Internet Access, and Communication Ports](#) on page 50
- [Documentation Resources](#) on page 54
- [Documentation Conventions](#) on page 54
- [IP Address Conventions](#) on page 57
- [Logging into the Appliance](#) on page 58
- [Logging into the Appliance to Set Up an Account](#) on page 60
- [Logging Out of the Appliance](#) on page 62
- [Using the Context Menu](#) on page 63

Sourcefire 3D System Appliances

A Sourcefire *appliance* is either a traffic-sensing managed *device* or a managing *Defense Center*.

Physical devices are fault-tolerant, purpose-built network appliances available in a range of throughputs and with a range of capabilities. Defense Centers serve as central management points for these devices, and automatically aggregate and correlate the events they generate. There are several *models* of each physical appliance type; these models are further grouped into *series* and *family*.

You can also host 64-bit virtual Defense Centers and devices on VMware ESX/ESXi. Virtual Defense Centers can manage physical devices, and physical Defense Centers can manage virtual devices.

Many Sourcefire 3D System capabilities are appliance dependent. For more information, see the following sections:

- [Defense Centers](#) on page 39
- [Managed Devices](#) on page 39
- [Understanding Appliance Series, Models, and Capabilities](#) on page 40

Defense Centers

The Defense Center provides a centralized management point and event database for your Sourcefire 3D System deployment. Defense Centers, which can be physical or virtual, aggregate and correlate intrusion, file, malware, discovery, connection, and performance data. This allows you to monitor the information that your devices report in relation to one another, and to assess and control the overall activity that occurs on your network.

Key features of the Defense Center include:

- device, license, and policy management
- event and contextual information displayed in tables, graphs, and charts
- health and performance monitoring
- external notification and alerting
- correlation and remediation features for real-time threat response
- custom and template-based reporting

For many physical Defense Centers, a high availability (redundancy) feature can help you ensure continuity of operations.

Managed Devices

Physical Sourcefire devices are fault-tolerant, purpose-built network appliances available in a range of throughputs. You can also host virtual devices. Devices deployed passively help you gain insight into your network traffic. Deployed inline, you can use Sourcefire devices to affect the flow of traffic based on multiple criteria. You must manage Sourcefire devices with a Defense Center.

Depending on model and license, managed devices:

- gather detailed information about your organization's hosts, operating systems, applications, users, files, networks, and vulnerabilities
- block or allow network traffic based on various network-based criteria, as well as other criteria including applications, users, URLs, IP address reputations, and the results of intrusion or malware inspections
- have switching, routing, DHCP, NAT, and VPN capabilities, as well as configurable bypass interfaces, fast-path rules, and strict TCP enforcement
- have clustering (redundancy) to help you ensure continuity of operations, and stacking to combine resources from multiple devices

Understanding Appliance Series, Models, and Capabilities

Version 5.2 of the Sourcefire 3D System is available on two series of physical appliances, as well as virtual appliances. Many Sourcefire 3D System capabilities are appliance dependent. For more information, see:

- [Series 2 Appliances](#) on page 40
- [Series 3 Appliances](#) on page 41
- [Virtual Appliances](#) on page 41
- [Appliances Delivered with Version 5.2](#) on page 41
- [Supported Capabilities by Appliance Model](#) on page 42

Series 2 Appliances

Series 2 is the second series of Sourcefire physical appliances. Because of resource and architecture limitations, Series 2 devices support a restricted set of Sourcefire 3D System features.

Although Sourcefire does not deliver Version 5.2 on Series 2 appliances other than 3D500/1000/2000 devices, you can reimage the following Series 2 devices and Defense Centers to Version 5.2:

- 3D2100/2500/3500/4500
- 3D6500
- 3D9900
- DC500/1000/3000

WARNING! There is no update path from Version 4.10.x to Version 5.2. Reimaging results in the loss of **all** configuration and event data on the appliance. You **cannot** import this data onto an appliance after a reimage. For more information, see the *Sourcefire 3D System Installation Guide*.

When running Version 5.2, Series 2 devices automatically have most of the capabilities associated with a Protection license: intrusion detection and prevention, file control, and basic access control. However, Series 2 devices cannot perform Security Intelligence filtering, advanced access control, or advanced malware protection. You also cannot enable other licensed capabilities on a Series 2 device. With the exception of the 3D9900, which supports fast-path rules, stacking, and tap mode, Series 2 devices do not support any of the hardware-based features associated with Series 3 devices: switching, routing, NAT, and so on.

When running Version 5.2, DC1000 and DC3000 Series 2 Defense Centers support all the features of the Sourcefire 3D System; the DC500 has more limited capabilities.

Series 3 Appliances

Series 3 is the third series of Sourcefire physical appliances. All 7000 Series and 8000 Series devices are Series 3 appliances. 8000 Series devices are more powerful and support a few features that 7000 Series devices do not.

Virtual Appliances

You can host 64-bit virtual Defense Centers and devices on VMware ESX/ESXi. Virtual Defense Centers can manage up to 25 physical or virtual devices; physical Defense Centers can manage virtual devices.

Regardless of the licenses installed and applied, virtual appliances do not support any of the system's hardware-based features: redundancy and resource sharing, switching, routing, and so on. Also, virtual devices do not have web interfaces.

Appliances Delivered with Version 5.2

The following table lists the appliances that Sourcefire delivers with Version 5.2 of the Sourcefire 3D System.

Version 5.2 Sourcefire Appliances

MODELS/FAMILY	SERIES	TYPE
Series 2 devices: • 3D500/1000/2000	Series 2	device
70xx Family: • 3D7010/7020/7030	Series 3 (7000 Series)	device
71xx Family: • 3D7110/7120 • 3D7115/7125	Series 3 (7000 Series)	device
81xx Family: • 3D8120/8130/8140	Series 3 (8000 Series)	device
82xx Family: • 3D8250 • 3D8260/8270/8290	Series 3 (8000 Series)	device
virtual devices	none	device
Series 3 Defense Centers: • DC750/1500/3500	Series 3	Defense Center
virtual Defense Centers	none	Defense Center

Although Sourcefire does not deliver Version 5.2 on Series 2 appliances other than 3D500/1000/2000 devices, you can reimage the following Series 2 devices and Defense Centers to Version 5.2:

- 3D2100/2500/3500/4500
- 3D6500
- 3D9900
- DC500/1000/3000

Reimaging results in the loss of **all** configuration and event data on the appliance. See the *Sourcefire 3D System Installation Guide* for more information.

Supported Capabilities by Appliance Model

Many Sourcefire 3D System capabilities are appliance dependent. The table below matches the major capabilities of the system with the appliances that support those capabilities, assuming you have the correct licenses installed and applied.

The Defense Center column for device-based capabilities (such as stacking, switching, and routing) indicates whether that Defense Center can manage and configure devices to perform their functions. For example, you can use a Series 2 DC1000 to manage NAT on Series 3 devices.

Supported Capabilities by Appliance Model

FEATURE	SERIES 2 DEVICE	SERIES 2 DEFENSE CENTER	SERIES 3 DEVICE	SERIES 3 DEFENSE CENTER	VIRTUAL DEVICE	VIRTUAL DEFENSE CENTER
network discovery: host, application, and user	X	X	X	X	X	X
geolocation data	X	DC1000 DC3000	X	X	X	X
intrusion detection and prevention (IPS)	X	X	X	X	X	X
Security Intelligence filtering		DC1000, DC3000	X	X	X	X
access control: basic network control	X	X	X	X	X	X
access control: applications		X	X	X	X	X

Supported Capabilities by Appliance Model (Continued)

FEATURE	SERIES 2 DEVICE	SERIES 2 DEFENSE CENTER	SERIES 3 DEVICE	SERIES 3 DEFENSE CENTER	VIRTUAL DEVICE	VIRTUAL DEFENSE CENTER
access control: users		DC1000, DC3000	X	X	X	X
access control: literal URLs		X	X	X	X	X
access control: URL filtering by category and reputation		DC1000, DC3000	X	X	X	X
file control: by file type	X	X	X	X	X	X
network-based advanced malware protection (AMP)		DC1000, DC3000	X	X	X	X
FireAMP integration	n/a	X	n/a	X	n/a	X
fast-path rules	3D9900	X	8000 Series	X		X
strict TCP enforcement		X	X	X		X
configurable bypass interfaces	X	X	except where hardware limited	X		X
tap mode	3D9900	X	X	X		X
switching and routing		X	X	X		X
NAT policies		X	X	X		X
VPN		X	X	X		X
high availability	n/a	DC1000, DC3000	n/a	DC1500, DC3500	n/a	
device stacking	3D9900	X	3D8140, 82xx Family	X		X

Supported Capabilities by Appliance Model (Continued)

FEATURE	SERIES 2 DEVICE	SERIES 2 DEFENSE CENTER	SERIES 3 DEVICE	SERIES 3 DEFENSE CENTER	VIRTUAL DEVICE	VIRTUAL DEFENSE CENTER
device clustering		X	X	X		X
clustered stacks		X	3D8140, 82xx Family	X		X
interactive CLI			X		X	

Sourcefire 3D System Components

The topics that follow describe some of the key capabilities of the Sourcefire 3D System that contribute to your organization's security, acceptable use policy, and traffic management strategy:

- [Redundancy and Resource Sharing](#) on page 44
- [Network Traffic Management](#) on page 45
- [FireSIGHT](#) on page 46
- [Access Control](#) on page 46
- [Intrusion Detection and Prevention](#) on page 47
- [File Tracking, Control, and Malware Protection](#) on page 48
- [Application Programming Interfaces](#) on page 49

TIP! Many Sourcefire 3D System features are appliance model, license, and user role dependent. This documentation includes information about which Sourcefire 3D System licenses and devices are required for each feature, and which user roles have permission to complete each procedure. For more information, see [Documentation Conventions](#) on page 54.

Redundancy and Resource Sharing

The redundancy and resource-sharing features of the Sourcefire 3D System allow you to ensure continuity of operations and to combine the processing resources of multiple physical devices.

Defense Center High Availability

To ensure continuity of operations, a Defense Center *high availability* feature allows you to designate redundant DC1000, DC1500, DC3000, or DC3500 Defense Centers to manage devices. Event data streams from managed devices

to both Defense Centers; certain configuration elements are maintained on both Defense Centers. If one Defense Center fails, you can monitor your network without interruption using the other Defense Center.

Device Stacking

Device stacking allows you to increase the amount of traffic inspected on a network segment by connecting two to four physical devices in a stacked configuration. When you establish a stacked configuration, you combine the resources of each stacked device into a single, shared configuration.

Device Clustering

Device clustering (sometimes called device high availability) allows you to establish redundancy of networking functionality and configuration data between two or more Series 3 devices or stacks. Clustering two or more peer devices or stacks results in a single logical system for policy applies, system updates, and registration. With device clustering, the system can fail over either manually or automatically.

In most cases, you can achieve Layer 3 redundancy without clustering devices by using the Sourcefire Redundancy Protocol (SFRP). SFRP allows devices to act as redundant gateways for specified IP addresses. With network redundancy, you can configure two or more devices or stacks to provide identical network connections, ensuring connectivity for other hosts on the network.

Network Traffic Management

The Sourcefire 3D System's network traffic management features allow managed devices to act as part of your organization's network infrastructure. You can configure Series 3 devices to serve in a switched, routed, or hybrid (switched and routed) environment; to perform network address translation (NAT); and to build secure virtual private network (VPN) tunnels.

Switching

You can configure the Sourcefire 3D System in a Layer 2 deployment so that it provides packet switching between two or more network segments. In a Layer 2 deployment, you configure switched interfaces and virtual switches on managed devices to operate as standalone broadcast domains. A virtual switch uses the MAC address from a host to determine where to send packets.

Routing

You can configure the Sourcefire 3D System in a Layer 3 deployment so that it routes traffic between two or more interfaces. In a Layer 3 deployment, you configure routed interfaces and virtual routers on managed devices to receive and forward traffic. The system routes packets by making packet forwarding decisions according to the destination IP address. Routers obtain the destination from the

outgoing interface based on the forwarding criteria, and access control rules designate the security policies to apply.

When you configure virtual routers, you can define static routes. In addition, you can configure Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) dynamic routing protocols. You can also configure a combination of static routes and RIP or static routes and OSPF. You can set up DHCP relay for each virtual router you configure.

If you use both virtual switches and virtual routers in your Sourcefire appliance configuration, you can configure associated hybrid interfaces to bridge traffic between them. These utilities analyze traffic to determine its type and the appropriate response (route, switch, or otherwise).

NAT

In a Layer 3 deployment, you can configure network address translation (NAT). You can expose an internal server to an external network, or allow an internal host or server to connect to an external application. You can also configure NAT to hide private network addresses from an external network by using a block of IP addresses, or by using a limited block of IP addresses and port translation.

VPN

A virtual private network (VPN) is a network connection that establishes a secure tunnel between endpoints via a public source, such as the Internet or other network. You can configure the Sourcefire 3D System to build secure VPN tunnels from the virtual routers of Series 3 devices to remote devices or other third-party VPN endpoints.

FireSIGHT

FireSIGHT™ is Sourcefire's discovery and awareness technology that collects information about hosts, operating systems, applications, users, files, networks, geolocation information, and vulnerabilities, in order to provide you with a complete view of your network.

You can use the Defense Center's web interface to view and analyze data collected by FireSIGHT. You can also use this data to help you perform access control and modify intrusion rule states.

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. You can use a policy that does

not include *access control rules* to handle traffic in one of the following ways, using what is called the *default action*:

- block all traffic from entering your network
- trust all traffic to enter your network without further inspection
- allow all traffic to enter your network, and inspect the traffic with a network discovery policy only
- allow all traffic to enter your network, and inspect the traffic with intrusion and network discovery policies

You can include access control rules in an access control policy to further define how traffic is handled by targeted devices, from simple IP address matching to complex scenarios involving different users, applications, ports, and URLs. For each rule, you specify a rule *action*, that is, whether to trust, monitor, block, or inspect matching traffic with an intrusion or file policy.

For each access control policy, you can create a custom HTML page that users see when the system blocks their HTTP requests. Optionally, you can display a page that warns users, but also allows them to click a button to continue to the originally requested site.

As part of access control, the Security Intelligence feature allows you to blacklist—deny traffic to and from—specific IP addresses before the traffic is subjected to analysis by access control rules.

Access control includes intrusion detection and prevention, file control, and advanced malware protection. For more information, see the next sections.

Intrusion Detection and Prevention

Intrusion detection and prevention allows you to monitor your network traffic for security violations and, in inline deployments, to block or alter malicious traffic.

Intrusion prevention is integrated into access control, where you can associate an intrusion policy with specific access control rules. If network traffic meets the conditions in a rule, you can analyze the matching traffic with an intrusion policy. You can also associate an intrusion policy with the default action of an access control policy.

An intrusion policy contains a variety of components, including:

- rules that inspect the protocol header values, payload content, and certain packet size characteristics
- rule state configuration based on FireSIGHT recommendations
- advanced settings, such as preprocessors and other detection and performance features
- preprocessor rules that allow you to generate events for associated preprocessors and preprocessor options

File Tracking, Control, and Malware Protection

To help you identify and mitigate the effects of malware, the Sourcefire 3D System's file control, network file trajectory, and advanced malware protection components can detect, track, and optionally block the transmission of files (including malware files) in network traffic.

File Control

File control allows managed 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, including PDFs, many Microsoft Office documents, and other types. When a managed device detects one of these file types, the Defense Center obtains the file's disposition from the *Sourcefire cloud*. The managed device uses this information to block or allow the file.

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.

FireAMP Integration

FireAMP is Sourcefire's enterprise-class, advanced malware analysis and protection solution that discovers, understands, and blocks advanced malware outbreaks, advanced persistent threats, and targeted attacks.

If your organization has a FireAMP subscription, individual users install *FireAMP Connectors* on their computers and mobile devices (also called *endpoints*). These lightweight agents communicate with the Sourcefire cloud, which in turn communicates with the Defense Center. After you configure the Defense Center to connect to the cloud, you can use the Defense Center web interface to view endpoint-based malware events generated as a result of scans, detections, and quarantines on the endpoints in your organization.

Use the *FireAMP portal* (<http://amp.sourcefire.com/>) to configure your FireAMP deployment. The portal helps you quickly identify and quarantine malware. You can identify outbreaks when they occur, track their trajectories, understand their effects, and learn how to successfully recover. You can also use FireAMP to create custom protections, block execution of certain applications based on group policy, and create custom whitelists.

Network File Trajectory

The network file trajectory feature allows you to track a file's transmission path across a network. The system uses SHA-256 hash values to track files; so, to track a file, the system must either:

- calculate the file's SHA-256 hash value and perform a malware cloud lookup using that value
- receive endpoint-based threat and quarantine data about that file, using the Defense Center's integration with your organization's FireAMP subscription

Each file has an associated trajectory map, which contains a visual display of the file's transfers over time as well as additional information about the file.

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 the Support Site.

eStreamer

The Event Streamer (eStreamer) allows you to stream several kinds of event data from a Sourcefire appliance to a custom-developed client application. After you create a client application, you can connect it to an eStreamer server (Defense Center or managed device), start the eStreamerservice, and begin exchanging data.

eStreamer integration requires custom programming, but allows you to request specific data from an appliance. If, for example, you display network host data within one of your network management applications, you could write a program to retrieve host criticality or vulnerability data from the Defense Center and add that information to your display.

External Database Access

The database access feature allows you to query several database tables on a Sourcefire Defense Center, using a third-party client that supports JDBC SSL connections.

You can use an industry-standard reporting tool such as Crystal Reports, Actuate BIRT, or JasperSoft iReport to design and submit queries. Or, you can configure your own custom application to query Sourcefire data. For example, you could build a servlet to report intrusion and discovery event data periodically or refresh an alert dashboard.

Host Input

The host input feature allows you to augment the information in the network map by importing data from third-party sources using scripts or command-line files.

The web interface also provides some host input functionality; you can modify operating system or application protocol identities, validate or invalidate vulnerabilities, and delete various items from the network map, including clients and server ports.

Remediation

The system includes an API that allows you to create *remediations* that your Defense Center can automatically launch when conditions on your network violate an associated correlation policy or compliance white list. This can not only automatically mitigate attacks when you are not immediately available to address them, but can also ensure that your system remains compliant with your organization's security policy. In addition to remediations that you create, the Defense Center ships with several predefined remediation modules.

Security, Internet Access, and Communication Ports

To safeguard the Defense Center, you must install the Defense Center on a protected internal network. Although the Defense Center is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it from outside the firewall.

If the Defense Center and the managed device reside on the same network, you can connect the management interface on the device to the same protected internal network as the Defense Center. This allows you to securely control the device from the Defense Center and aggregate the event data generated on the managed device's network segment. By using the Defense Center's filtering capabilities, you can analyze and correlate data from attacks across your network to evaluate how well your security policies are being implemented.

Note, however, that Sourcefire appliances are configured to directly connect to the Internet. Specific features of the Sourcefire 3D System require this direct connection, and others support use of a proxy server. Additionally, the system requires that certain ports remain open for basic intra-appliance communication, as well as to allow you to access appliances' web interfaces. By default, several other ports are open to allow the system to take advantage of additional features and functionality.

For more information, see:

- [Internet Access Requirements](#) on page 50
- [Open Communication Ports Requirements](#) on page 52

Internet Access Requirements

By default, Sourcefire appliances are configured to directly connect to the Internet. Specific features of the Sourcefire 3D System require this direct connection, and others support use of a proxy server; see [Configuring Network](#)

[Settings](#) on page 2009.

TIP! You can manually upload system software, intrusion rule, GeoDB, and VDB updates to appliances.

To ensure continuity of operations, both Defense Centers in a high availability pair must have Internet access. For specific features, the primary Defense Center contacts the Internet, then shares information with the secondary during the synchronization process. Therefore, if the primary fails, you should promote the secondary to primary as described in [Monitoring and Changing High Availability Status](#) on page 202.

The following table describes the Internet access requirements of the Sourcefire 3D System.

Sourcefire 3D System Internet Access Requirements

FOR..	INTERNET ACCESS IS REQUIRED TO...	HIGH AVAILABILITY CONSIDERATIONS	PROXY?
RSS Feed dashboard widget	download RSS feed data from an external source, including Sourcefire.	Feed data is not synchronized.	yes
Security Intelligence feeds	download Security Intelligence feed data from an external source, including the Sourcefire Intelligence Feed.	The primary Defense Center downloads feed data and shares it with the secondary. In case of primary failure, you must switch roles.	yes
URL filtering data	download cloud-based URL category and reputation data for access control, and perform lookups for uncategorized URLs.	The primary Defense Center downloads URL filtering data and shares it with the secondary. In case of primary failure, you must switch roles.	yes
malware cloud lookups (Malware licensed)	perform cloud lookups to determine if files detected in network traffic contain malware.	Paired Defense Centers perform cloud lookups independently, although file policies are synchronized.	no

Sourcefire 3D System Internet Access Requirements (Continued)

FOR..	INTERNET ACCESS IS REQUIRED TO...	HIGH AVAILABILITY CONSIDERATIONS	PROXY?
FireAMP integration (FireAMP subscription)	receive endpoint-based malware events from the Sourcefire cloud.	Cloud connections are not synchronized. Configure them on both Defense Centers.	no
system, intrusion rule, GeoDB, and VDB updates	download or schedule the download of an intrusion rule, GeoDB, VDB, or system update directly to the appliance.	Rule, GeoDB, and VDB updates are synchronized; system updates are not. All appliances that download updates must have Internet access.	yes
obtaining whois information using the IP address context menu	obtain whois information.	Any appliance requesting whois information must have Internet access.	yes

Open Communication Ports Requirements

The Sourcefire 3D System requires that ports 443 (inbound) and 8305 (inbound and outbound) remain open for basic intra-appliance communication, as well as to allow you to access appliances' web interfaces.

By default, several other ports are open to allow the system to take advantage of additional features and functionality. The following table lists these ports. Note that DHCP on ports 67 and 68 is disabled by default.

Sourcefire 3D System Open Communication Ports Requirements

PORTS	DESCRIPTION	PROTOCOL	DIRECTION	THIS PORT IS OPEN TO...
22	SSH/SSL	TCP	Bidirectional	allow a secure remote connection to the appliance.
25	SMTP	TCP	Outbound	send email notices and alerts from the appliance.
53	DNS	TCP	Outbound	use DNS.
67, 68	DHCP	UDP	Outbound	use DHCP. Disabled by default.

Sourcefire 3D System Open Communication Ports Requirements (Continued)

PORTS	DESCRIPTION	PROTOCOL	DIRECTION	THIS PORT IS OPEN TO...
80	HTTP	TCP	Outbound or Bidirectional	allow the RSS Feed dashboard widget to connect to a remote web server (outbound). Adding inbound access allows the Defense Center to update custom and third-party Security Intelligence feeds via HTTP, and to download URL filtering information.
161, 162	SNMP	UDP	Bidirectional (161); Outbound (162)	provide access if you enabled SNMP polling (inbound) and SNMP traps (outbound).
389, 636	LDAP	TCP	Outbound	track user activity and for authentication.
443	HTTPS/AMQP	TCP	Inbound or Bidirectional	access the appliance. Required. Adding outbound access allows the Defense Center to download or receive software updates, VDB and GeoDB updates, URL filtering information, secure Security Intelligence feeds, and endpoint-based (FireAMP) malware events.
514	syslog	UDP	Outbound	send alerts to a remote syslog server.
623	SOL/LOM	UDP	Bidirectional	allow you to perform Lights-Out Management (LOM) using a Serial Over LAN (SOL) connection on a Series 3 appliance.
1500, 2000	database access	TCP	Inbound	access the Defense Center if external database access is enabled.
1812, 1813	RADIUS	UDP	Outbound or Bidirectional	use RADIUS. Adding inbound access ensures that RADIUS authentication and accounting function correctly. Ports 1812 and 1813 are the default, but you can configure RADIUS to use other ports; see Configuring RADIUS Connection Settings on page 1886.
3306	Sourcefire User Agent	TCP	Inbound	allow communication between the Defense Center and Sourcefire User Agents.

Sourcefire 3D System Open Communication Ports Requirements (Continued)

PORTS	DESCRIPTION	PROTOCOL	DIRECTION	THIS PORT IS OPEN TO...
8302	eStreamer	TCP	Bidirectional	use an eStreamer client.
8305	device management	TCP	Bidirectional	communicate between the Defense Center and managed devices. Required.
8307	host input client	TCP	Bidirectional	allow communication between the Defense Center and host input clients.
32137	malware cloud lookups	TCP	Outbound	allow the Defense Center to perform cloud lookups to determine if a file detected in network traffic contains malware, and to track files' trajectories.

Documentation Resources

The Sourcefire 3D System documentation set includes online help and PDF files. You can reach the online help in two ways:

- by clicking the context-sensitive help links on each page
- by selecting **Help > Online**

The online help includes information about the tasks you can complete on the web interface, including procedural and conceptual information about user management, system management, and event analysis.

The Documentation CD contains PDF versions of:

- the *Sourcefire 3D System User Guide*, which includes the same content as the online help, but in an easy-to-print format
- the *Sourcefire 3D System Installation Guide*, which includes information about installing Sourcefire appliances as well as hardware specifications and safety information
- the *Sourcefire 3D System Virtual Installation Guide*, which includes information about installing, managing, and troubleshooting virtual devices and virtual Defense Centers
- various API guides and supplementary material

You can access the most up-to-date versions of the PDF documentation on the Sourcefire Support site (<https://support.sourcefire.com/>).

Documentation Conventions

This documentation includes information about which Sourcefire 3D System licenses and appliance models are required for each feature, and which user roles

have permission to complete each procedure. For more information, see the following sections:

- [License Conventions](#) on page 55
- [Supported Device and Defense Center Conventions](#) on page 56
- [Access Conventions](#) on page 56

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:

FireSIGHT

A FireSIGHT license is included with your Defense Center and is required to perform host, application, and user discovery. The FireSIGHT license on your Defense Center determines how many individual hosts and users you can monitor with the Defense Center and its managed devices, as well as how many users you can use to perform user control.

If your Defense Center was previously running Version 4.10.x, you may be able to use legacy RNA Host and RUA User licenses instead of a FireSIGHT license.

Protection

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

Control

A Control license allows managed devices to perform user and application control. It also allows devices to perform switching and routing (including DHCP relay), NAT, and to cluster devices and stacks. A Control license requires a Protection license.

URL Filtering

A URL Filtering license allows managed 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 Protection and Control licenses.

Malware

A Malware license allows managed devices to perform network-based advanced malware protection (AMP), that is, to detect and block malware in files transmitted over your network. It also allows you to view trajectories, which track files transmitted over your network. A Malware license requires a Protection license.

VPN

A VPN license allows you to build secure VPN tunnels among the virtual routers on Sourcefire managed devices, or from managed devices to remote devices or other third-party VPN endpoints. A VPN license requires Protection and Control licenses.

Because licensed capabilities are often additive, this documentation only provides the highest required license for each feature. For example, if a feature requires FireSIGHT, Protection, and Control licenses, only Control is listed.

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.”

Note that because of architecture and resource limitations, not all licenses can be applied to all managed devices. In general, you cannot license a capability that a device does not support; see [Supported Capabilities by Appliance Model](#) on page 42. For more information on how your licenses affect the features you can use, including information on using legacy RNA Host and RUA User licenses, see [Understanding Licensing](#) on page 2036.

Supported Device and Defense Center Conventions

The Supported Devices statement at the beginning of a section indicates that a feature is supported only on the specified device series, family, or model. For example, stacking is only supported on Series 3 devices. If a section does not have a Supported Devices statement, the feature is supported on all devices, or the section does not apply to managed devices.

For more information on platforms supported by this release, see [Understanding Appliance Series, Models, and Capabilities](#) on page 40.

Access Conventions

The Access statement at the beginning of each procedure in this documentation indicates the predefined user role required to perform the procedure. A forward slash separating roles indicates that any of the listed roles can perform the

procedure. The following table defines common terms that appear in the Access statement.

Access Conventions

ACCESS TERM	INDICATES
Access Admin	User must have the Access Control Admin role
Admin	User must have the Administrator role
Any	User can have any role
Any/Admin	User can have any role, but only the Administrator role has unrestricted access (such as the ability to view other users' data saved as private)
Any Security Analyst	User can have either the Security Analyst or Security Analyst (Read Only) role
Database	User must have the External Database role
Discovery Admin	User must have the Discovery Admin role
Intrusion Admin	User must have the Intrusion Admin role
Maint	User must have the Maintenance User role
Network Admin	User must have the Network Admin role
Security Analyst	User must have the Security Analyst role
Security Approver	User must have the Security Approver role

Users with custom roles may have permission sets that differ from those of the predefined roles. When a predefined role is used to indicate access requirements for a procedure, a custom role with similar permissions also has access. For more information on custom user roles, see [Managing Custom User Roles](#) on page 1908.

IP Address Conventions

You can use IPv4 Classless Inter-Domain Routing (CIDR) notation to define address blocks in many places in the Sourcefire 3D System. In those places that allow IPv6 addresses (for example, in intrusion policies) you can also use the similar IPv6 prefix length notation.

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.

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 Sourcefire 3D System 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 Sourcefire 3D System uses 10.0.0.0/8.

In other words, although Sourcefire recommends the standard method of using a network IP address on the bit boundary when using CIDR or prefix length notation, the Sourcefire 3D System does not require it.

Logging into the Appliance

LICENSE: Any

The Defense Center has a web-based interface that you can use to perform administrative, management, and analysis tasks. Physical managed devices have a limited web-based interface that you can use to perform initial setup and basic analysis and configuration tasks. Virtual managed devices do not have web interfaces. You can access the web interface by logging into the appliance using a web browser. For information on browser requirements, refer to the release notes for this version of the Sourcefire 3D System.

If you are the first user to log into the appliance after it is installed, you must log in using the administrative (**admin**) user account to complete the initial setup process, which is described in the *Sourcefire 3D System Installation Guide*. After you create other user accounts as described in [Adding New User Accounts](#) on

page 1899, you and other users should use those accounts to log in to the web interface.

IMPORTANT! Because Sourcefire appliances audit user activity based on user accounts, make sure that users log into the system with the correct account.

After you log into the appliance, the features you can access are controlled by the privileges granted to your user account. However, the procedures for logging into and out of the appliance remain the same. If your organization uses SecurID® tokens when logging in, append the token to your SecurID PIN and use that as your password to log in. For example, if your PIN is 1111 and the SecurID token is 222222, type 1111222222.

WARNING! If you supply incorrect credentials multiple times, your shell access account may be locked. Even after the account is locked, you are prompted to supply your credentials again. If you supply correct credentials and the login is refused, contact your system administrator rather than repeatedly attempting to log in.

The first time you visit the appliance home page during a web session, you can view information about your last login session for that appliance. You can see the following information about your last login:

- the day of the week, month, date, and year of the login
- the appliance-local time of the login in 24-hour notation
- the host and domain name last used to access the appliance

By default, your session automatically logs you out after 1 hour of inactivity, unless you are otherwise configured to be exempt from session timeout. Users with the Administrator role can change the session timeout interval in the system policy. For more information, see [Managing User Login Settings](#) on page 1903 and [Configuring User Interface Settings](#) on page 1994.

Note that some processes that take a significant amount of time may cause your web browser to display a message that a script has become unresponsive. If this occurs, make sure you allow the script to continue until it finishes.

To log into the appliance via the web interface:

ACCESS: Any

1. Direct your browser to `https://hostname/`, where *hostname* corresponds to the host name of the appliance.

The Login page appears.

2. In the **Username** and **Password** fields, type your user name and password. User names are case sensitive.
If your company uses SecurID, append the SecurID token to the end of your SecurID PIN and use that as your password when you log in. You must have already generated your SecurID PIN before you can log into the Sourcefire 3D System.
3. Click **Login**.
The default start page appears. If you selected a new home page for your user account, that page is displayed instead. See [Specifying Your Home Page](#) on page 2217 for more information.
The menus and menu options listed at the top of the page are based on the privileges for your user account. However, the links on the default home page include options that span the range of user account privileges. If you click a link that requires different privileges from those granted to your account, the following warning message is displayed:

You are attempting to view an unauthorized page. This activity has been logged.

You can either select a different option from the available menus or click **Back** in your browser window.

To log into a Series 3 or virtual device via the command line:

ACCESS: CLI Basic Configuration

1. Open an SSH connection to the appliance at *hostname*, where *hostname* corresponds to the host name of the appliance.
The **login as:** command prompt appears.
2. Type your user name and press Enter.
The **Password:** prompt appears.
3. Type your password and press Enter.
The login banner appears, followed by the > prompt.
You can use any of the commands allowed by your level of command line access. See the [Command Line Reference](#) on page 2242 for more information on available CLI commands.

Logging into the Appliance to Set Up an Account

LICENSE: Any

Some user accounts may be authenticated through an external authentication server. If your organization allows you to log on to the Sourcefire 3D System using LDAP or RADIUS credentials, the first time you log into the appliance using your external user credentials, the appliance associates those credentials with a set of permissions by creating a local user record. The permissions for that local user

record can then be modified, unless they are granted through group or list membership, as follows:

- If the default role for externally authenticated user accounts is set to a specific access role, you can log into the appliance using your external account credentials without any additional configuration by the system administrator.
- If an account is externally authenticated and by default receives no access privileges, you can log in but cannot access any functionality. You (or your system administrator) can then change the permissions to grant the appropriate access to user functionality.

If you are a shell access user, the system does not create a local user account for you on the appliance. Shell access is controlled entirely through either the shell access filter or PAM login attribute set for an LDAP server, or the shell access list on a RADIUS server.

Shell users should log in using usernames with all lowercase letters. LDAP usernames can include underscores (_), periods (.), and hyphens (-), but otherwise only alphanumeric characters are supported.

If your organization uses SecurID tokens when logging in, append the token to your SecurID PIN and use that as your password to log in. For example, if your PIN is 1111 and the SecurID token is 222222, type 1111222222.

IMPORTANT! If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account. For more information, see [Modifying User Privileges and Options](#) on page 1912.

To create an externally authenticated account on the appliance:

ACCESS: Any

1. Direct your browser to `https://hostname/`, where *hostname* corresponds to the host name of the appliance.
The Login page appears.
2. Type your **Username** and **Password**.

IMPORTANT! If your company uses SecurID, append the SecurID token to your SecurID PIN and use that as your password when you log in.

3. Click **Login**.

The page that appears depends on the default access role for external authentication:

- If a default access role is selected in the authentication object or the system policy, the default start page appears. If you selected a new home page for your user account, then that page is displayed instead. See [Specifying Your Home Page](#) on page 2217 for more information.

The menus and menu options that are available to you at the top of the page are based on the privileges for your user account. However, the links on the default home page include options that span the range of user account privileges. If you click a link that requires different privileges from those granted to your account, the following warning message is displayed:

You are attempting to view an unauthorized page. This activity has been logged.

You can either select a different option from the available menus or click **Back** in your browser window.

- If no default access role is selected, the Login page reappears, with the following error message:

unable to authorize access. If you continue to have difficulty accessing this device, please contact the system administrator.

Note that when you use a RADIUS server using attribute matching as your authentication method, your first login attempt is rejected as your user account is created. You must log in a second time.

Logging Out of the Appliance

LICENSE: Any

When you are no longer actively using the web interface, Sourcefire recommends that you log out, even if you are only stepping away from your web browser for a short period of time. Logging out ends your web session and ensures that no one can use the appliance with your credentials.

By default, your session automatically logs you out after 1 hour of inactivity, unless you are otherwise configured to be exempt from session timeout. Users with the Administrator role can change the session timeout interval in the system policy. For more information, see [Managing User Login Settings](#) on page 1903 and [Configuring User Interface Settings](#) on page 1994.

To log out of the appliance:

ACCESS: Any

- ▶ Click **Logout** on the toolbar.

Using the Context Menu

LICENSE: feature dependent

For your convenience, certain pages in the web interface support a pop-up context menu that you can use as a shortcut for accessing other features in the Sourcefire 3D System. The contents of the menu depend on the *hotspot* where you access it—not only the page but also the specific data.

For example, *IP address hotspots* in event views, intrusion event packet views, the dashboard, and the Context Explorer provide additional options. Use the IP address context menu by right-clicking on the hotspot to learn more about the host associated with that address, including any available whois and host profile information. Except on the DC500 Defense Center, which does not support Security Intelligence filtering, you can also add an individual IP address to the Security Intelligence global whitelist or blacklist.

As another example, *SHA-256 value hotspots* in event views and the dashboard allow you to add a file's SHA-256 hash value to the global malware whitelist. Note that this functionality is also not supported on the DC500 Defense Center.

The following list describes many of the options available in the context menu on various pages of the web interface. On pages or locations where the Sourcefire context menu is not supported, the normal context menu for your browser appears.

Access Control Policy Editor

The access control policy editor contains hotspots over each access control rule. You can use the context menu to insert new rules and categories; cut, copy, and paste rules; set the rule state; and edit the rule.

NAT Policy Editor

The NAT policy editor contains hotspots over each NAT rule. You can use the context menu to insert new rules; cut, copy, and paste rules; set the rule state; and edit the rule.

Intrusion Rule Editor

The intrusion rule editor contains hotspots over each intrusion rule. You can use the context menu to edit the rule, set the rule state (including disabling the rule), configure thresholding and suppression options, and view rule documentation.

Event Viewer

Event pages (drill-down pages and table views) contain hotspots over each events, IP address, and certain detected files' SHA-256 hash values. For most event types, you can use the context menu to view related information in the Context Explorer, or drill down into event information in a new window. In places where an event field contains text too long to fully display in the event view, such as a file's SHA-256 hash value, a vulnerability description, or a URL, you can use the context menu to view the full text.

For intrusion events, you can also use the context menu to perform similar tasks to those in the intrusion rule editor: edit the triggering rule, set the rule state (including disabling the rule), configure thresholding and suppression options, and view rule documentation.

Packet View

Intrusion event packet views contain IP address hotspots. Note that the packet view uses a left-click context menu instead of a right-click menu.

Dashboard

Many dashboard widgets contain hotspots to view related information in the Context Explorer. Dashboard widgets can also contain IP address and SHA-256 value hotspots.

Context Explorer

The Context Explorer contains hotspots over its charts, tables, and graphs. If you want to examine data from graphs or lists in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. You can also view related host, user, application, file, and intrusion rule information.

Note that the Context Explorer uses a left-click context menu, which also contains filtering and other options unique to the Context Explorer. For detailed information, see [Drilling Down on Context Explorer Data](#) on page 148.

To access the context menu:

ACCESS: Any

1. On a hotspot-enabled page in the web interface, hover your pointer over a hotspot.
Except in the Context Explorer, a **right-click for menu** message appears.
2. Invoke the context menu:
 - In the Context Explorer or packet view, left-click your pointing device.
 - On all other hotspot-enabled pages, right-click your pointing device.A pop-up context menu appears with options appropriate for the hotspot.

3. Select one of the options by left-clicking the name of the option.
A new browser window opens based on the option you selected.

CHAPTER 2

USING DASHBOARDS

The Sourcefire 3D System dashboard provides you with at-a-glance views of current system status, including data about the events collected and generated by the system. You can also use the dashboard to see information about the status and overall health of the appliances in your deployment. Only certain user roles (Administrator, Maintenance User, Security Analyst, Security Analyst [Read Only], and custom roles with the Dashboards permission) have access to the dashboard. Other roles see as their default start pages a page relevant to the role; for example, a Discovery Admin sees the Network Discovery page.

A dashboard has one or more tabs, each of which can display one or more widgets in a three-column layout. Widgets are small, self-contained components that provide insight into different aspects of the Sourcefire 3D System. The Sourcefire 3D System is delivered with several predefined widgets. For example, the Appliance Information widget tells you the appliance name, model, remote manager, and currently running version of the Sourcefire 3D System 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.

The dashboard is a complex, highly customizable monitoring feature. Another way to view many types of system data is the Context Explorer, which presents information using intrusion, connection, and discovery data in a set of preset visual contexts that you change, only temporarily, with filters to add granularity. In contrast to the exhaustive data available in the Sourcefire 3D System dashboard, the Context Explorer offers a broad, brief, and colorful picture of how your monitored network looks and acts. For more information on the Context Explorer, see [Using the Context Explorer](#) on page 117.

Each type of appliance is delivered with a default dashboard, named Summary Dashboard. This dashboard provides the casual user with general FireSIGHT,

intrusion, geolocation, and system status information for your Sourcefire 3D System deployment. Note that because some widgets are useful only for specific types of appliances, the Summary Dashboard differs depending on whether you are using a Defense Center or managed device.

By default, the home page for your appliance displays the Summary Dashboard, although you can configure your appliance to display a different default home page.

TIP! If you change the home page, you can access dashboards by selecting **Overview > Dashboards**. For more information, see [Viewing Dashboards](#) on page 109.

Note that the data displayed depends on such factors as how you license and deploy your managed devices, whether you configure features that provide the data and, in the case of Series 2 appliances, whether the appliance supports a feature that provides the data. For example, because neither the DC500 Defense Center nor Series 2 devices support URL filtering by category and reputation, the DC500 Defense Center does not display data for this feature and Series 2 devices do not detect this data.

In addition to the Summary Dashboard, the Defense Center is delivered with the following predefined dashboards:

- The Application Statistics dashboard provides detailed information about application activity and intrusion events on your monitored network. You can use this dashboard to track which applications produce the most traffic, allowed and denied connections, and intrusion events, as well as the number of unique applications in use and the estimated risk and business relevance of those applications.
- The Connection Summary dashboard uses connection data to create tables and charts of the activity on your monitored network. You can use this dashboard to track the ports, applications, and initiator and responder IPs associated with connections and traffic on your network, the overall volume of connections and traffic, and geolocation information. You must log connections for this dashboard to generate data; see [Understanding Connection Data](#) on page 540. Note that the output of this widget depends on your connection logging configuration.

TIP! Widgets on this dashboard list total traffic in kilobytes (KB). The total traffic in KB is equal to the traffic in KB/s multiplied by the total seconds covered by the selected time window.

- The Detailed Dashboard provides advanced users with detailed information about their Sourcefire 3D System deployment and includes multiple widgets that summarize collected intrusion event, network discovery, compliance, correlation, traffic, and system status data, as well as providing information about Sourcefire news and product updates. You can use this dashboard to monitor a very broad variety of network information at once.
- The Files Dashboard provides detailed information about the files (including malware files) detected on your network by managed devices, as well as malware detected using a subscription-based FireAMP strategy. Note that you must have a Malware license and enable malware detection for this dashboard to include network-based malware data. Note also that neither the DC500 nor Series 2 devices support advanced malware detection, so the DC500 cannot display this data and Series 2 devices do not detect it. For more information, see [Working with Malware Protection and File Control](#) on page 1197.
- The URL Statistics dashboard provides detailed information about allowed and denied traffic from your monitored network to external URLs, sorted by URL category and reputation. Note that you must have a URL Filtering license and enable URL Filtering for this dashboard to include URL category and reputation data. Note also that neither the DC500 nor Series 2 devices support URL filtering by reputation and category, so the DC500 cannot display this data and Series 2 devices do not detect it. See [Adding URL Conditions](#) on page 506.
- The User Statistics dashboard provides detailed information about user activity and intrusion events on your monitored network. You can use this dashboard to track allowed and denied connections, traffic, and intrusion events associated with users on your network, as well as the number of unique users on the network. Because this dashboard depends on user awareness data, for this dashboard to display meaningful statistics you must configure at least one Sourcefire User Agent and a Defense Center-Active Directory LDAP server connection; see [Obtaining User Data from LDAP Servers](#) on page 1298.

You can use the predefined dashboards, modify the predefined dashboards, or create a custom dashboard to suit your needs. You can share custom dashboards among all users of an appliance, or you can create a custom dashboard solely for your own use. You can also set a custom dashboard as your default dashboard.

Some drill-down pages and table views of events include a **Dashboard** toolbar link that you can click to view a relevant predefined dashboard. The [Event Table Dashboard Links](#) table lists which event views correspond to which predefined

dashboards. Note that if you delete a predefined dashboard or tab, the associated Dashboard links do not function.

Event Table Dashboard Links

TABLE	DASHBOARD LINK
Connection Events (Analysis > Connections)	Connection Summary
Intrusion Events (Analysis > Intrusion > Events)	Summary (Intrusion Events tab)
Malware Events (Analysis > Malware)	Files (Malware tab)
File Events (Analysis > Files)	Files (Files tab)
Applications (Analysis > Hosts > Applications)	Application Statistics
Application Details (Analysis > Hosts > Application Details)	Application Statistics
Users (Analysis > Users > Users)	User Statistics
User Activity (Analysis > Users > User Activity)	User Statistics
Correlation Events (Analysis > Correlation > Correlation Events)	Detailed (Correlation tab)
White List Events (Analysis > Correlation > White List Events)	Detailed (Correlation tab)

For more information on dashboards and their contents, see the following sections:

- [Understanding Dashboard Widgets](#) on page 70
- [Understanding the Predefined Widgets](#) on page 74
- [Working with Dashboards](#) on page 106

Understanding Dashboard Widgets

LICENSE: Any

A dashboard has one or more tabs, each of which can display one or more widgets in a three-column layout. The Sourcefire 3D System is delivered with many predefined dashboard widgets, each of which provides insight into a different aspect of the Sourcefire 3D System. Widgets are grouped into three categories:

- *Analysis & Reporting widgets* display data about the events collected and generated by the Sourcefire 3D System.
- *Miscellaneous widgets* display neither event data nor operations data. Currently, the only widget in this category displays an RSS feed.
- *Operations widgets* display information about the status and overall health of the Sourcefire 3D System.

The dashboard widgets that you can view depend on the type of appliance you are using and on your user role. In addition, each dashboard has a set of preferences that determines its behavior. You can minimize and maximize widgets, add and remove widgets from tabs, as well as rearrange the widgets on a tab.

IMPORTANT! For widgets that display event counts over a time range, the total number of events may not reflect the number of events for which detailed data is available in the event viewer. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment. For more information, see [Logging Connection, File, and Malware Information](#) on page 515.

For more information, see:

- [Understanding Widget Availability](#) on page 71
- [Understanding Widget Preferences](#) on page 73
- [Understanding the Predefined Widgets](#) on page 74
- [Working with Dashboards](#) on page 106

Understanding Widget Availability

LICENSE: Any

The Sourcefire 3D System is delivered with several predefined dashboard widgets. The dashboard widgets that you can view depend on the type of appliance you are using and on your user role:

- An *invalid* widget is one that you cannot view because you are using the wrong type of appliance.
- An *unauthorized* widget is one that you cannot view because you do not have the necessary account privileges.

For example, the Current Sessions widget is available on all appliances, but only to users with Administrator account privileges, while the Appliance Status widget is available only on the Defense Center for users with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) account privileges.

Although you cannot add an unauthorized or invalid widget to a dashboard, if you import a dashboard created either on a different kind of appliance or by a user with different access privileges, that dashboard may contain unauthorized or invalid widgets. These widgets are disabled and display error messages that indicate the reason why you cannot view them.

Also note that widgets cannot display data to which an appliance has no access. For example, managed devices cannot access correlation events, intrusion events, discovery events, and so on. If you import a dashboard onto a managed device that contains a Custom Analysis widget configured to display one of those data types, the widget displays an error message. Individual widgets also display error messages when those widgets have timed out or are otherwise experiencing problems.

The content of a widget can differ depending on the type of appliance you are using. For example, the Custom Analysis widget on a Defense Center can display discovery information, but this feature is not available when you configure the Custom Analysis widget on a managed device. Note that you can sort any content generated in table format by clicking on the table column header.

You can delete or minimize unauthorized and invalid widgets, as well as widgets that display no data, keeping in mind that modifying a widget on a shared dashboard modifies it for all users of the appliance. For more information, see [Minimizing and Maximizing Widgets](#) on page 115 and [Deleting Widgets](#) on page 116.

The [Sourcefire Appliances and Dashboard Widget Availability](#) table lists the valid widgets for each appliance. An X indicates that the appliance can display the widget.

Sourcefire Appliances and Dashboard Widget Availability

WIDGET	DEFENSE CENTER	ANY MANAGED DEVICE
Appliance Information	X	X
Appliance Status	X	
Correlation Events	X	
Current Interface Status	X	X
Current Sessions	X	X
Custom Analysis	X	
Disk Usage	X	X
Interface Traffic	X	X
Intrusion Events	X	
Network Compliance	X	
Product Licensing	X	
Product Updates	X	X
RSS Feed	X	X
System Load	X	X
System Time	X	X
White List Events	X	

The [User Roles and Dashboard Widget Availability](#) table lists the user account privileges required to view each widget. An X indicates the user can view the widget. Only user accounts with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) access can use dashboards.

Users with custom roles may have access to any combination of widgets, or none at all, as their user roles permit.

User Roles and Dashboard Widget Availability

WIDGET	ADMINISTRATOR	MAINTENANCE USER	SECURITY ANALYST	SECURITY ANALYST (RO)
Appliance Information	X	X	X	X
Appliance Status	X	X	X	
Correlation Events	X		X	X
Current Interface Status	X	X	X	X
Current Sessions	X			
Custom Analysis	X		X	X
Disk Usage	X	X	X	X
Interface Traffic	X	X	X	X
Intrusion Events	X		X	X
Network Compliance	X		X	X
Product Licensing	X	X		
Product Updates	X	X		
RSS Feed	X	X	X	X
System Load	X	X	X	X
System Time	X	X	X	X
White List Events	X		X	X

Understanding Widget Preferences

LICENSE: Any

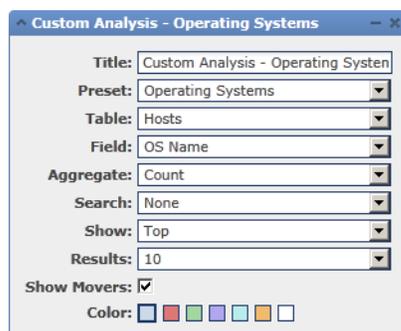
Each widget has a set of preferences that determines its behavior.

Widget preferences can be simple. For example, the following graphic shows the 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.



Widget preferences can also be more complex. For example, the following graphic shows the preferences for the Custom Analysis widget, which is a highly customizable widget that allows you to display detailed information on the events collected and generated by the Sourcefire 3D System.



To modify a widget's preferences:

ACCESS: Admin/Any Security Analyst/Maint

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.
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](#) on page 74.
3. On the widget title bar, click the hide preferences icon (▲) to hide the preferences section.

Understanding the Predefined Widgets

LICENSE: Any

The Sourcefire 3D System is delivered with several predefined widgets that, when used on dashboards, can provide you with at-a-glance views of current system status, including data about the events collected and generated by the system, as well as information about the status and overall health of the appliances in your deployment.

For detailed information on the widgets delivered with the Sourcefire 3D System, see the following sections:

- [Understanding the Appliance Information Widget](#) on page 75
- [Understanding the Appliance Status Widget](#) on page 76
- [Understanding the Correlation Events Widget](#) on page 77
- [Understanding the Current Interface Status Widget](#) on page 77
- [Understanding the Current Sessions Widget](#) on page 78
- [Understanding the Custom Analysis Widget](#) on page 79
- [Understanding the Disk Usage Widget](#) on page 96
- [Understanding the Interface Traffic Widget](#) on page 97
- [Understanding the Intrusion Events Widget](#) on page 97
- [Understanding the Network Compliance Widget](#) on page 99
- [Understanding the Product Licensing Widget](#) on page 101
- [Understanding the Product Updates Widget](#) on page 101
- [Understanding the RSS Feed Widget](#) on page 103
- [Understanding the System Load Widget](#) on page 104
- [Understanding the System Time Widget](#) on page 104
- [Understanding the White List Events Widget](#) on page 105

IMPORTANT! The dashboard widgets you can view depend on the type of appliance you are using and on your user role. For more information, see [Understanding Widget Availability](#) on page 71.

Understanding the Appliance Information Widget

LICENSE: Any

The Appliance Information widget provides a snapshot of the appliance. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.



The screenshot shows a window titled "Appliance Information" with a list of system parameters and their values. The parameters are grouped into "Versions" and other system details.

Name	katsura
IPv4 Address	10.10.0.2 (eth0)
IPv6 Address	Disabled
Model	Defense Center 3500 (66)
Versions	
Software	5.0.0-652
OS	Sourcefire Linux OS 5.0.0-27
Snort	2.9.2-41
Rule Update	2011-08-30-001-dev
Geolocation Update	None
Rulepack	753
Module Pack	1253
VDB	70.2017

The widget provides:

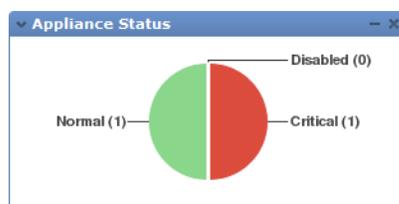
- the name, IPv4 address, IPv6 address, and model of the appliance
- the versions of the Sourcefire 3D System software, operating system, Snort, rule update, rule pack, module pack, vulnerability database (VDB), and geolocation update installed on the appliance
- for managed appliances, the name and status of the communications link with the managing appliance
- for Defense Centers in a high availability pair, the name, model, and Sourcefire 3D System software and operating system versions of the peer Defense Center, as well as how recently the Defense Centers made contact

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](#) on page 73.

Understanding the Appliance Status Widget

LICENSE: Any

The Appliance Status widget indicates the health of the appliance and of any appliances it is managing. Note that because the Defense Center does not automatically apply a health policy to managed devices, you must manually apply a health policy to devices or their status appears as **Disabled**. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.



You can configure the widget to display appliance status as a pie chart or in a table by modifying the widget preferences.

A table titled "Appliance Status" showing the count of appliances in different health states. The table has two columns: "Type" and a count column. The "Type" column includes icons for error (red X), warning (yellow triangle), normal (green checkmark), and info (blue question mark).

Type	Count
Managed Device	1
Defense Center	1

The preferences also control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

You can click a section on the pie chart or one of the numbers on the appliance status table to go to the Health Monitor page and view the compiled health status

of the appliance and of any appliances it is managing. For more information, see [Using the Health Monitor](#) on page 2163.

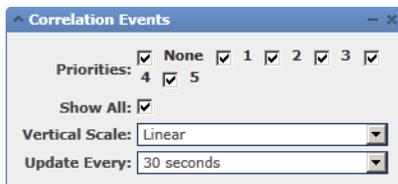
Understanding the Correlation Events Widget

LICENSE: FireSIGHT

The Correlation Events widget shows the average number of correlation events per second, by priority, over the dashboard time range. It appears by default on the Correlation tab of the Detailed Dashboard.



You can configure the widget to display correlation events of different priorities by modifying the widget preferences, as well as to select a linear (incremental) or logarithmic (factor of ten) scale.



Select one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority. Select **Show All** to display an additional graph for all correlation events, regardless of priority. The preferences also control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

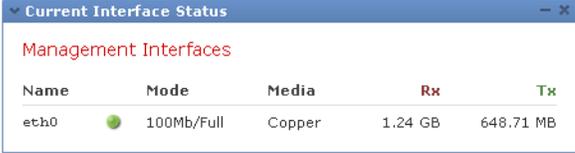
You can click a graph to view correlation events of a specific priority, or click the **All** graph to view all correlation events. In either case, the events are constrained by the dashboard time range; accessing correlation events via the dashboard changes the events (or global) time window for the appliance. For more information on correlation events, see [Viewing Correlation Events](#) on page 1521.

Understanding the Current Interface Status Widget

LICENSE: Any

The Current Interface Status widget shows the status of all enabled interfaces for the appliance, grouped by type: management, inline, passive, switched, routed, stacking, and unused. Note that Defense Centers have only management

interfaces. This widget does not appear by default on any of the predefined dashboards.



The screenshot shows a window titled "Current Interface Status" with a sub-header "Management Interfaces". Below the header is a table with columns: Name, Mode, Media, Rx, and Tx. The table contains one row for the interface "eth0".

Name	Mode	Media	Rx	Tx
eth0	100Mb/Full	Copper	1.24 GB	648.71 MB

For each interface, the widget provides:

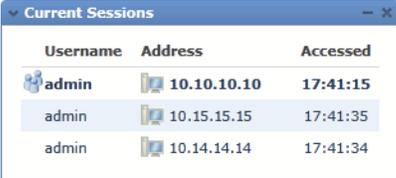
- the name of the interface
- the link state of the interface, represented by a green ball (up) or a gray ball (down)
- the link mode (for example, 100Mb full duplex, or 10Mb half duplex) of the interface
- the type of interface, that is, copper or fiber
- the amount of data received (Rx) and transmitted (Tx) by the interface

The widget preferences control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

Understanding the Current Sessions Widget

LICENSE: Any

The Current Sessions widget shows which users are currently logged into the appliance, the IP address associated with the machine where the session originated, and the last time each user accessed a page on the appliance (based on the local time for the appliance). The user that represents you, that is, the user currently viewing the widget, is marked with a user icon (👤) and rendered in bold type. Sessions are pruned from this widget's data within one hour of logoff or inactivity. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.



The screenshot shows a window titled "Current Sessions" with a table containing columns: Username, Address, and Accessed. The first row shows a session for "admin" with IP "10.10.10.10" accessed at "17:41:15". The second row shows a session for "admin" with IP "10.15.15.15" accessed at "17:41:35". The third row shows a session for "admin" with IP "10.14.14.14" accessed at "17:41:34". The first row is highlighted and has a user icon next to the username.

Username	Address	Accessed
admin	10.10.10.10	17:41:15
admin	10.15.15.15	17:41:35
admin	10.14.14.14	17:41:34

On the Current Sessions widget, you can:

- click any user name to manage user accounts on the User Management page; see [Managing User Accounts](#) on page 1897
- click the host icon (🖨️) next to any IP address to view the host profile for the associated machine; see [Using Host Profiles](#) on page 1334 (Defense Center with network discovery only)
- click any IP address or access time to view the audit log constrained by that IP address and by the time that the user associated with that IP address logged on to the web interface; see [Viewing Audit Records](#) on page 2188

The widget preferences control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

Understanding the Custom Analysis Widget

LICENSE: Any

The Custom Analysis widget is a highly customizable widget that allows you to display detailed information on the events collected and generated by the Sourcefire 3D System.

The Custom Analysis widget is delivered with numerous widget presets, which are groups of configurations that are predefined by Sourcefire. The presets serve as examples and can provide quick access to information about your deployment. You can use these presets or create a custom configuration.

When you configure the widget preferences, you must select which table and individual field you want to display, as well as the aggregation method that configures how the widget groups the data it displays.

For example, you can configure the Custom Analysis widget to display a list of recent intrusion events by configuring the widget to display data from the **Intrusion Events** table. Selecting the **Classification** field and aggregating this data by **Count** tells you how many events of each type were generated.

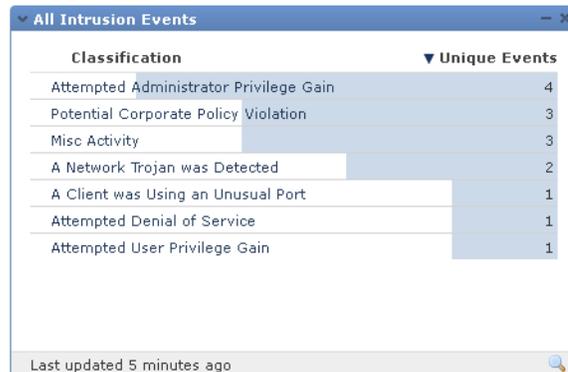


Classification	Count
A Client was Using an Unusual Port	15,003
Potential Corporate Policy Violation	955
Attempted User Privilege Gain	42
Attempted Administrator Privilege Gain	18
Misc Activity	16
A Network Trojan was Detected	5
Attempted Denial of Service	1

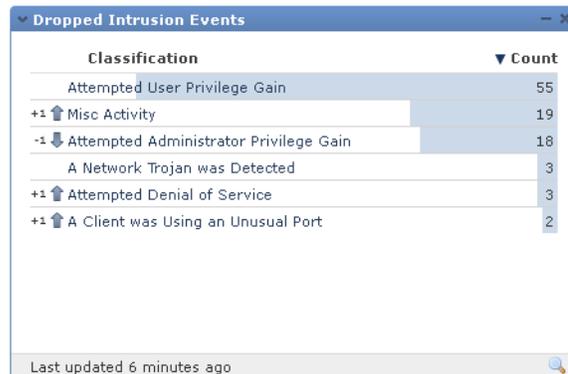
Last updated 1 minute ago

On the other hand, aggregating by **Unique Events** tells you how many unique intrusion events of each type have occurred (for example, how many detections

of network trojans, potential violations of corporate policy, attempted denial-of-service attacks, and so on).



Optionally, you can further constrain the widget using a saved search, either one of the predefined searches delivered with your appliance or a custom search that you created. For example, constraining the first example (intrusion events using the **Classification** field, aggregated by **Count**) using the **Dropped Events** search tells you how many intrusion events of each type were dropped.



The colored bars in the widget background show the relative number of occurrences of each event; you should read the bars from right to left. You can change the color of the bars as well as the number of rows that the widget displays. You can also configure the widget to display the most frequently occurring events or the least frequently occurring events.

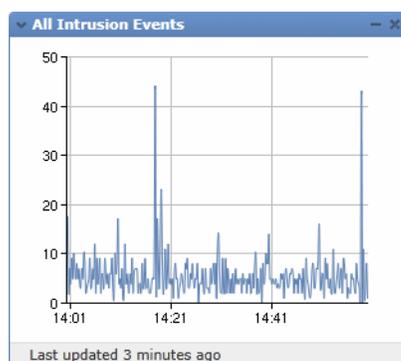
The direction icon (▼) indicates and controls the sort order of the display. A downward-pointing icon indicates descending order; an upward-pointing icon indicates ascending order. To change the sort order, click the icon.

Next to each event, the widget can display one of three icons to indicate any changes from the most recent results:

- The new event icon (+) signifies that the event is new to the results.
- The up arrow icon (↑) indicates that the event has moved up in the standings since the last time the widget updated. A number indicating how many places the event has moved up appears next to the icon.
- The down arrow icon (↓) indicates that the event has moved down in the standings since the last time the widget updated. A number indicating how many places the event has moved down appears next to the icon.

The widget displays the last time it updated, based on the local time of the appliance. The widget updates with a frequency that depends on the dashboard time range. For example, if you set the dashboard time range to an hour, the widget updates every five minutes. On the other hand, if you set the dashboard time range to a year, the widget updates once a week. To determine when the dashboard will update next, hover your pointer over the **Last updated** notice in the bottom left corner of the widget.

If you want information on events or other collected data over time, you can configure the Custom Analysis widget to display a line graph, such as one that displays the total number of intrusion events generated in your deployment over time. For graphs over time, you can choose the time zone that the widget uses as well as the color of the line.



Finally, you can choose a custom title for the widget.

From Custom Analysis widgets, you can invoke event views (that is, workflows) that provide detailed information about the events displayed in the widget. To do so, click the event for which you want more information.

You can also right-click any IP address in the custom analysis widget to display a context menu that allows you to obtain more information on the associated host,

as well as add it to the global blacklist or whitelist for Security Intelligence filtering.

IMPORTANT! Depending on how you configure them, Custom Analysis widgets may place a drain on an appliance's resources; a red-shaded Custom Analysis widget indicates that its use is harming system performance. If the widget continues to stay red over time, you should remove the widget.

For more information, see the following sections:

- [Configuring the Custom Analysis Widget](#) on page 82
- [Viewing Associated Events from the Custom Analysis Widget](#) on page 94
- [Custom Analysis Widget Limitations](#) on page 96
- [Using the Context Menu](#) on page 63

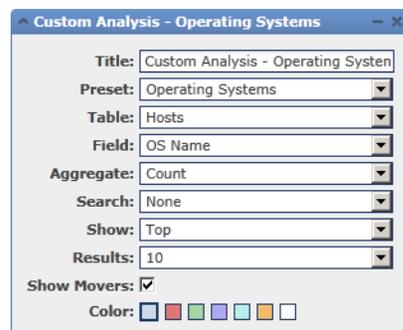
Configuring the Custom Analysis Widget

LICENSE: Any

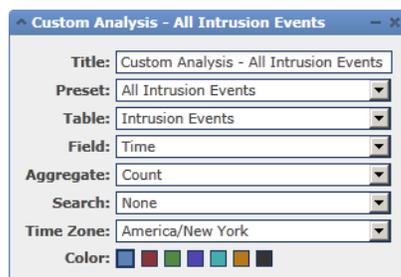
Like all widgets, the Custom Analysis widget has preferences that determine its behavior. To configure a Custom Analysis widget, show the preferences as described in [Understanding Widget Preferences](#) on page 73.

A different set of preferences appears depending on whether you configure the widget to show relative occurrences of events (that is, a bar graph), or configure the widget to show a graph over time (that is, a line graph).

To configure the widget to show a bar graph, select any value except **Time** from the **Field** drop-down list.



To configure the widget to show a line graph, select **Time** from the **Field** drop-down list.



The following table describes the various preferences you can set in the Custom Analysis widget.

Custom Analysis Widget Preferences

USE THIS PREFERENCE...	TO CONTROL...
Title	<p>the title of the widget.</p> <p>If you do not specify a title, the appliance uses the configured event type as the widget title.</p>
Preset	<p>the preset for the widget.</p> <p>The Custom Analysis widget is delivered with numerous presets, which are widget configurations predefined by Sourcefire. The presets serve as examples and can provide quick access to information about your deployment. You can use these presets or you can create a custom configuration.</p> <p>For a detailed list of presets, see the Custom Analysis Widget Presets table on page 85.</p>
Table	<p>the table of events which contains the event data the widget displays.</p>
Field	<p>the specific field of the event type you want to display.</p> <p>TIP! To display a graph over time, select Time.</p>
Aggregate	<p>the aggregation method for the widget.</p> <p>The aggregation method configures how the widget groups the data it displays. For most event types, the default aggregation criterion is Count.</p>

Custom Analysis Widget Preferences (Continued)

USE THIS PREFERENCE...	TO CONTROL...
Filter	<p>a user-defined application filter that you want to use to further constrain the data that the widget displays.</p> <p>You can only use application filters if you are displaying data from the Application Statistics or Intrusion Event Statistics by Application tables. For more information on application filters, see Working with Application Filters on page 176.</p>
Search	<p>the saved search you want to use to further constrain the data that the widget displays.</p> <p>You do not have to specify a search, although some presets use predefined searches.</p>
Show	<p>whether you want to display the most frequently occurring events (Top) or the least frequently occurring events (Bottom).</p>
Results	<p>the number of result rows you want to display.</p> <p>You can display from 10 to 25 result rows, in increments of five.</p>
Show Movers	<p>whether you want to display the icons that indicate changes from the most recent results.</p>
Time Zone	<p>which time zone you want to use to display results.</p> <p>The time zone appears whenever you select a time-based field.</p>
Color	<p>the color of the bars in the widget background that show the relative number of occurrences of each result.</p>

The following table describes the available presets for the Custom Analysis widget. It also indicates which, if any, Defense Center predefined dashboard uses each preset. Note the following:

- Predefined dashboards on managed devices do not include Custom Analysis widgets.
- The DC500 Defense Center does not display and Series 2 devices do not detect data for features they do not support. See the [Supported Capabilities by Appliance Model table](#) on page 42 for a summary of Series 2 appliance features.

Custom Analysis Widget Presets

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
All Intrusion Events	Displays a graph of the total number of intrusion events on your monitored network over the dashboard time range.	Detailed Dashboard Summary Dashboard	Protection
All Intrusion Events (Not Dropped)	Displays the most frequently occurring types of intrusion events, by classification, where the packet was not dropped as part of the event.	Detailed Dashboard	Protection
Allowed Connections by Application	Displays allowed application connections on your monitored network, grouped by application.	Application Statistics	FireSIGHT
Allowed Connections by Application Risk	Displays allowed application connections on your monitored network, grouped by application risk level.	Application Statistics	FireSIGHT
Allowed Connections by Business Relevance	Displays allowed application connections on your monitored network, grouped by estimated relevance to business activity.	Application Statistics	FireSIGHT
Allowed Connections by URL Category	Displays allowed application connections on your monitored network, grouped by URL category.	URL Statistics	URL Filtering
Allowed Connections by URL Reputation	Displays allowed application connections on your monitored network, grouped by URL reputation.	URL Statistics	URL Filtering
Allowed Connections by User	Displays allowed application connections on your monitored network, grouped by connecting user.	User Statistics	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Application Protocols Introducing Malware	Displays the number of malware files transmitted over your network, grouped by the application protocol used to transmit the files.	Files Dashboard	Malware
Application Protocols Transferring Files	Displays the number of files transmitted over your network, grouped by the application protocol used to transmit the files.	Files Dashboard	Protection
Client Applications Introducing Malware	Displays the applications, or parent files, that accessed or created malware detected by FireAMP Connectors.	Files Dashboard	FireAMP subscription
Client Applications Transferring Files	Displays the applications, or parent files, that transmitted files over your network.	Files Dashboard	Protection
Clients	Displays clients on your monitored network, by type.	Detailed Dashboard	FireSIGHT
Connections by Application	Displays applications on your monitored network, based on the number of detected connections.	Connection Summary	FireSIGHT
Connections by Destination Continent	Displays continents to which connections were sent from your monitored network, based on the number of connections.	Connection Summary	FireSIGHT
Connections by Destination Country	Displays countries to which connections were sent from your monitored network, based on the number of connections.	Connection Summary	FireSIGHT
Connections by Initiator IP	Displays host IP addresses on your monitored network, based on the number of connections where that IP address on a host initiated the session.	Connection Summary	FireSIGHT
Connections by Port	Displays ports on your monitored network, based on the number of detected connections.	Connection Summary	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Connections by Responder IP	Displays host IP addresses on your monitored network, based on the number of connections where the responder in that session was that IP address on a host. The output of this widget varies according to your connection logging configuration.	Connection Summary	FireSIGHT
Connections by Source Continent	Displays continents communicating with your monitored network, based on the number of connections initiated from each continent.	Connection Summary	FireSIGHT
Connections by Source Country	Displays countries communicating with your monitored network, based on the number of connections initiated from each country.	Connection Summary	FireSIGHT
Connections by URL Category	Displays all application connections on your monitored network, grouped by URL category.	Summary Dashboard	URL Filtering
Connections by URL Reputation	Displays all application connections on your monitored network, grouped by URL reputation.	Summary Dashboard	URL Filtering
Connections over Time	Displays a graph of the total number of connections on your monitored network, over the dashboard time range.	Connection Summary	FireSIGHT
Denied Connections by Application	Displays denied connections on your monitored network, grouped by application.	Application Statistics	FireSIGHT
Denied Connections by URL Category	Displays denied connections on your monitored network, grouped by URL category.	URL Statistics	URL Filtering
Denied Connections by URL Reputation	Displays denied connections on your monitored network, grouped by URL reputation.	URL Statistics	URL Filtering
Denied Connections by User	Displays denied connections on your monitored network, grouped by connecting user.	User Statistics	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Dropped Events by Application	Displays dropped intrusion events, grouped by application.	Application Statistics	Protection + FireSIGHT
Dropped Events by User	Displays dropped intrusion events, grouped by user.	User Statistics	Protection + FireSIGHT
Dropped Intrusion Events	Displays counts for intrusion events, by classification, where the packet was dropped.	Detailed Dashboard Summary Dashboard	Protection
File Actions	Displays the number of files transmitted over your network, grouped by the file rule actions used to handle the files.	Files Dashboard	Protection or Malware
File Categories	Displays the number of files transmitted over your network, grouped by file category.	Files Dashboard	Protection
File Dispositions	Displays the number of files detected in network traffic as a result of Malware Cloud Lookup file rules, grouped by malware disposition.	Files Dashboard	Malware
File Names	Displays the number of files transmitted over your network, grouped by file name.	Files Dashboard	Protection
File Transfers over Time	Displays a graph of the total number of file transfers detected in network traffic by the system, over the dashboard time range.	Files Dashboard	Protection
File Types	Displays the number of files transmitted over your network, grouped by file type.	Files Dashboard	Protection
File Types Infected with Malware	Displays the number of malware detected either in network traffic by the system or by FireAMP Connectors, grouped by file type.	Files Dashboard	Malware
Hosts Receiving Files	Displays the number of files received (downloaded) by host IP addresses on your network, grouped by IP address.	Files Dashboard	Protection

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Hosts Receiving Malware	Displays the number of malware files received by host IP addresses on your network, grouped by IP address.	Files Dashboard	Malware license or FireAMP subscription
Hosts Sending Files	Displays the number of files sent (uploaded) from host IP addresses on your network, grouped by IP address.	Files Dashboard	Protection
Hosts Sending Malware	Displays the number of malware files sent from host IP addresses on your network, grouped by IP address.	Files Dashboard	Malware
Impact Level x Events by Application	Displays number of events of estimated impact level x (where x is a number 0-4), grouped by application.	Application Statistics	Protection + FireSIGHT
Impact Level x Events by Application Protocol	Displays number of events of estimated impact level x (where x is a number 1-2), grouped by application protocol.	Summary Dashboard	Protection + FireSIGHT
Impact Level x Events by User	Displays number of events of estimated impact level x (where x is a number 0-4), grouped by user.	User Statistics	Protection + FireSIGHT
Intrusion Events Requiring Analysis	Displays a count of intrusion events requiring analysis, based on event classification.	Detailed Dashboard	Protection + FireSIGHT
Intrusion Events by Destination Continent	Displays continents targeted by intrusion events, based on the number of events associated with each continent.	Summary Dashboard	FireSIGHT
Intrusion Events by Destination Country	Displays countries targeted by intrusion events, based on the number of events associated with each country.	Summary Dashboard	FireSIGHT
Intrusion Events by Source Continent	Displays continents where intrusion events originated, based on the number of events originated from each continent.	Summary Dashboard	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Intrusion Events by Source Country	Displays countries where intrusion events originated, based on the number of events originated from each country.	Summary Dashboard	FireSIGHT
Intrusion Events to High Criticality Hosts	Displays intrusion events, based on the number of intrusion events occurring on high criticality hosts.	Detailed Dashboard	Protection + FireSIGHT
Malware Intrusions	Displays intrusion events, based on the number of intrusion events occurring in connections transmitting malware.	Files Dashboard	Malware
Malware Threats	Displays the number of malware threats detected either in network traffic by the system or by FireAMP Connectors, grouped by threat name.	Files Dashboard	Malware license or FireAMP subscription
Operating Systems	Displays operating systems, based on the number of hosts running each operating system within your network.	Detailed Dashboard	FireSIGHT
Processes Introducing Malware	Displays the system processes that accessed or created malware detected by FireAMP Connectors.	Files Dashboard	Malware license or FireAMP subscription
Risky Applications with Low Business Relevance	Displays all application connections on your monitored network that have both high application risk level and low estimated business relevance.	Summary Dashboard	FireSIGHT
Servers	Displays servers, by number of hosts.	Detailed Dashboard	FireSIGHT
Threat Detections over Time	Displays a graph of the total number of malware threats detected either in network traffic by the system or by FireAMP Connectors, over the dashboard time range.	Files Dashboard	Malware license or FireAMP subscription
Top Attackers	Displays attacking host IP addresses on your monitored network, based on the number of intrusion events where the listed IP address was the attacker in the connection that caused the event.	Summary Dashboard	Protection

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Top Client Applications Seen	Displays client applications on your monitored network, based on total kilobytes of data transmitted by the client application.	Summary Dashboard	FireSIGHT
Top Operating Systems Seen	Displays operating systems on your monitored network, based on the number of network hosts with the operating system.	Summary Dashboard	FireSIGHT
Top Server Applications Seen	Displays server applications on your monitored network, based on the number of hosts running the service.	Summary Dashboard	FireSIGHT
Top Targets	Displays host IP addresses on your monitored network, based on the number of intrusion events where that address was targeted in the connection that caused the event.	Summary Dashboard	Protection
Top Web Applications Seen	Displays web applications on your monitored network, based on total kilobytes of data transmitted by the client application.	Summary Dashboard	FireSIGHT
Total Events by Application	Displays applications on your monitored network, based on the number of intrusion events generated by the application.	Application Statistics	Protection + FireSIGHT
Total Events by Application Protocol	Displays application protocols on your monitored network, based on the number of intrusion events associated with the application protocol.	Summary Dashboard	Protection + FireSIGHT
Total Events by User	Displays users on your monitored network, based on the number of intrusion events generated by each user's activity.	Summary Dashboard User Statistics	Protection + FireSIGHT
Traffic by Application	Displays applications on your monitored network, based on total kilobytes of data transmitted on your monitored network by the application over the dashboard time range.	Application Statistics Connection Summary Detailed Dashboard	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Traffic by Application Category	Displays application categories on your monitored network, based on total kilobytes of data transmitted on your monitored network by applications in each category over the dashboard time range.	Application Statistics Summary Dashboard	FireSIGHT
Traffic by Application Risk	Displays estimated risk levels of applications on your monitored network, based on total kilobytes of data transmitted on your monitored network by applications at each level over the dashboard time range.	Summary Dashboard	FireSIGHT
Traffic by Business Relevance	Displays estimated business relevance levels of applications on your monitored network, based on total kilobytes of data transmitted on your monitored network by applications at each level over the dashboard time range.	Summary Dashboard	FireSIGHT
Traffic by Destination Continent	Displays continents contacted from your monitored network, based on total kilobytes of data transmitted on your monitored network to each continent over the dashboard time range.	Connection Summary	FireSIGHT
Traffic by Destination Country	Displays countries contacted from your monitored network, based on total kilobytes of data transmitted on your monitored network to each country over the dashboard time range.	Connection Summary	FireSIGHT
Traffic by Initiator IP	Displays host IP addresses on your monitored network, based on total kilobytes of data transmitted on your monitored network from the IP address over the dashboard time range.	Connection Summary Detailed Dashboard	FireSIGHT
Traffic by Initiator User	Displays users on your monitored network, based on total kilobytes of data received by the hosts where those users are logged in.	Detailed Dashboard Summary Dashboard	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Traffic by Port	Displays responder ports on your monitored network, based on total kilobytes of data transmitted on your monitored network via each port over the dashboard time range. The output of this widget varies according to your connection logging configuration.	Connection Summary	FireSIGHT
Traffic by Responder IP	Displays IP addresses on your monitored network, based on total kilobytes of data received by the IP addresses (on hosts) over the dashboard time range. The output of this widget varies according to your connection logging configuration.	Connection Summary Detailed Dashboard	FireSIGHT
Traffic by Source Continent	Displays continents transmitting data to your monitored network, based on total kilobytes of data on your monitored network transmitted from each continent over the dashboard time range.	Connection Summary	FireSIGHT
Traffic by Source Country	Displays countries transmitting data to your monitored network, based on total kilobytes of data on your monitored network transmitted from each country over the dashboard time range.	Connection Summary	FireSIGHT
Traffic by URL Category	Displays application URL categories on your monitored network, based on total kilobytes of data exchanged with URLs of each category over the dashboard time range.	URL Statistics	URL Filtering
Traffic by URL Reputation	Displays application URL reputation types on your monitored network, based on total kilobytes of data exchanged with URLs of each reputation over the dashboard time range.	URL Statistics	URL Filtering
Traffic by User	Displays users on your monitored network, based on total kilobytes of data exchanged by each user over the dashboard time range.	None	FireSIGHT

Custom Analysis Widget Presets (Continued)

PRESET	DESCRIPTION	PREDEFINED DASHBOARDS	LICENSES
Traffic over Time	Displays a graph of total kilobytes of data transmitted on your monitored network over the dashboard time range.	Connection Summary Detailed Dashboard	FireSIGHT
Unique Applications over Time	Displays a graph of total unique applications detected on your monitored network over the dashboard time range.	Application Statistics Summary Dashboard	FireSIGHT
Unique Users over Time	Displays a graph of total unique users detected on your monitored network over the dashboard time range.	User Statistics	FireSIGHT
Users Affected by Malware	Displays the number of threats detected either in network traffic by the system or by FireAMP Connectors, grouped by user.	Files Dashboard	Malware + FireSIGHT, or FireAMP subscription
Users Transferring Files	Displays the number of files being transmitted over your network, grouped by sender.	Files Dashboard	Malware + FireSIGHT
Web Applications Introducing Malware	Displays web applications on your monitored network that accessed or created malware detected by FireAMP Connectors.	Files Dashboard	Malware license or FireAMP subscription
Web Applications Transferring Files	Displays the number of files transmitted over your network, grouped by the web application used to transmit the files.	Files Dashboard	Malware license or FireAMP subscription
White List Violations	Displays hosts with white list violations, by violation count.	Detailed Dashboard	FireSIGHT

Viewing Associated Events from the Custom Analysis Widget

LICENSE: Any

Depending on the kind of data that a Custom Analysis widget is configured to display, you can invoke an event view (that is, a workflow) that provides detailed information about the events displayed in the widget.

When you invoke an event view from the dashboard, the events appear in the default workflow for that event type, constrained by the dashboard time range. This also changes the appropriate time window for the appliance, depending on

how many time windows you have configured and on what type of event you are trying to view.

For example, if you configure multiple time windows on your Defense Center and then access health events from a Custom Analysis widget, the events appear in the default health events workflow, and the health monitoring time window changes to the dashboard time range.

As another example, if you configure a single time window and then access any type of event from the Custom Analysis widget, the events appear in the default workflow for that event type, and the global time window changes to the dashboard time range.

For more information on time windows, see [Default Time Windows](#) on page 2219 and [Specifying Time Constraints in Searches](#) on page 1775.

To view associated events from the Custom Analysis Widget:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ You have two options, depending on how you configured the widget:
 - On widgets configured to show relative occurrences of events (that is, bar graphs), click any event to view associated events constrained by the widget preferences, as well as by that event. You can also click the view all icon () in the lower right corner of the widget to view all associated events, constrained by the widget preferences.
 - On widgets configured to show connection data over time, click the view all icon in the lower right corner of the widget to view all associated events, constrained by the widget preferences.

For information on working with specific event types, see the following sections:

- [Working with Security Intelligence Lists and Feeds](#) on page 162
- [Viewing Audit Records](#) on page 2188
- [Viewing Intrusion Events](#) on page 602
- [Viewing Discovery and Host Input Events](#) on page 1396
- [Viewing File Events](#) on page 1223
- [Viewing Malware Events](#) on page 1232
- [Viewing Hosts](#) on page 1402
- [Viewing Host Attributes](#) on page 1412
- [Viewing Servers](#) on page 1418
- [Viewing Application Details](#) on page 1429
- [Viewing Sourcefire Vulnerabilities](#) on page 1435
- [Viewing Third-Party Vulnerabilities](#) on page 1441
- [Viewing Connection Data Graphs and Tables](#) on page 555
- [Viewing Users](#) on page 1447

- [Viewing User Activity Events](#) on page 1454
- [Viewing Correlation Events](#) on page 1521
- [Viewing White List Events](#) on page 1573
- [Viewing White List Violations](#) on page 1579
- [Viewing Health Events](#) on page 2175
- [Viewing the Rule Update Log](#) on page 2082
- [Working with Active Scan Results](#) on page 1716
- [Using Geolocation](#) on page 1816
- [Understanding Custom Tables](#) on page 1780

Custom Analysis Widget Limitations

LICENSE: Any

There are some important points to keep in mind when using the Custom Analysis widget.

If you are configuring the widget on a shared dashboard, remember that not all users can view data of all event types, depending on the user's account privileges. For example, Maintenance Users cannot view discovery events.

Similarly, if you are using a dashboard imported from another appliance, remember that not all appliances have access to data of all event types. For example, managed devices do not store correlation data. If your dashboard includes a Custom Analysis widget that displays data you cannot see, the widget indicates that you are unauthorized to view the data. Note, however, that you (and any other users who share the dashboard) can modify the preferences of the widget to display data that you can see, or even delete the widget. If you want to make sure that this does not happen, save the dashboard as private.

Remember that only you can access searches that you have saved as private. If you configure the widget on a shared dashboard and constrain its events using a private search, the widget resets to not using the search when another user logs in. This affects your view of the widget as well. If you want to make sure that this does not happen, save the dashboard as private.

You enable or disable the Custom Analysis widget from the Dashboard settings in your system policy. For more information, see [Configuring Dashboard Settings](#) on page 1977.

Understanding the Disk Usage Widget

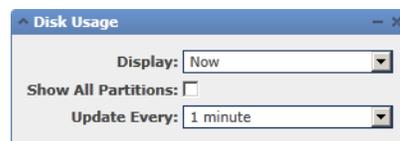
LICENSE: Any

The Disk Usage widget indicates the percentage of space used on each partition of the appliance's hard drive. It also shows the capacity of each partition. This

widget appears by default on the Status tabs of the Default Dashboard and the Summary Dashboard.



You can configure the widget to display only admin (/) and /volume partition usage, or you can show these plus the /boot partition usage by modifying the widget preferences.

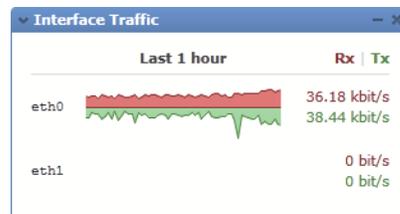


The widget preferences also control how often the widget updates, as well as whether it displays the current disk usage or collected disk usage statistics over the dashboard time range. For more information, see [Understanding Widget Preferences](#) on page 73.

Understanding the Interface Traffic Widget

LICENSE: Any

The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's enabled interfaces over the dashboard time range. It does not appear by default on any of the predefined dashboards.



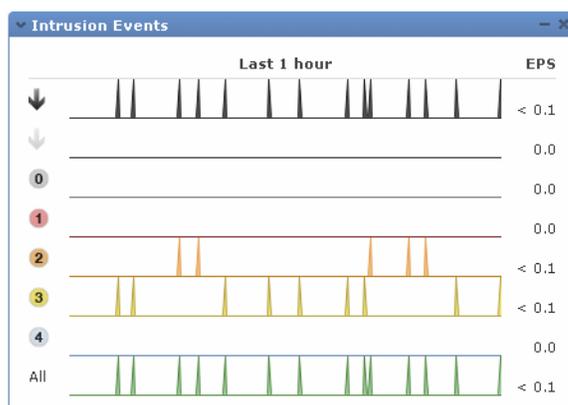
The widget preferences control how often the widget updates. On managed devices, the preferences also control whether the widget displays the traffic rate for unused interfaces (by default, the widget only displays the traffic rate for active interfaces). For more information, see [Understanding Widget Preferences](#) on page 73.

Understanding the Intrusion Events Widget

LICENSE: Protection

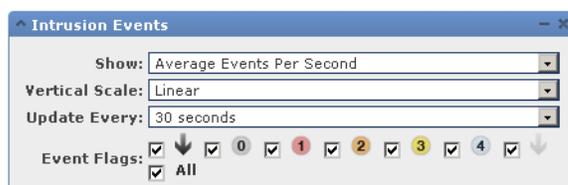
The Intrusion Events widget shows the intrusion events that occurred over the dashboard time range, organized by priority. This includes statistics on intrusion

events with dropped packets and different impacts. This widget appears by default on the Intrusion Events tab of the Summary Dashboard.



On managed devices, the widget can display statistics for dropped (or, on passively deployed devices, would have dropped) intrusion events, all intrusion events, or both. Note that you must enable local event storage or the widget will not have any data to display. Note also that the total rate represented by **All** does not include the dropped event rate.

On the Defense Center, but not managed devices, you can configure the widget to display intrusion events with dropped/would have dropped packets and different impacts by modifying the widget preferences. You can display dropped and would have dropped events on Defense Centers and devices. The following graphic shows the Defense Center version of the widget preferences.



In the widget preferences, you can:

- on a Defense Center, select one or more **Event Flags** check boxes to display separate graphs for events with dropped packets, would have dropped packets, or specific impacts; select **All** to display an additional graph for all intrusion events, regardless of impact or rule state; see [Using Impact Flags to Evaluate Events](#) on page 641 for more information
- select **Show** to choose **Average Events Per Second** or **Total Events**
- select **Vertical Scale** to choose **Linear** (incremental) or **Logarithmic** (factor of ten) scale

The preferences also control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

On the Intrusion Events widget, you can:

- on a Defense Center, click a graph corresponding to dropped packets, to would have dropped packets, or to a specific impact to view intrusion events of that type
- click the graph corresponding to dropped events to view dropped events
- click the graph corresponding to would have dropped events to view would have dropped events
- click the **All** graph to view all intrusion events

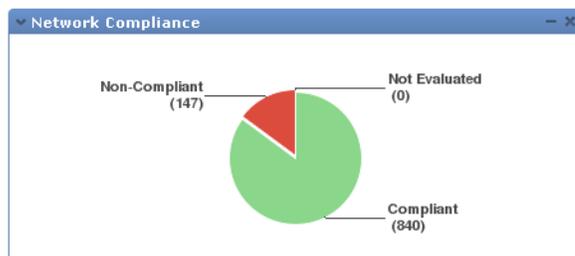
Note that the resulting event view is constrained by the dashboard time range; accessing intrusion events via the dashboard changes the events (or global) time window for the appliance. For more information on intrusion events, see [Viewing Intrusion Events](#) on page 602.

Note also that packets in a passive deployment are not dropped, regardless of the rule state or the inline drop behavior of the intrusion policy. For more information, see [Setting Rule States](#) on page 743 and [Setting Drop Behavior in an Inline Deployment](#) on page 689.

Understanding the Network Compliance Widget

LICENSE: FireSIGHT

The Network Compliance widget summarizes your hosts' compliance with the white lists you configured (see [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530). By default, the widget displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated, for all compliance white lists in active correlation policies. This widget appears by default on the Correlation tab of the Detailed Dashboard.



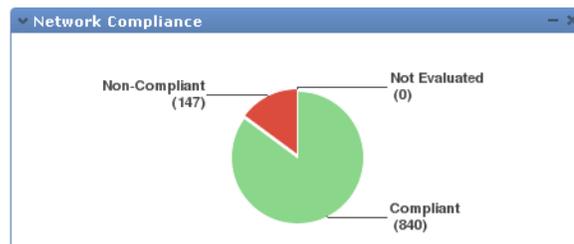
You can configure the widget to display network compliance either for all white lists or for a specific white list by modifying the widget preferences.



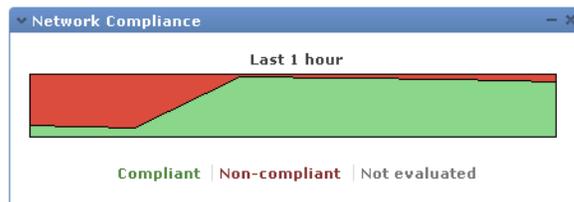
If you choose to display network compliance for all white lists, the widget considers a host to be non-compliant if it is not compliant with any white list in an active correlation policy.

You can also use the widget preferences to specify which of three different styles you want to use to display network compliance.

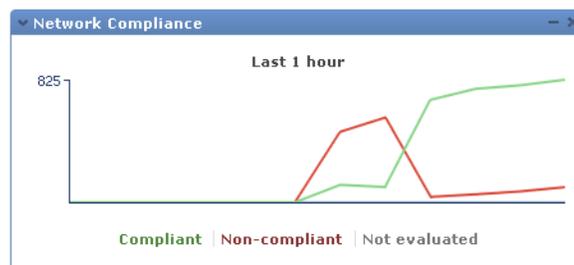
The **Network Compliance** style (the default) displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated. You can click the pie chart to view the host violation count, which lists the hosts that violate at least one white list. For more information, see [Viewing White List Violations](#) on page 1579.



The **Network Compliance over Time (%)** style displays a stacked area graph showing the relative proportion of hosts that are compliant, non-compliant, and that have not yet been evaluated, over the dashboard time range.



The **Network Compliance over Time** style displays a line graph that shows the number of hosts that are compliant, non-compliant, and that have not yet been evaluated, over the dashboard time range.



The preferences control how often the widget updates. You can check the **Show Not Evaluated** box to hide events which have not been evaluated. For more information, see [Understanding Widget Preferences](#) on page 73.

Understanding the Product Licensing Widget

LICENSE: Any

The Product Licensing widget shows the device and feature licenses currently installed on the Defense Center. It also indicates the number of items (such as hosts or users) licensed and the number of remaining licensed items allowed. It does not appear by default on any of the predefined dashboards.



License Type	Licensed	Remaining	%
3D8250 Control	100	99	99%
3D8250 Protection	100	99	99%
3D8250 URL Filtering	100	99	99%
DC3500 FireSIGHT Host	300,000	290,579	96%
DC3500 FireSIGHT User	300,000	299,998	99%

Expiring Licenses

License Type	Expires	Licensed
3D8250 URL Filtering	2012-05-19	100

The top section of the widget displays all device and feature licenses installed on the Defense Center, including temporary licenses, while the Expiring Licenses section displays only temporary and expired licenses. For example, if you have two feature licenses for FireSIGHT Hosts, one of which is a permanent license and allows 750 hosts, and another that is temporary and allows an additional 750 hosts, the top section of the widget displays a FireSIGHT Hosts feature license with 1500 licensed hosts, while the Expiring Licenses section displays a FireSIGHT Hosts feature license with 750 hosts.

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](#) on page 73.

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 Sourcefire 3D System](#) on page 2036.

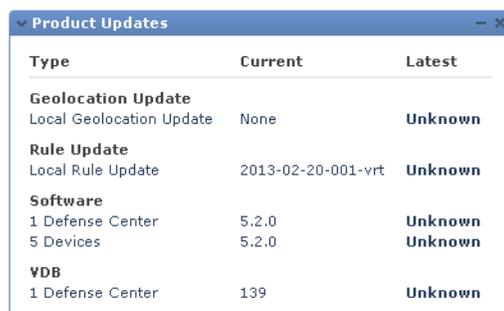
Understanding the Product Updates Widget

LICENSE: Any

The Product Updates widget provides you with a summary of the software (Sourcefire 3D System 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. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

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](#) on page 1929.

The widget also provides you with links to pages where you can update the software; the Defense Center version of the widget provides you with similar links so you can update the software on your managed devices.



Type	Current	Latest
Geolocation Update		
Local Geolocation Update	None	Unknown
Rule Update		
Local Rule Update	2013-02-20-001-vrt	Unknown
Software		
1 Defense Center	5.2.0	Unknown
5 Devices	5.2.0	Unknown
VDB		
1 Defense Center	139	Unknown

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](#) on page 73.

On the Product Updates widget, you can:

- manually update an appliance by clicking the current version of the Sourcefire 3D System software, rule update, geolocation update, or VDB:
 - to update the system software, geolocation database, or VDB, see [Updating System Software](#) on page 2054.
 - to import the newest rule update, see [Importing Rule Updates and Local Rule Files](#) on page 2072.
- create a scheduled task to download the latest version of the Sourcefire 3D System software, rule update, or VDB by clicking either the latest version or the **Unknown** link in the Latest column; see [Scheduling Tasks](#) on page 1929.

Understanding the RSS Feed Widget

LICENSE: Any

The RSS Feed widget adds an RSS feed to a dashboard. By default, the widget shows a feed of Sourcefire security news. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.



You can also configure the widget to display a preconfigured feed of Sourcefire company news, the Snort.org blog, or the Vulnerability Research Team (VRT) blog, or you can create a custom connection to any other RSS feed by specifying its URL in the widget preferences.



Feeds update every 24 hours (although you can manually update the feed), and the widget displays the last time the feed was updated based on the local time of the appliance. Keep in mind that the appliance must have access to the Sourcefire web site (for the two preconfigured feeds) or to any custom feed you configure.

When you configure the widget, you can also choose how many stories from the feed you want to show in the widget, as well as whether you want to show descriptions of the stories along with the headlines; keep in mind that not all RSS feeds use descriptions.

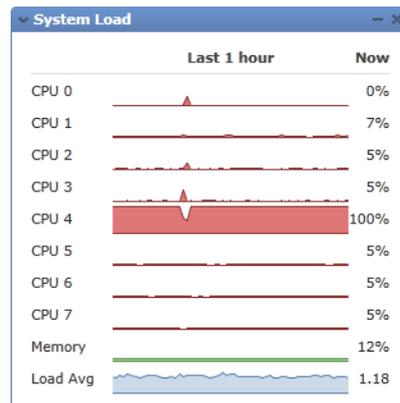
On the RSS Feed widget, you can:

- click one of the stories in the feed to view the story
- click the **more** link to go to the feed's web site
- click the update icon (🔄) to manually update the feed

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. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

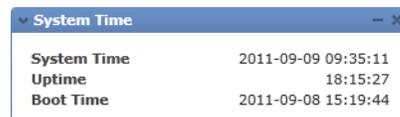


You can configure the widget to show or hide the load average by modifying the widget preferences. The preferences also control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

Understanding the System Time Widget

LICENSE: Any

The System Time widget shows the local system time, uptime, and boot time for the appliance. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

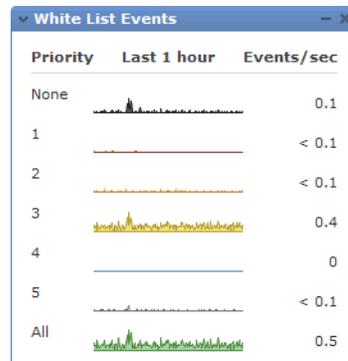


You can configure the widget to hide the boot time by modifying the widget preferences. The preferences also control how often the widget synchronizes with the appliance's clock. For more information, see [Understanding Widget Preferences](#) on page 73.

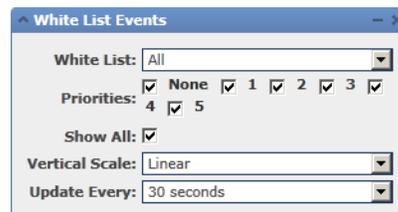
Understanding the White List Events Widget

LICENSE: FireSIGHT

The White List Events widget shows the average events per second by priority, over the dashboard time range. It appears by default on the Correlation tab of the Default Dashboard.



You can configure the widget to display white list events of different priorities by modifying the widget preferences.



In the widget preferences, you can:

- select one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority
- select **Show All** to display an additional graph for all white list events, regardless of priority
- select **Vertical Scale** to choose **Linear** (incremental) or **Logarithmic** (factor of ten) scale

The preferences also control how often the widget updates. For more information, see [Understanding Widget Preferences](#) on page 73.

You can click a graph to view white list events of a specific priority, or click the **All** graph to view all white list events. In either case, the events are constrained by the dashboard time range; accessing white list events via the dashboard changes the events (or global) time window for the Defense Center. For more information on white list events, see [Viewing White List Events](#) on page 1573.

Working with Dashboards

LICENSE: Any

You can view and modify the widgets that appear on the dashboard.

You manage dashboards on the Dashboard Management page (see [Viewing Dashboards](#) on page 109). You can create, view, modify, export, and delete dashboards.

Name	Owner	Private	Default	
Application Statistics Provides traffic and intrusion event statistics by application	admin	No	No	   
Connection Summary Provides tables and charts of the activity on your monitored network segment organized by different criteria	admin	No	No	   
Detailed Dashboard Provides a detailed view of activity on the appliance	admin	No	No	   
Files Dashboard Provides an overview of Malware and File Events	admin	No	No	   
Summary Dashboard Provides a summary of activity on the appliance	admin	No	Yes	   
URL Statistics Provides traffic statistics by URL category and reputation	admin	No	No	   
User Statistics Provides traffic and intrusion event statistics by user	admin	No	No	   

For each dashboard, the page indicates the owner (that is, the user who created it) and whether a dashboard is private. Note that, unless you have Administrator access, you can only see your own private dashboards; you cannot view or modify private dashboards created by other users.

Finally, the page indicates which dashboard is the default. You specify the default dashboard in your user preferences; for more information, see [Specifying Your Default Dashboard](#) on page 2223.

For more information on working with dashboards, see:

- [Creating a Custom Dashboard](#) on page 106
- [Viewing Dashboards](#) on page 109
- [Modifying Dashboards](#) on page 111
- [Deleting a Dashboard](#) on page 116
- [Exporting Objects](#) on page 2226

Creating a Custom Dashboard

LICENSE: Any

When you create a new dashboard, you can choose to base it on any existing dashboard, whether user-created or predefined by Sourcefire. This makes a copy of the preexisting dashboard; you can modify this copy to suit your needs. Optionally, you can create a blank new dashboard by choosing not to base your dashboard on any preexisting dashboards.

You must also specify (or disable) the tab change and page refresh intervals. These settings determine how often the dashboard cycles through its tabs and how often the entire dashboard page refreshes.

Refreshing the entire dashboard allows you to see any preference or layout changes that were made to a shared dashboard by another user, or that you made to a private dashboard on another computer, since the last time the dashboard refreshed. This may be useful, for example, in a network operations center (NOC) where a dashboard is displayed at all times. If you want to make changes to the dashboard, you can make the changes at a local computer. Then, the dashboard in the NOC automatically refreshes at the interval you specify and displays your changes without you having to manually refresh the dashboard in the NOC. Note that you do not need to refresh the entire dashboard to see data updates; individual widgets update according to their preferences.

Finally, you can choose to associate the new dashboard with your user account by saving it as a private dashboard. If you choose not to save the dashboard as private, all other users of the appliance can view it.

Keep in mind that because not all user roles have access to all dashboard widgets, users with fewer permissions viewing a dashboard created by a user with more permissions may not be able to use all of the widgets on the dashboard. Although the unauthorized widgets still appear on the dashboard, they are disabled.

You should also keep in mind that any user with dashboard access, regardless of role, can modify shared dashboards. If you want to make sure that only you can modify a particular dashboard, save it as private.

TIP! Instead of creating a new dashboard, you can export a dashboard from another appliance, then import it onto your appliance. You can then edit the imported dashboard to suit your needs. Note that the dashboard widgets you can view depend on the type of appliance you are using and on your user role; for example, a dashboard created on the Defense Center and imported onto a managed device may display some invalid, disabled widgets. For more information, see [Importing and Exporting Objects](#) on page 2225.

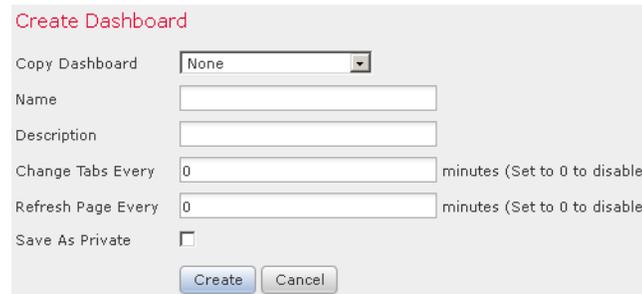
To create a new dashboard:

ACCESS: Admin/Any Security Analyst/Maint

1. Select **Overview > Dashboards > Management**.
The Dashboard Management page appears.

2. Click **Create Dashboard**.

The Create Dashboard page appears.



3. Use the **Copy Dashboard** drop-down list to select the dashboard on which you want to base the new dashboard.

You can select any predefined or user-defined dashboard. Optionally, select **None** (the default) to create a blank dashboard.

4. Type a name and optional description for the dashboard.

5. In the **Change Tabs Every** field, specify (in minutes) how often the dashboard should change tabs.

Unless you pause the dashboard or your dashboard has only one tab, this setting advances your view to the next tab at the interval you specify. To disable tab cycling, enter 0 in the **Change Tabs Every** field.

6. In the **Refresh Page Every** field, specify (in minutes) how often the current dashboard tab should refresh with new data. This value must be greater than the **Change Tabs Every** setting.

Unless you pause the dashboard, this setting will refresh the entire dashboard at the interval you specify. To disable the periodic page refresh, enter 0 in the **Refresh Page Every** field.

Note that this setting is separate from the update interval available on many individual widgets; although refreshing the dashboard page resets the update interval on individual widgets, widgets will update according to their individual preferences even if you disable the **Refresh Page Every** setting.

7. Optionally, select the **Save As Private** check box to associate the dashboard with your user account and to prevent other users from viewing and modifying the dashboard.

8. Click **Save**.

Your dashboard is created and appears in the web interface. You can now tailor it to suit your needs by adding tabs and widgets (and, if you based it on a preexisting dashboard, by rearranging and deleting widgets). For more information, see [Modifying Dashboards](#) on page 111.

Viewing Dashboards

LICENSE: Any

By default, the home page for your appliance displays the default dashboard. If you do not have a default dashboard defined, the home page shows the Dashboard Management page, where you can choose a dashboard to view. At any time, to view the default dashboard you have configured for your appliance, select **Overview > Dashboards**; to view details of all available dashboards, select **Overview > Dashboards > Management**.

TIP! You can configure your appliance to display a different default home page, including pages that are not dashboard pages. You can also change the default dashboard. For more information, see [Specifying Your Home Page](#) on page 2217 and [Specifying Your Default Dashboard](#) on page 2223.

Each 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 Sourcefire 3D System software.

Keep in mind that for enterprise deployments of the Sourcefire 3D System, changing the time range to a long period may not be useful for widgets like the Custom Analysis widget, depending on how often newer events replace older events.

You can also pause a dashboard, which allows you to examine the data provided by the widgets without the display changing and interrupting your analysis. Pausing a dashboard has the following effects:

- Individual widgets stop updating, regardless of any **Update Every** widget preference.
- Dashboard tabs stop cycling, regardless of the **Cycle Tabs Every** setting in the dashboard properties.
- Dashboard pages stop refreshing, regardless of the **Refresh Page Every** setting in the dashboard properties.
- Changing the time range has no effect.

When you are finished with your analysis, you can unpause the dashboard. Unpausing the dashboard causes all appropriate widgets on the page to update to reflect the current time range. In addition, dashboard tabs resume cycling and the dashboard page resumes refreshing according to the settings you specified in the dashboard properties.

If you experience connectivity problems or other issues that interrupt the flow of system information to the dashboard, the dashboard automatically pauses and an error notice appears until the problem is resolved.

IMPORTANT! Your session normally logs you out after 1 hour of inactivity (or another configured interval), regardless of whether the dashboard is paused. If you plan to passively monitor the dashboard for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings. For more information, see [Managing User Login Settings](#) on page 1903 and [Configuring User Interface Settings](#) on page 1994.

To view a dashboard:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ Select **Overview > Dashboards**. You have two options, depending on whether you have a default dashboard defined:
 - If you have a default dashboard defined, it appears. To view a different dashboard, use the **Overview > Dashboards** menu.
 - If you do not have a default dashboard defined, the Dashboard Management page appears. Click **View** next to the dashboard you want to view.

The dashboard you selected appears.

To change the dashboard time range:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ From the **Show the Last** drop-down list, choose a dashboard time range. Unless the dashboard is paused, all appropriate widgets on the page update to reflect the new time range.

To pause the dashboard:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ On the time range control, click the pause icon (⏸). The dashboard is paused until you unpauses it.

To unpauses the dashboard:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ On the time range control of a paused dashboard, click the play icon (▶). The dashboard is unpaused.

Modifying Dashboards

LICENSE: Any

A dashboard has one or more tabs. You can add, delete, and rename tabs. Note that you cannot change the order of dashboard tabs.

Each tab can display one or more widgets in a three-column layout. You can minimize and maximize widgets, add and remove widgets from tabs, as well as rearrange the widgets on a tab.

You can also change the basic dashboard properties, which include its name and description, the tab cycle and page refresh intervals, and whether you want to share the dashboard with other users.

Note that any user with dashboard access, regardless of role, can modify shared dashboards. If you want to make sure that only you can modify a particular dashboard, make sure to set it as a private dashboard in the dashboard properties.

Every configuration of the Custom Analysis widget in the Sourcefire predefined dashboards corresponds to a preset for that widget. If you change or delete one of these widgets, you can restore it by creating a new Custom Analysis widget based on the appropriate preset. For more information, see

TIP! Every configuration of the Custom Analysis widget in the Sourcefire predefined dashboards corresponds to a system preset for that widget. If you change or delete one of these widgets, you can restore it by creating a new Custom Analysis widget based on the appropriate preset. For more information, see [Configuring the Custom Analysis Widget](#) on page 82.

For more information, see the following sections:

- [Changing Dashboard Properties](#) on page 111
- [Adding Tabs](#) on page 112
- [Deleting Tabs](#) on page 113
- [Renaming Tabs](#) on page 113
- [Adding Widgets](#) on page 113
- [Rearranging Widgets](#) on page 115
- [Minimizing and Maximizing Widgets](#) on page 115
- [Deleting Widgets](#) on page 116

Changing Dashboard Properties

LICENSE: Any

Use the following procedure to change the basic dashboard properties, which include its name and description, the tab cycle and page refresh intervals, and whether you want to share the dashboard with other users.

To change a dashboard's properties:

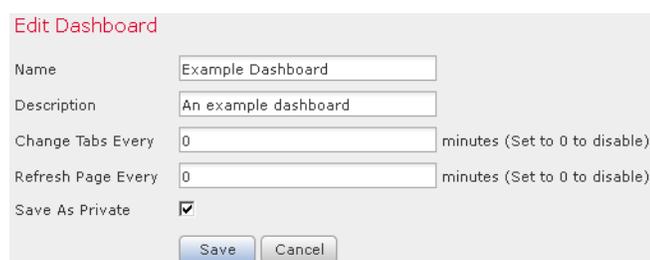
ACCESS: Admin/Any Security Analyst/Maint

1. Select **Overview > Dashboards > Management**.

The Dashboard Management page appears.

2. Click the edit icon (✎) next to the dashboard whose properties you want to change.

The Edit Dashboard page appears. See [Creating a Custom Dashboard](#) on page 106 for information on the various configurations you can change.



3. Make changes as needed and click **Save**.

The dashboard is changed.

Adding Tabs

LICENSE: Any

Use the following procedure to add a tab to a dashboard.

To add a tab to a dashboard:

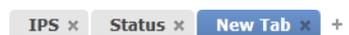
ACCESS: Admin/Any Security Analyst/Maint

1. View the dashboard where you want to add a tab.
For more information, see [Viewing Dashboards](#) on page 109.

2. To the right of the existing tabs, click the add tab icon (+).
A pop-up window appears, prompting you to name the tab.

3. Type a name for the tab (maximum 25 characters) and click **OK**, or simply click **OK** to accept the default name. Note that you can rename the tab at any time; see [Renaming Tabs](#) on page 113.

The new tab is added.



You can now add widgets to the new tab. For more information, see [Adding Widgets](#) on page 113.

Deleting Tabs

LICENSE: Any

Use the following procedure to delete a dashboard tab and all its widgets. You cannot delete the last tab from a dashboard; each dashboard must have at least one tab.

To delete a tab from a dashboard:

ACCESS: Admin/Any Security Analyst/Maint

1. View the dashboard where you want to delete a tab.
For more information, see [Viewing Dashboards](#) on page 109.
2. On the tab you want to delete, click the delete icon (✕).
3. Confirm that you want to delete the tab.
The tab is deleted.

Renaming Tabs

LICENSE: Any

Use the following procedure to rename a dashboard tab.

To rename a tab:

ACCESS: Admin/Any Security Analyst/Maint

1. View the dashboard where you want to rename a tab.
For more information, see [Viewing Dashboards](#) on page 109.
2. Click the tab you want to rename.
3. Click the tab title.
A pop-up window appears, prompting you to rename the tab.
4. Type a name for the tab (maximum 25 characters) and click **OK**.
The tab is renamed.

Adding Widgets

LICENSE: Any

To add a widget to a dashboard, you must first decide to which tab you want to add the widget. When you add a widget to a tab, the appliance automatically adds it to the column with the fewest widgets. If all columns have an equal number of

widgets, the new widget is added to the leftmost column. You can add a maximum of 15 widgets to a dashboard tab.

TIP! After you add widgets, you can move them to any location on the tab. You cannot, however, move widgets from tab to tab. For more information, see [Rearranging Widgets](#) on page 115.

To add a widget to a dashboard:

ACCESS: Admin/Any Security Analyst/Maint

1. View the dashboard where you want to add a widget.
For more information, see [Viewing Dashboards](#) on page 109.
2. Select the tab where you want to add the widget.
3. Click **Add Widgets**.

The Add Widgets page appears.

The screenshot shows the 'Add Widgets' interface. On the left, a 'Categories' sidebar lists 'All Categories (16)', 'Analysis & Reporting (5)', 'Miscellaneous (1)', and 'Operations (10)'. On the right, three widget cards are displayed:

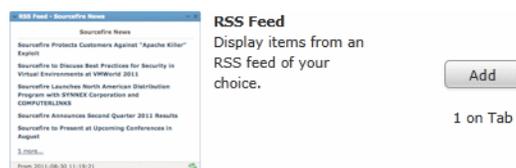
- Appliance Information:** This widget displays local appliance information including software versions, Remote Management, and High Availability status. It includes a table with fields like Name, IP# Address, Model, Versions (Software, S/N, VDB), and their values. An 'Add' button and '1 on Tab' are shown.
- Appliance Status:** This widget displays the current Health Monitoring appliance status. It features a pie chart showing 'Normal (1)' in green and 'Warning (1)' in yellow. An 'Add' button and '1 on Tab' are shown.
- Correlation Events:** This widget displays Correlation events. It shows a line graph with 'Priority' on the y-axis (None, 1, 2, 3, 4, 5, All) and 'Events/sec' on the x-axis. An 'Add' button and '1 on Tab' are shown.

The widgets that you can add depend on the type of appliance you are using and on your user role. They are organized according to function: Analysis & Reporting, Miscellaneous, and Operations. You can view the widgets in each category by clicking on the category name, or you can view all widgets by clicking **All Categories**.

4. Click **Add** next to the widgets you want to add.

TIP! To add multiple widgets of the same type (for example, you may want to add multiple RSS Feed widgets, or multiple Custom Analysis widgets), click **Add** again.

The widget is immediately added to the dashboard. The Add Widgets page indicates how many widgets of each type are on the tab, including the widget you just added.



5. Optionally, when you are finished adding widgets, click **Done** to return to the dashboard.

The tab where you added the widgets appears again, reflecting the changes you made.

Rearranging Widgets

LICENSE: Any

You can change the location of any widget on a tab. Note, however, that you cannot move widgets from tab to tab. If you want a widget to appear on a different tab, you must delete it from the existing tab and add it to the new tab.

To move a widget:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ 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:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ Click the minimize icon () in a widget's title bar.

To maximize a widget:

ACCESS: Admin/Any Security Analyst/Maint

- ▶ Click the maximize icon () in a minimized widget's title bar.

Deleting Widgets

LICENSE: Any

Delete a widget if you no longer want to view it on a tab.

To delete a widget:

ACCESS: Admin/Any Security Analyst/Maint

1. Click the close icon () in the title bar of the widget.
2. Confirm that you want to delete the widget.
The widget is deleted from the tab.

Deleting a Dashboard

LICENSE: Any

Delete a dashboard if you no longer need to use it.

If you delete your default dashboard, you must define a new default or the appliance will force you to select a dashboard to view every time you attempt to view a dashboard. For more information, see [Specifying Your Default Dashboard](#) on page 2223.

To delete a dashboard:

ACCESS: Admin/Any Security Analyst/Maint

1. Select **Overview > Dashboards > Management**.
The Dashboard Management page appears.
2. Click the delete icon () next to the dashboard you want to delete.
3. Confirm that you want to delete the dashboard.
The dashboard is deleted.

CHAPTER 3

USING THE CONTEXT EXPLORER

The Sourcefire 3D System Context Explorer displays detailed, interactive graphical information in context about the status of your monitored network, including data on applications, application statistics, connections, geolocation, intrusion events, hosts, servers, users, files (including malware files), and relevant URLs. Up to six distinct sections present this data in the form of vivid line, bar, pie, and donut graphs, accompanied by detailed lists.

You can easily create and apply custom filters to fine-tune your analysis, and you can examine data sections in more detail by simply clicking or hovering your cursor over graph areas. You can also configure the explorer's time range to reflect a period as short as the last hour or as long as the last year. Only users with the Administrator, Security Analyst, or Security Analyst (Read Only) user roles have access to the Context Explorer.

The Sourcefire 3D System dashboard is highly customizable and compartmentalized and updates in real time. In contrast, the Context Explorer is manually updated, designed to provide broader context for its data, and has a single, consistent layout designed for active user exploration.

You use the dashboard to monitor real-time activity on your network and appliances according to your own specific needs. Conversely, you use the Context Explorer to investigate a predefined set of recent FireSIGHT data in granular detail and clear context: for example, if you notice that only 15% of hosts on your network use Linux, but account for almost all YouTube traffic, you can quickly apply filters to view data only for Linux hosts, only for YouTube-associated application data, or both. Unlike the compact, narrowly focused dashboard widgets, the Context Explorer sections are designed to provide striking visual representations of system activity in a format useful to both expert and casual users of the Sourcefire 3D System.

Note that the data displayed depends on such factors as how you license and deploy your managed devices, whether you configure features that provide the data and, in the case of Series 2 appliances, whether the appliance supports a feature that provides the data. For example, because neither the DC500 Defense Center nor Series 2 devices support URL filtering by category and reputation, the DC500 Defense Center does not display data for this feature and Series 2 devices do not detect this data.

The following table summarizes some of the key differences between the dashboard and the Context Explorer.

Comparison: Dashboard and Context Explorer

FEATURE	DASHBOARD	CONTEXT EXPLORER
Displayable data	Anything monitored by the Sourcefire 3D System	Applications, application statistics, geolocation, intrusion events, files (including malware files), hosts, servers, users, and URLs
Customizability	<ul style="list-style-type: none"> Selection of widgets for a dashboard is customizable Individual widgets can be customized to varying degrees 	<ul style="list-style-type: none"> Cannot change base layout Applied filters appear in explorer URL and can be bookmarked for later use
Data update frequency	Automatic (default); user-configured	Manual
Data filtering	Possible for some widgets (must edit widget preferences)	Possible for all parts of the explorer, with support for multiple filters
Graphical context	Some widgets (particularly Custom Analysis) can display data in graph form	Extensive graphical context for all data, including uniquely detailed donut graphs
Links to relevant web interface pages	In some widgets	In every section
Time range of displayed data	User-configured	User-configured

For more information on the related Sourcefire 3D System dashboard, see [Using Dashboards](#) on page 66.

Understanding the Context Explorer

LICENSE: FireSIGHT

The Context Explorer comprises seven distinct sections that together offer a complete overview of FireSIGHT data on your monitored network. The first section, a line chart of traffic and event counts over time, provides an at-a-glance picture of recent trends in your network's activity.

The other sections are sets of interactive graphs and lists that provide greater detail for network, application, intrusion, file, geolocation, and URL data. Except for the traffic and events time graph, you can view or hide any section. You can also apply filters to constrain the data that appears in all sections; see [Working with Filters in the Context Explorer](#) on page 150 for more information.

For in-depth information on the content and function of Context Explorer sections, see the following topics:

- [Understanding the Traffic and Intrusion Event Counts Time Graph](#) on page 119
- [Understanding the Network Information Section](#) on page 120
- [Understanding the Application Information Section](#) on page 126
- [Understanding the Intrusion Information Section](#) on page 131
- [Understanding the Files Information Section](#) on page 135
- [Understanding the Geolocation Information Section](#) on page 140
- [Understanding the URL Information Section](#) on page 143

For information on how to configure the Context Explorer as a whole, see the following topics:

- [Refreshing the Context Explorer](#) on page 146
- [Setting the Context Explorer Time Range](#) on page 147
- [Minimizing and Maximizing Context Explorer Sections](#) on page 147
- [Drilling Down on Context Explorer Data](#) on page 148

For information on configuring and using Context Explorer filters, see the following topics:

- [Working with Filters in the Context Explorer](#) on page 150
- [Adding and Applying Filters](#) on page 151
- [Creating Filters with the Context Menu](#) on page 155
- [Bookmarking Filters](#) on page 156

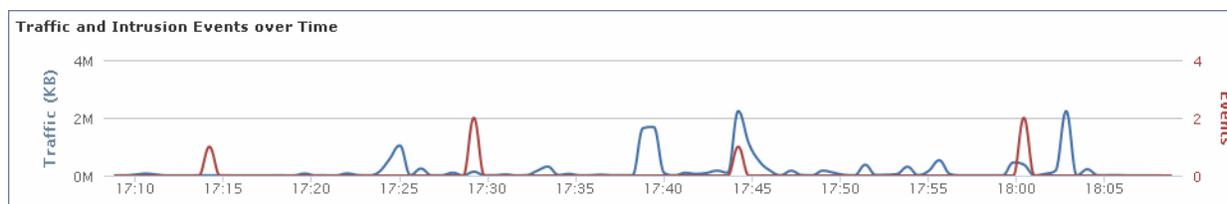
Understanding the Traffic and Intrusion Event Counts Time Graph

LICENSE: FireSIGHT

At the top of the Context Explorer is a line chart of traffic and intrusion events over time. The X-axis plots time intervals (which range from five minutes to one

month, depending on the selected time window). The Yaxis plots traffic in kilobytes (blue line) and intrusion event count (red line).

Note that the smallest X-axis interval is five minutes. To accommodate this, the system will round the beginning and ending points in your selected time range down to the nearest five-minute interval.



By default, this section shows all network traffic and all generated intrusion events for the selected time range. If you apply filters, the chart changes to display only traffic and intrusion events associated with the criteria specified in the filters. For example, filtering on the **OS Name** of **windows** causes the time graph to display only traffic and events associated with hosts using Windows operating systems.

If you filter the Context Explorer on intrusion event data (such as a **Priority of High**), the blue Traffic line is hidden to allow greater focus on intrusion events alone.

You can hover your pointer over any point on the graph lines to view exact information about traffic and event counts. Hovering your pointer over one of the colored lines also brings that line to the forefront of the graph, providing clearer context.



This section draws data primarily from the Intrusion Events and Connection Events tables.

Understanding the Network Information Section

LICENSE: FireSIGHT

The Network Information section of the Context Explorer contains six interactive graphs that display an overall picture of connection traffic on your monitored network: sources, destinations, users, and security zones associated with traffic, a breakdown of operating systems used by hosts on the network, as well as a

proportional view of access control actions your Sourcefire 3D System has performed on network traffic.

For more information on the graphs in the Network Information section, see the following topics:

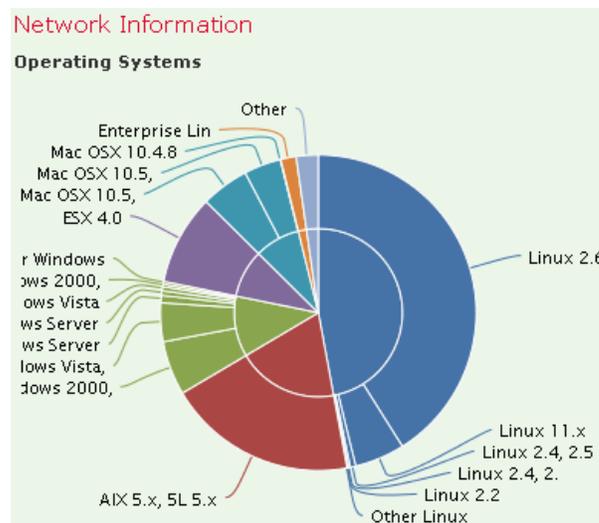
- [Viewing the Operating Systems Graph](#) on page 121
- [Viewing the Traffic by Source IP Graph](#) on page 122
- [Viewing the Traffic by Source User Graph](#) on page 122
- [Viewing the Connections by Access Control Action Graph](#) on page 124
- [Viewing the Traffic by Destination IP Graph](#) on page 124
- [Viewing the Traffic by Ingress/Egress Security Zone Graph](#) on page 125

Viewing the Operating Systems Graph

LICENSE: FireSIGHT

The Operating Systems graph, in donut form, displays a proportional representation of operating systems detected on hosts on your monitored network. The inner ring divides by OS name (such as **windows** or **Linux**), while the outer ring further divides that data by specific operating system version (such as **windows Server 2008** or **Linux 11.x**). Some closely related operating systems (such as Windows 2000, Windows XP, and Windows Server 2003) are grouped together. Very scarce or unrecognized operating systems are grouped under **Other**.

Note that this graph reflects all available data regardless of date and time constraints. If you change the explorer time range, the graph does not change.



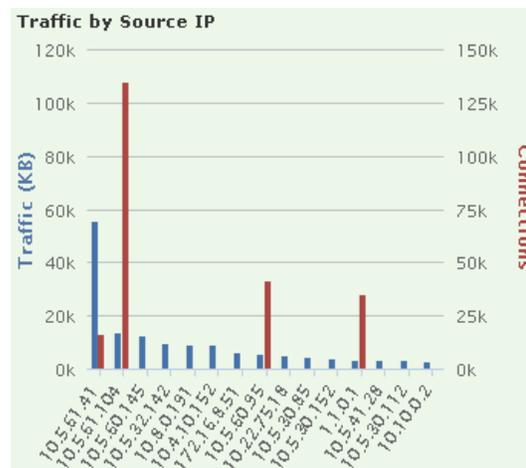
Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts table.

Viewing the Traffic by Source IP Graph

LICENSE: FireSIGHT

The Traffic by Source IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source IP addresses on your monitored network. For each source IP listed, blue bars represent traffic data and red bars represent connection data.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

IMPORTANT! If you filter on intrusion event information, the Traffic by Source IP graph is hidden.

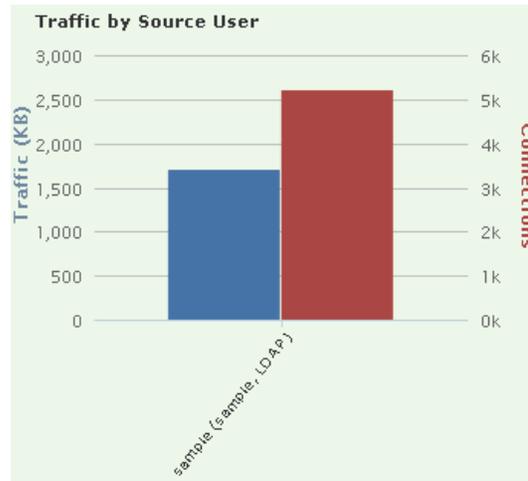
This graph draws data primarily from the Connection Events table.

Viewing the Traffic by Source User Graph

LICENSE: FireSIGHT

The Traffic by Source User graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source

users on your monitored network. For each source IP listed, blue bars represent traffic data and red bars represent connection data.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

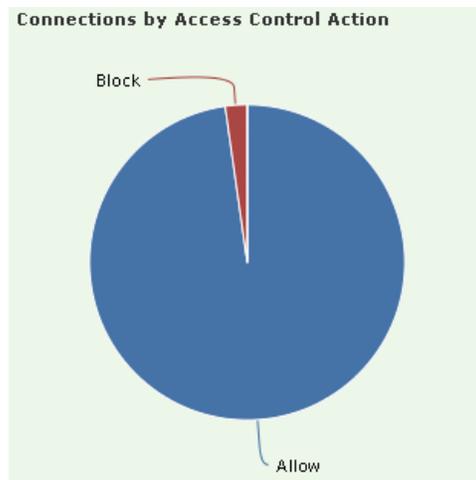
IMPORTANT! If you filter on intrusion event information, the Traffic by Source User graph is hidden.

This graph draws data primarily from the Connection Events table. Note that it displays only users reported by the Sourcefire User Agent.

Viewing the Connections by Access Control Action Graph

LICENSE: FireSIGHT

The Connections by Access Control Action graph, in pie form, displays a proportional view of access control actions (such as **B**lock or **A**llow) that your Sourcefire 3D System deployment has taken on monitored traffic.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

IMPORTANT! If you filter on intrusion event information, the Traffic by Source User graph is hidden.

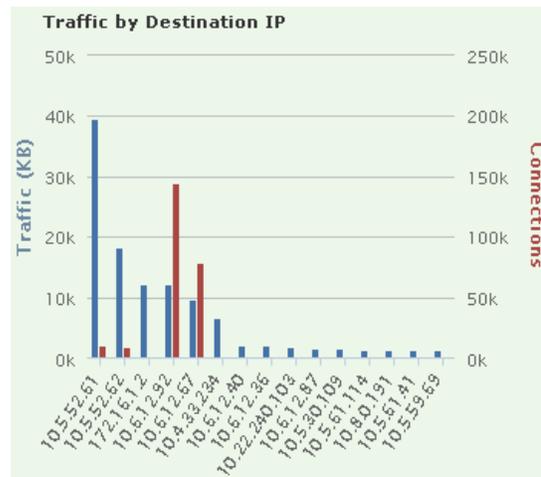
This graph draws data primarily from the Connection Events table.

Viewing the Traffic by Destination IP Graph

LICENSE: FireSIGHT

The Traffic by Destination IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active

destination IP addresses on your monitored network. For each destination IP listed, blue bars represent traffic data and red bars represent connection data.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

IMPORTANT! If you filter on intrusion event information, the Traffic by Destination IP graph is hidden.

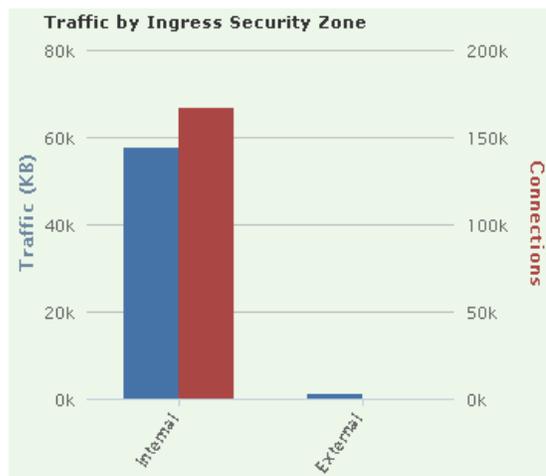
This graph draws data primarily from the Connection Events table.

Viewing the Traffic by Ingress/Egress Security Zone Graph

LICENSE: FireSIGHT

The Traffic by Ingress/Egress Security Zone graph, in bar form, displays counts of incoming or outgoing network traffic (in kilobytes per second) and unique connections for each security zone configured on your monitored network. You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

For each security zone listed, blue bars represent traffic data and red bars represent connection data. For information about security zones, see [Working with Security Zones](#) on page 187.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information

TIP! To constrain the graph so it displays only traffic by egress security zone, hover your pointer over the graph, then click **Egress** on the toggle button that appears. Click **Ingress** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Ingress view.

IMPORTANT! If you filter on intrusion event information, the Traffic by Ingress/Egress Security Zone graph is hidden.

This graph draws data primarily from the Connection Events table.

Understanding the Application Information Section

LICENSE: FireSIGHT

The Application Information section of the Context Explorer contains three interactive graphs and one table-format list that display an overall picture of application activity on your monitored network: traffic, intrusion events, and hosts associated with applications, further organized by the estimated risk or business relevance assigned to each application. The Application Details list provides an interactive list of each application and its risk, business relevance, category, and host count.

For all instances of "application" in this section, the Application Information graph set, by default, specifically examines application protocols (such as DNS or SSH).

You can also configure the Application Information section to specifically examine client applications (such as PuTTY or Firefox) or web applications (such as Facebook or Pandora).

For more information on the graphs and list in the Application Information section, see the following topics:

- [Viewing the Traffic by Risk/Business Relevance and Application Graph](#) on page 127
- [Viewing the Intrusion Events by Risk/Business Relevance and Application Graph](#) on page 128
- [Viewing the Hosts by Risk/Business Relevance and Application Graph](#) on page 129
- [Viewing the Application Details List](#) on page 130

To configure the **Application Information** section focus:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Context Explorer**.

The Context Explorer appears.

2. Hover your pointer over the **Application Protocol Information** section. (Note that if you previously changed this setting in the same Context Explorer session, the section title may appear as **Client Application Information** or **Web Application Information** instead.)

The section option buttons appear at the upper right.



3. Click **Application Protocol**, **Client Application**, or **Web Application**.

The Application Information section refreshes according to the option you selected.

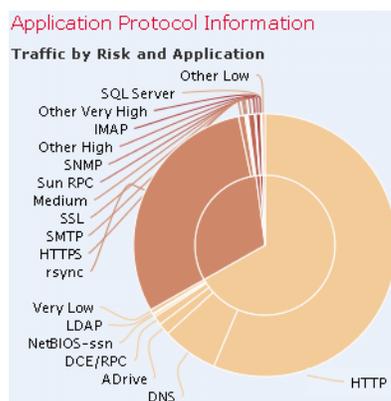
IMPORTANT! If you navigate away from the Context Explorer, this section reverts to its default state (Application Protocol).

Viewing the Traffic by Risk/Business Relevance and Application Graph

LICENSE: FireSIGHT

The Traffic by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of application traffic detected on your monitored network, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as **Medium** or **High**), while the outer ring further divides that data by specific application (such as **SSH** or **NetBIOS**). Scarcely detected applications are grouped under **Other**.

Note that this graph reflects all available data regardless of date and time constraints. If you change the explorer time range, the graph does not change.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays traffic by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

IMPORTANT! If you filter on intrusion event information, the Traffic by Risk/Business and Application graph is hidden.

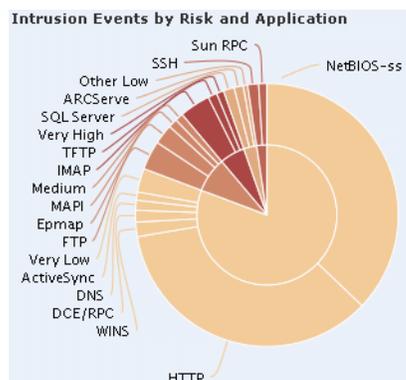
This graph draws data primarily from the Connection Events and Application Statistics tables.

Viewing the Intrusion Events by Risk/Business Relevance and Application Graph

LICENSE: FireSIGHT

The Intrusion Events by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of intrusion events detected on your monitored network and the applications associated with those events, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as **Medium**

or **High**), while the outer ring further divides that data by specific application (such as **SSH** or **NetBIOS**). Scarcely detected applications are grouped under **Other**.



Hover your pointer over any part of the donut graph to view more detailed information. Click any part of the graph to filter or drill down on that information, or (where applicable) to view application information.

TIP! To constrain the graph so it displays intrusion events by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

IMPORTANT! If you filter on URL information, the Intrusion Events by Risk/Business and Application graph is hidden.

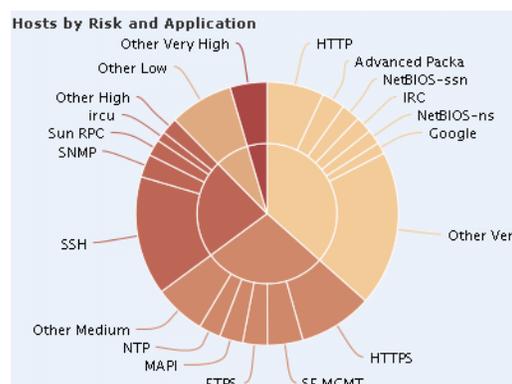
This graph draws data primarily from the Intrusion Events and Application Statistics tables.

Viewing the Hosts by Risk/Business Relevance and Application Graph

LICENSE: FireSIGHT

The Hosts by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of hosts detected on your monitored network and the applications associated with those hosts, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as **Medium** or

High), while the outer ring further divides that data by specific application (such as SSH or NetBIOS). Very scarce applications are grouped under **Other**.



Hover your pointer over any part of the donut graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays hosts by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

This graph draws data primarily from the Applications table.

Viewing the Application Details List

LICENSE: FireSIGHT

At the bottom of the Application Information section is the Application Details List, a table that provides estimated risk, estimated business relevance, category, and hosts count information for each application detected on your monitored network. The applications are listed in descending order of associated host count.

Application Details				
Application	Risk	Business Relevance	Category	Hosts
NetBIOS-ns	Very Low	High	network protocols/services	14
MySQL	Very Low	High	database	13
DNS	High	Medium	network protocols/services	11
Google Analytics	Very Low	High	web services provider	11
Shockwave	Low	Medium	multimedia (TV/video)	11
SMTP	Medium	High	email	11

The Application Details List table is not sortable, but you can click on any table entry to filter or drill down on that information, or (where applicable) to view

application information. This table draws data primarily from the Applications table.

Note that this list reflects all available data regardless of date and time constraints. If you change the explorer time range, the list does not change.

Understanding the Intrusion Information Section

LICENSE: Protection

The Intrusion Information section of the Context Explorer contains six interactive graphs and one table-format list that display an overall picture of intrusion events on your monitored network: impact levels, attack sources, target destinations, users, priority levels, and security zones associated with intrusion events, as well as a detailed list of intrusion event classifications, priorities, and counts.

IMPORTANT! If you filter on URL information, the entire Intrusion Information Section is hidden.

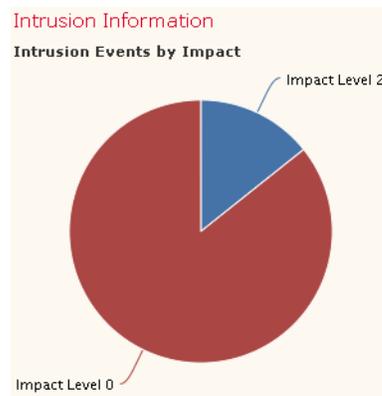
For more information on the graphs and list in the Network Information section, see the following topics:

- [Viewing the Intrusion Events by Impact Graph](#) on page 132
- [Viewing the Top Attackers Graph](#) on page 132
- [Viewing the Top Users Graph](#) on page 133
- [Viewing the Intrusion Events by Priority Graph](#) on page 133
- [Viewing the Top Targets Graph](#) on page 134
- [Viewing the Top Ingress/Egress Security Zones Graph](#) on page 134
- [Viewing the Intrusion Event Details List](#) on page 135

Viewing the Intrusion Events by Impact Graph

LICENSE: Protection

The Intrusion Events by Impact graph, in pie form, displays a proportional view of intrusion events on your monitored network, grouped by estimated impact level (from 0 to 4).



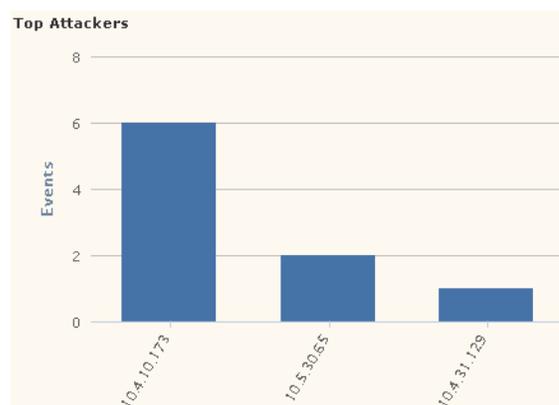
Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the IDS Statistics and Intrusion Events tables.

Viewing the Top Attackers Graph

LICENSE: Protection

The Top Attackers graph, in bar form, displays counts of intrusion events for the top attacking host IP addresses (causing those events) on your monitored network.



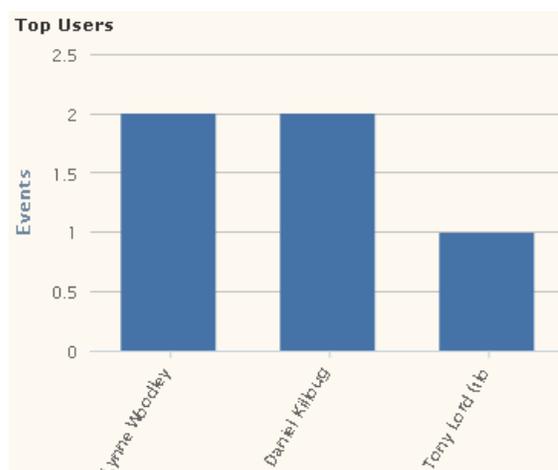
Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

Viewing the Top Users Graph

LICENSE: Protection

The Top Users graph, in bar form, displays users on your monitored network that are associated with the highest intrusion event counts, by event count.



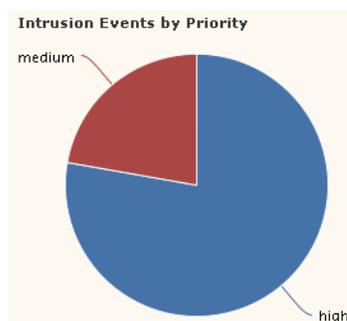
Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the IDS User Statistics and Intrusion Events tables. Note that it displays only users reported by the Sourcefire User Agent.

Viewing the Intrusion Events by Priority Graph

LICENSE: Protection

The Intrusion Events by Priority graph, in pie form, displays a proportional view of intrusion events on your monitored network, grouped by estimated priority level (such as **H**igh, **M**edium, or **L**ow).



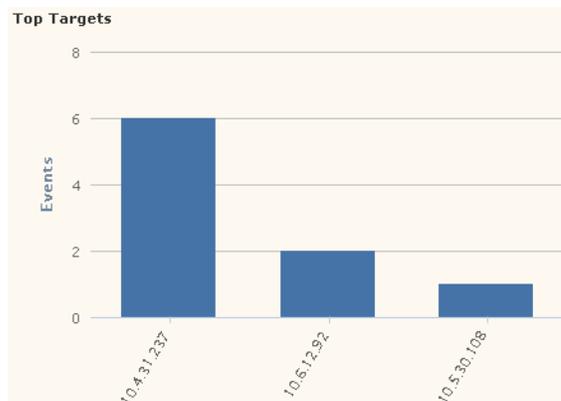
Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

Viewing the Top Targets Graph

LICENSE: Protection

The Top Targets graph, in bar form, displays counts of intrusion events for the top target host IP addresses (targeted in the connections causing those events) on your monitored network.



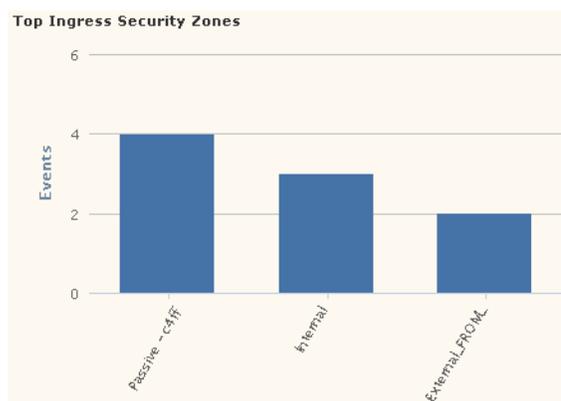
Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

Viewing the Top Ingress/Egress Security Zones Graph

LICENSE: Protection

The Top Ingress/Egress Security Zones graph, in bar form, displays counts of intrusion events associated with each security zone (ingress or egress, depending on graph settings) configured on your monitored network. For information about security zones, see [Working with Security Zones](#) on page 187.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays only traffic by egress security zone, hover your pointer over the graph, then click **Egress** on the toggle button that appears. Click **Ingress** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Ingress view.

This graph draws data primarily from the Intrusion Events table.

You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

Viewing the Intrusion Event Details List

LICENSE: Protection

At the bottom of the Intrusion Information section is the Intrusion Event Details List, a table that provides classification, estimated priority, and event count information for each intrusion event detected on your monitored network. The events are listed in descending order of event count.

Intrusion Event Details			
Event	Classification	Priority	Events
ICMP-INFO unassigned type 1 undefined	Misc Activity	low	7
DNS dns response for rfc1918 10/8 addr	Potential Corporate Policy Violation	high	7
DNS dns response for rfc1918 192.168/1	Potential Corporate Policy Violation	high	4
DOS tcpdump tcp LDP print zero length r	Attempted Denial of Service	medium	2

The Intrusion Event Details List table is not sortable, but you can click on any table entry to filter or drill down on that information. This table draws data primarily from the Intrusion Events table.

Understanding the Files Information Section

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Files Information section of the Context Explorer contains six interactive graphs that display an overall picture of file and malware events on your monitored network. Five of the graphs display the file types, file names, and malware dispositions of the files detected in network traffic, as well as the hosts sending (uploading) and receiving (downloading) those files. The final graph displays the malware threats detected on your network and, if you have a

FireAMP subscription, on the endpoints where your users installed FireAMP Connectors.

IMPORTANT! If you filter on intrusion information, the entire Files Information Section is hidden.

Note that you must have a Malware license and enable malware detection for Files Information graphs to include network-based malware data. Note also that neither the DC500 Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

For more information on the graphs in the Files Information section, see the following topics:

- [Viewing the Top File Types Graph](#) on page 136
- [Viewing the Top File Names Graph](#) on page 137
- [Viewing the Files by Disposition Graph](#) on page 137
- [Viewing the Top Hosts Sending Files Graph](#) on page 138
- [Viewing the Top Hosts Receiving Files Graph](#) on page 139
- [Viewing the Top Malware Detections Graph](#) on page 140

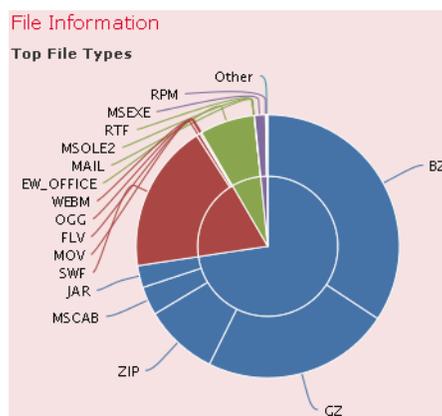
Viewing the Top File Types Graph

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Top File Types graph, in donut form, displays a proportional view of the file types detected in network traffic (outer ring), grouped by file category (inner ring).



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license and enable malware detection for this graph to include network-based malware data. Note also that neither the DC500 Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

This graph draws data primarily from the File Events table.

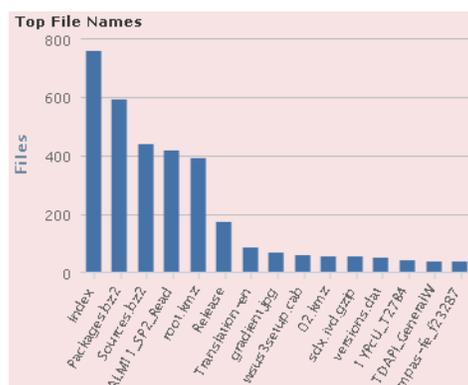
Viewing the Top File Names Graph

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Top File Names graph, in bar form, displays counts of the top unique file names detected in network traffic.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license and enable malware detection for this graph to include network-based malware data. Note also that neither the DC500 Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

This graph draws data primarily from the File Events table.

Viewing the Files by Disposition Graph

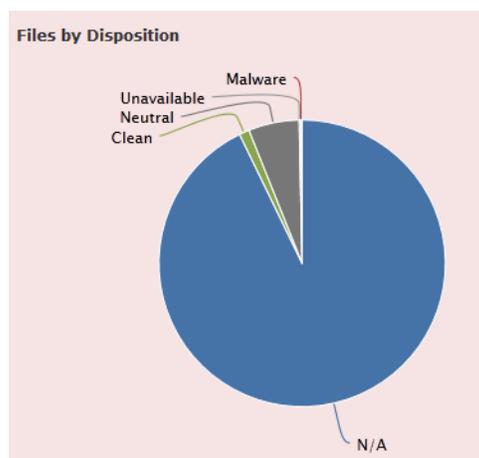
LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Top File Types graph, in pie form, displays a proportional view of the malware dispositions for files detected in network traffic. Note that only files for which the Defense Center performed a malware cloud lookup (which requires a Malware license) have dispositions. Files that did not trigger a cloud lookup have a disposition of N/A. The disposition `unavailable` indicates that the Defense Center could not perform a malware cloud lookup. See [Understanding Malware](#)

[Protection and File Control](#) on page 1199 for descriptions of the other dispositions.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license and enable malware detection for this graph to include network-based malware data. Note also that neither the DC500 Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

This graph draws data primarily from the File Events table.

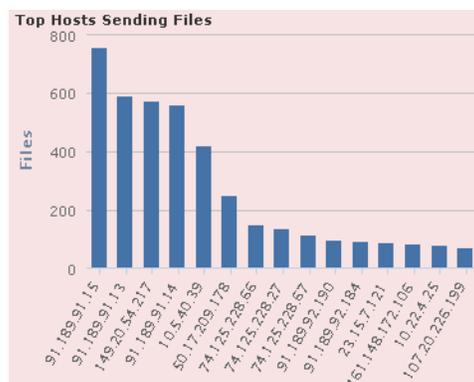
Viewing the Top Hosts Sending Files Graph

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Top Hosts Sending Files graph, in bar form, displays counts of the number of files detected in network traffic for the top file-sending host IP addresses.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays only hosts sending malware, hover your pointer over the graph, then click **Malware** on the toggle button that appears. Click **Files** to return to the default files view. Note that navigating away from the Context Explorer also returns the graph to the default files view.

Note that you must have a Malware license and enable malware detection for this graph to include network-based malware data. Note also that neither the DC500 Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

This graph draws data primarily from the File Events table.

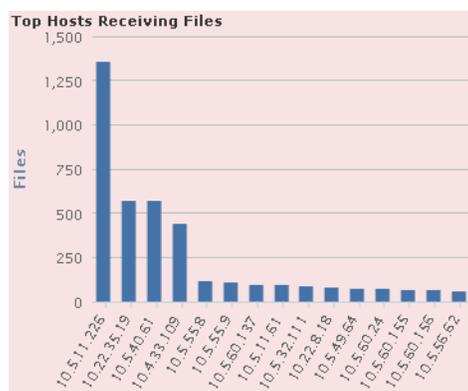
Viewing the Top Hosts Receiving Files Graph

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Top Hosts Receiving Files graph, in bar form, displays counts of the number of files detected in network traffic for the top file-receiving host IP addresses.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays only hosts receiving malware, hover your pointer over the graph, then click **Malware** on the toggle button that appears. Click **Files** to return to the default files view. Note that navigating away from the Context Explorer also returns the graph to the default files view.

Note that you must have a Malware license and enable malware detection for this graph to include network-based malware data. Note also that neither the DC500

Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

This graph draws data primarily from the File Events table.

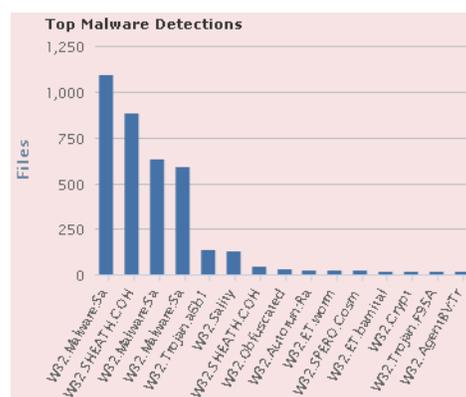
Viewing the Top Malware Detections Graph

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Top Malware Detections graph, in bar form, displays counts of the top malware threats detected on your network and, if you have a FireAMP subscription, on the endpoints where your users installed FireAMP Connectors.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware license and enable malware detection for this graph to include network-based malware data. Note also that neither the DC500 Defense Center nor Series 2 devices support advanced malware detection, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Working with Malware Protection and File Control](#) on page 1197.

This graph draws data primarily from the File Events and Malware Events tables.

Understanding the Geolocation Information Section

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

The Geolocation Information section of the Context Explorer contains three interactive donut graphs that display an overall picture of countries with which hosts on your monitored network are exchanging data: unique connections by initiator or responder country, intrusion events by source or destination country, and file events by sending or receiving country.

For more information on the graphs in the Geolocation Information section, see the following topics:

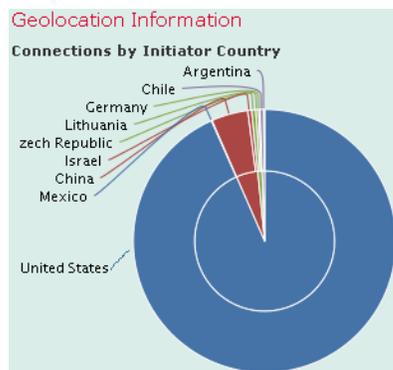
- [Viewing the Connections by Initiator/Responder Country Graph](#) on page 141
- [Viewing the Intrusion Events by Source/Destination Country Graph](#) on page 142
- [Viewing the File Events by Sending/Receiving Country Graph](#) on page 142

Viewing the Connections by Initiator/Responder Country Graph

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

The Connections by Initiator/Responder Country graph, in donut form, displays a proportional view of the countries involved in connections on your network as either the initiator (the default) or the responder. The inner ring groups these countries together by continent. For information about geolocation information, see [Using Geolocation](#) on page 1816. For information about connection data, see [Working With Connection Data](#) on page 539.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays only countries acting as the responder in connections, hover your pointer over the graph, then click **Responder** on the toggle button that appears. Click **Initiator** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Initiator view.

This graph draws data primarily from the Connection Summary Data table.

Viewing the Intrusion Events by Source/Destination Country Graph

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

The Intrusion Events by Source/Destination Country graph, in donut form, displays a proportional view of the countries involved in intrusion events on your network as either the source of the event (the default) or the destination. The inner ring groups these countries together by continent. For information about geolocation information, see [Using Geolocation](#) on page 1816. For information about intrusion event data, see [Working with Intrusion Events](#) on page 593.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays only countries acting as the destinations of intrusion events, hover your pointer over the graph, then click **Destination** on the toggle button that appears. Click **Source** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Source view.

This graph draws data primarily from the Intrusion Events table.

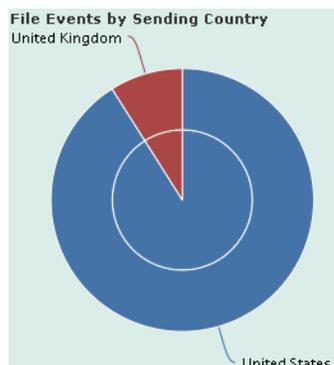
Viewing the File Events by Sending/Receiving Country Graph

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

The File Events by Sending/Receiving Country graph, in donut form, displays a proportional view of the countries detected in file events on your network as either sending (the default) or receiving files. The inner ring groups these countries together by continent. For information about geolocation information,

see [Using Geolocation](#) on page 1816. For information about file event data, see [Working with File Events](#) on page 1222.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

TIP! To constrain the graph so it displays only countries receiving files, hover your pointer over the graph, then click **Receiver** on the toggle button that appears. Click **Sender** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Sender view.

This graph draws data primarily from the File Events table.

Understanding the URL Information Section

LICENSE: FireSIGHT or URL Filtering

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The URL Information section of the Context Explorer contains three interactive bar graphs that display an overall picture of URLs with which hosts on your monitored network are exchanging data: traffic and unique connections associated with URLs, sorted by individual URL, URL category, and URL reputation.

IMPORTANT! If you filter on intrusion event information, the entire URL Information Section is hidden.

Note that you must have a URL Filtering license and enable URL Filtering for URL filtering graphs to include URL category and reputation data. Note also that neither the DC500 Defense Center nor Series 2 devices support URL filtering by reputation and category, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Adding URL Conditions](#) on page 506.

For more information on the graphs in the URL Information section, see the following topics:

- [Viewing the Traffic by URL Graph](#) on page 144
- [Viewing the Traffic by URL Category Graph](#) on page 145
- [Viewing the Traffic by URL Reputation Graph](#) on page 145

Viewing the Traffic by URL Graph

LICENSE: FireSIGHT or URL Filtering

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Traffic by URL graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most requested URLs on your monitored network. For each URL listed, blue bars represent traffic data and red bars represent connection data.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

IMPORTANT! If you filter on intrusion event information, the Traffic by URL graph is hidden.

Note that you must have a URL Filtering license and enable URL Filtering for URL filtering graphs to include URL category and reputation data. Note also that neither the DC500 Defense Center nor Series 2 devices support URL filtering by reputation and category, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Enabling Sourcefire Cloud Communications](#) on page 2032.

This graph draws data primarily from the Connection Events table.

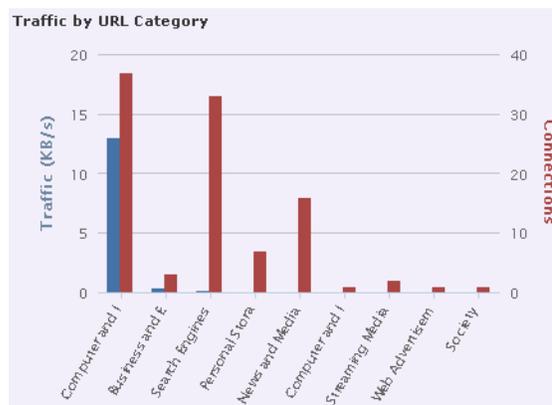
Viewing the Traffic by URL Category Graph

LICENSE: URL Filtering

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Traffic by URL Category graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL categories (such as **Search Engines** or **Streaming Media**) on your monitored network. For each URL category listed, blue bars represent traffic data and red bars represent connection data.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

IMPORTANT! If you filter on intrusion event information, the Traffic by URL Category graph is hidden.

Note that you must have a URL Filtering license and enable URL Filtering for URL filtering graphs to include URL category and reputation data. Note also that neither the DC500 Defense Center nor Series 2 devices support URL filtering by reputation and category, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Adding URL Conditions](#) on page 506.

This graph draws data primarily from the URL Statistics and Connection Events tables.

Viewing the Traffic by URL Reputation Graph

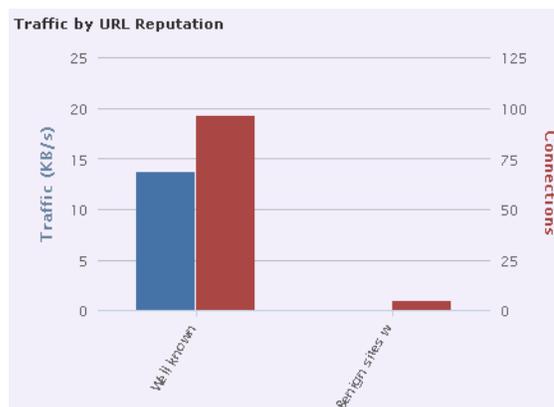
LICENSE: URL Filtering

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Traffic by URL Reputation graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL reputation groups (such as **well known** or **Benign sites with security**

risks) on your monitored network. For each URL reputation listed, blue bars represent traffic data and red bars represent connection data.



Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

IMPORTANT! If you filter on intrusion event information, the Traffic by URL Reputation graph is hidden.

Note that you must have a URL Filtering license and enable URL Filtering for URL filtering graphs to include URL category and reputation data. Note also that neither the DC500 Defense Center nor Series 2 devices support URL filtering by reputation and category, so the DC500 Defense Center cannot display this data and Series 2 devices do not detect it. See [Adding URL Conditions](#) on page 506.

This graph draws data primarily from the URL Statistics and Connection Events tables.

Refreshing the Context Explorer

LICENSE: FireSIGHT

The Context Explorer does not automatically update the information it displays. To incorporate new data, you must manually refresh the explorer.

Note that, although reloading the Context Explorer itself (by refreshing the browser program or navigating away from, then back to, the Context Explorer) refreshes all displayed information, this does not preserve any changes you made to section configuration (such as the Ingress/Egress graphs and the Application Information section) and may cause delays in loading.

To refresh the Context Explorer:

ACCESS: Admin/Any Security Analyst

- ▶ On the Context Explorer, click **Reload** at the upper right.
The explorer updates to display the latest information within your selected time range. Note that the **Reload** button is grayed out until your refresh is finished.

Setting the Context Explorer Time Range

LICENSE: FireSIGHT

You can configure the Context Explorer time range to reflect a period as short as the last hour (the default) or as long as the last year. Note that when you change the time range, the Context Explorer does not automatically update to reflect the change. To apply the new time range, you must manually refresh the explorer.

Changes to the time range persist even if you navigate away from the Context Explorer or end your login session.

To change the Context Explorer time range:

ACCESS: Admin/Any Security Analyst

1. From the **Show the last** drop-down list, select a time range.
2. Optionally, to view data from the new time range, click **Reload**.
All sections of the Context Explorer update to reflect the new time range.

TIP! Clicking **Apply Filters** also applies any time range updates.

Minimizing and Maximizing Context Explorer Sections

LICENSE: FireSIGHT

You can minimize and hide one or more sections of the Context Explorer. This is useful if you want to focus on only certain sections, or if you want a simpler view. You cannot minimize the Traffic and Intrusion Event Counts Time Graph.

Note that Context Explorer sections retain the minimized or maximized states that you configure even if you refresh the page or log out of the appliance.

To minimize a Context Explorer section:

ACCESS: Admin/Any Security Analyst

- ▶ Click the minimize icon () in a section's title bar.

To maximize a Context Explorer section:

ACCESS: Admin/Any Security Analyst

- ▶ Click the maximize icon () in a minimized section's title bar.

Drilling Down on Context Explorer Data

LICENSE: feature dependent

If you want to examine graph or list data in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. (Note that you cannot drill down on the Traffic and Intrusion Events over Time graph.) For example, drilling down on an IP address in the Traffic by Source IP graph displays the Connections with Application Details view of the Connection Events table, including only data associated with the source IP address you selected.

Depending on the type of data you examine, additional options can appear in the context menu. Data points that are associated with specific IP addresses offer the option to view host or whois information on the IP address you select. Data points associated with specific applications offer the option to view application information on the application you select. Data points associated with a specific user offer the option to view that user's user profile page. Data points associated with an intrusion event message offer the option to view the rule documentation for that event's associated intrusion rule, and data points associated with a specific IP address offer the option to blacklist or whitelist that address.

The context menu that you use to drill down on data also contains options to filter that data. For more information on filtering, see [Working with Filters in the Context Explorer](#) on page 150.

To drill down on data in the Context Explorer:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Context Explorer**.

The Context Explorer appears.

2. In any section but Traffic and Intrusion Events over Time, click a data point that you want to investigate.

The context menu pop-up window appears nearby. The following graphic shows the context menu window for NetBIOS in the Hosts by Risk and Application graph.



3. Depending on the data point you selected, you have several options:
 - To view more details of this data in a table view, select **Drill into Analysis**. A new window opens with a detailed table view of the data you selected.
 - If you selected a data point associated with a specific IP address and want more information about the associated host, select **View Host Information**. A new window opens with a host profile page for the IP address you selected. For more information on host attributes and host profiles, see [Using Host Profiles](#) on page 1334.
 - If you selected a data point with a specific IP address and want to make a whois search on that address, select **Whois**. A new window opens with the results of a whois query for the IP address you selected.
 - If you selected a data point associated with a specific application and want more information about that application, select **View Application Information**. A new window opens with information on the application you selected. For more information about application attributes, see [Understanding Application Detection](#) on page 1263.
 - If you selected a data point associated with a specific user and want more information about that user, select **View User Information**. A new window opens with a user profile page for the user you selected. For more information on user details, see [Understanding User Details and Host History](#) on page 1449.

- If you selected a data point associated with a specific intrusion event message and want more information about the associated intrusion rule, select **View Rule Documentation**.
A new window opens with a rule details page relevant to the event you selected. For more information on intrusion rule details, see [Viewing Rule Details](#) on page 723.
- If you selected a data point associated with a specific file and want to view that file's trajectory, select **View Network File Trajectory**.
A new window opens with the trajectory map for the selected file. For more information on using the network file trajectory feature, see [Working with Network File Trajectory](#) on page 1241.
- If you selected a data point associated with a specific IP address and want to add that IP address to the Security Intelligence global blacklist or whitelist, select the appropriate option: **Blacklist Now** or **Whitelist Now**. Confirm your choice in the pop-up window that appears.
The IP address is blacklisted or whitelisted. For more information, see [Working with the Global Whitelist and Blacklist](#) on page 166.
These options are not listed on the DC500 Defense Center, which does not support Security Intelligence data.

Working with Filters in the Context Explorer

LICENSE: FireSIGHT

Beyond the basic, wide-ranging data that the Context Explorer initially displays, you have the option to filter that data for a more granular contextual picture of activity on your network. Filters encompass all types of FireSIGHT data, support exclusion as well as inclusion, can be applied quickly by clicking on Context Explorer graph data points, and affect the entire explorer. You can apply up to 20 filters at once to create a highly specific portrait tailored to the needs of your network and organization. Filters that you apply are reflected in the Context Explorer URL so that you can bookmark useful filter sets in your browser program for later use.

For information on using filters in the Context Explorer, see the following topics:

- [Adding and Applying Filters](#) on page 151
- [Creating Filters with the Context Menu](#) on page 155
- [Bookmarking Filters](#) on page 156

Adding and Applying Filters

LICENSE: FireSIGHT, Protection, Control, URL Filtering, or Malware

SUPPORTED DEVICES: feature dependent

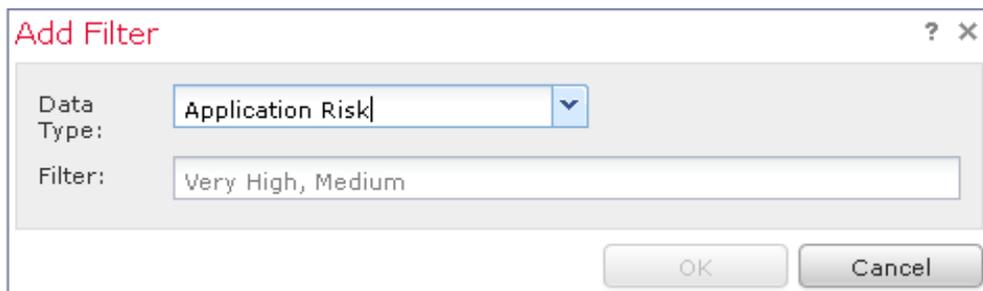
SUPPORTED DEFENSE CENTERS: feature dependent

You can add filters to Context Explorer data in several ways:

- from the Add Filter window
- from the context menu pop-up window, when you select a data point in the explorer
- from the Context Explorer icon (**sf**) or from text links that appear in certain detail view pages (Application Detail, Host Profile, Rule Detail, and User Profile). Clicking these links automatically opens and filters the Context Explorer according to the relevant data on the detail view page. For example, clicking the Context Explorer link on a user detail page for the user **jenkins** constrains the explorer to show only data associated with that user

This section focuses on creating filters from scratch with the Add Filter window. For information on using the context menu to create quick filters from Context Explorer graph and list data, see [Creating Filters with the Context Menu](#) on page 155.

The Add Filter window, which you access by clicking the plus icon (**+**) under **Filters** at the top left of the Context Explorer, contains only two fields: **Data Type** and **Filter**.



The screenshot shows a dialog box titled "Add Filter" with a question mark and close icon in the top right corner. It contains two main input fields: "Data Type" and "Filter". The "Data Type" field is a dropdown menu currently showing "Application Risk". The "Filter" field is a text input box containing the text "Very High, Medium". Below these fields are two buttons: "OK" and "Cancel".

The Data Type drop-down list contains over 25 different types of Sourcefire 3D System data you can use to constrain the Context Explorer. After you select a data type, you then enter a specific value for that type in the **Filter** field (for example, a value of Shopping for the type URL Category). To assist you, the Filter field presents several grayed-out example values for the data type you select. (These are erased when you enter data in the field.)

The [Filter Data Types](#) table lists the data types available as filters, with examples and brief definitions of each. Note that The DC500 Defense Center does not display and Series 2 devices do not detect data for features they do not support.

See the [Supported Capabilities by Appliance Model table](#) on page 42 for a summary of Series 2 appliance features.

Filter Data Types

TYPE	EXAMPLE VALUES	DEFINITION
Access Control Action	Allow, Block	Action taken by your access control policy to allow or block traffic
Application Category	web browser, email	General classification of an application's most essential function
Application Name	Facebook, HTTP	Name of an application
Application Risk	Very High, Medium	Estimated security risk of an application
Application Tag	encrypts communications, sends mail	Additional information about an application; applications can have any number of tags, including none
Application Type	Client, web Application	Type of an application: application protocol, client, or web application
Business Relevance	Very Low, High	Estimated relevance of an application to business activity (as opposed to recreation)
Continent	North America, Asia	Continent associated with a routable IP address detected on your monitored network
Country	Canada, Japan	Country associated with a routable IP address detected on your monitored network
Device	device1.example.com, 192.168.1.3	Name or IP address of a device on your monitored network
Event Classification	Potential Corporate Policy Violation, Attempted Denial of Service	Capsule description of an intrusion event, determined by the classification of the rule, decoder, or preprocessor that triggered it
Event Message	dns response, P2P	Message generated by an event, determined by the rule, decoder, or preprocessor that triggered it
File Disposition	Malware, Clean	Cloud-determined disposition of a file for which the Defense Center performed a malware cloud lookup
File Name	Packages.bz2	Name of a file detected in network traffic

Filter Data Types (Continued)

TYPE	EXAMPLE VALUES	DEFINITION
File SHA256	any 32-bit string	SHA-256 hash value of a file for which the Defense Center performed a malware cloud lookup
File Type	GZ, SWF, MOV	File type detected in network traffic
File Type Category	Archive, Multimedia, Executables	General category of file type detected in network traffic
IP Address	192.168.1.3, 2001:0db8:85a3::0000/24	IPv4 or IPv6 addresses, address ranges, or address blocks Note that searching for an IP address returns events where that address was either the source or the destination for the event
Impact Level	Impact Level 1, Impact Level 2	Estimated impact of an event on your monitored network
Inline Result	dropped, would have dropped	Whether traffic was dropped, would have been dropped, or was not acted upon by the system
Malware Threat Name	w32.Trojan.a6b1	The name of a malware threat
OS Name	windows, Linux	Name of an operating system
OS Version	xp, 2.6	Specific version of an operating system
Priority	high, low	Estimated urgency of an event
Security Zone	My Security Zone, Security Zone X	A set of interfaces through which traffic is analyzed and, in an inline deployment, passes
URL	example.com, loremipsum.net	URL requested by a host on your monitored network
URL Category	shopping, Travel	General category of a URL; URLs can have more than one category
URL Reputation	well known, high risk	Estimated potential security risk of a URL
User	wsmith, mtwain	Identity of a user logged in to a host on your monitored network

In the Filter field, you can input special search parameters such as * and ! essentially as you can in event searches. You can create exclusionary filters by prefixing filter parameters with the ! symbol. For more information on the search constraints typically supported by the Sourcefire 3D System, see [Using Wildcards and Symbols in Searches](#) on page 1775.

When multiple filters are active, values for the same data type are treated as OR search criteria: all data that matches at least one of the values appears. Values for different data types are treated as AND search criteria: to appear, data must match at least one value for each filtered data type. For example, data that appears for the filter set of **Application: 2channel**, **Application: Reddit**, and **User: edickinson** must be associated with the user **edickinson** **AND** either the application **2channel** **OR** the application **Reddit**.

After you confirm a data type and value for your filter, a filter widget appears at the top left of the page, displaying the new filter's data type and value.

Filters: Access Control Action **Allow** ✕ Application Type **Client** ✕

Because you may want to configure multiple filters before you apply them, and because the Context Explorer may take time to fully reload all sections, filters that you add are not automatically applied. To apply filters, you must click **Apply Filters**. Filters that are configured, but not yet applied, appear faded. You can have up to 20 filters at a time, and you can delete individual filters by clicking the delete icon (✕) on the filter's widget. If you want to delete all filters at once, you can click the **Clear** button.

Note that some filter types are incompatible with others: for example, URL-related filters (such as **URL**, **Category**, and **Reputation**) cannot be applied at the same time as intrusion event-related filters (such as **Device** and **Inline Result**) because the system cannot sort connection event data by intrusion event data. The system automatically prevents incompatible filters from simultaneously applying; when one filter type is more recently activated, filters of the incompatible type are hidden as long as the incompatibility exists.

Note that the data displayed depends on such factors as how you license and deploy your managed devices, whether you configure features that provide the data and, in the case of Series 2 appliances, whether the appliance supports a feature that provides the data. For example, because neither the DC500 Defense Center nor Series 2 devices support URL filtering by category and reputation, the DC500 Defense Center does not display data for this feature and Series 2 devices do not detect this data.

To create a new filter from the Add Filter window:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Context Explorer**.

The Context Explorer appears.

2. Under **Filters** at the top right, click the plus icon (+).
The Add Filter pop-up window appears.
3. From the **Data Type** drop-down list, select the data type you want to filter on.
The Filter field populates with example values for that data type.
4. In the **Filter** field, type the data type value you want to filter on.
5. Click **OK**.
Your filter is added. The Context Explorer reappears and a corresponding filter widget appears.
6. Optionally, repeat the previous steps to add more filters until you have the filter set you need. Note that because the Context Explorer does not automatically refresh, your filters are not applied when you add them.
7. Click **Apply Filters**.
Your filters are applied and the Context Explorer refreshes to reflect the filtered data.

To delete a filter:

ACCESS: Admin/Any Security Analyst

- ▶ Click the delete icon (✕) on any filter widget.
The filter is deleted.

To clear all filters:

ACCESS: Admin/Any Security Analyst

- ▶ Click the **Clear** button that appears to the right of the filter widgets.



All filters are cleared.

Note that this button does not appear if no filters have been created.

Creating Filters with the Context Menu

LICENSE: FireSIGHT

While exploring Context Explorer graph and list data, you can click on data points, then use the context menu to quickly create a filter based on that data, either inclusive or exclusive. If you use the context menu to filter on information of data type Application, User, or Intrusion Event Message, or any individual host, the filter widget includes a widget information icon that links to the relevant detail page for that data type (such as Application Detail for application data).

You can also use the context menu to investigate specific graph or list data in more detail. For information, see [Drilling Down on Context Explorer Data](#) on page 148.

To create a filter from the context menu:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Context Explorer**.

The Context Explorer appears.

2. In any explorer section except Traffic and Intrusion Events over Time, click a data point you want to filter on.

The context menu pop-up window appears nearby. The following graphic shows the context menu window for NetBIOS in the Hosts by Risk and Application graph.



3. You have two options:

- To add a filter for this data, click **Add Filter**.
The filter is added and its widget appears at upper left.
- To add an exclusion filter for this data, click **Add Exclude Filter**. The filter, when applied, displays all data **not** associated with the excluded value.
The filter is added and its widget appears at upper left. Exclude filters display an exclamation point before the filter value.

To view filter detail:

ACCESS: Admin/Any Security Analyst

- ▶ Click the information icon () on any eligible filter widget.

A new window opens with the detail page relevant to the filter's data type.

Bookmarking Filters

LICENSE: FireSIGHT

Filters function as a simple, agile tool to get the precise FireSIGHT data context you need at any given time. They are not intended as permanent configuration settings, and disappear when you navigate away from the Context Explorer or end your session. However, your organization may use certain filter combinations frequently. To preserve filter settings for later use, you can create a browser bookmark of the Context Explorer with your preferred filters applied. Because

applied filters are incorporated in the Context Explorer page URL, loading a bookmark of that page also loads the corresponding filters.

CHAPTER 4

USING OBJECTS AND SECURITY ZONES

For increased flexibility and web interface ease-of-use, the Sourcefire 3D System 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 objects for IP addresses and networks, Security Intelligence feeds and lists, port/protocol pairs, VLAN tags, URLs, application filters, file lists, and security zones. You can then use these objects in various places in the system's web interface, including access control policies, intrusion policy variables, intrusion rules, network discovery rules, event searches, reports, dashboards, and so on.

Grouping objects allows you to reference multiple objects with a single configuration. You can group network, port, VLAN tag, and URL objects.

IMPORTANT! In most cases, editing an object used in an access control, network discovery, or intrusion policy requires a policy reapply for your changes to take effect. Editing a security zone also requires that you reapply the appropriate device configurations.

For more information, see the following sections:

- [Using the Object Manager](#) on page 159
- [Working with Network Objects](#) on page 161
- [Working with Security Intelligence Lists and Feeds](#) on page 162
- [Working with Port Objects](#) on page 173
- [Working with VLAN Tag Objects](#) on page 174

- [Working with URL Objects](#) on page 175
- [Working with Application Filters](#) on page 176
- [Working with the Global Malware Whitelist](#) on page 180
- [Working with Security Zones](#) on page 187

Using the Object Manager

LICENSE: Any

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



The screenshot shows the Object Manager interface. On the left is a navigation pane with categories: Network, Security Intelligence, Port, VLAN Tag, and URL. Each category has sub-items for 'Individual Objects' and 'Object Groups'. The main area displays a table with columns 'Name' and 'Value'. There are two rows: 'IPv4 Private Networks' with values '192.168.0.0/16', '10.0.0.0/8', and '172.16.0.0/12'; and 'IPv6 Private Networks' with value 'fc00::/7'. Each row has edit and delete icons. At the top right are 'Add Network' and 'Filter' buttons. At the bottom, it says 'Displaying 1 - 2 of 2 rows' and 'Page 1 of 1'.

Name	Value	
IPv4 Private Networks	192.168.0.0/16 10.0.0.0/8 172.16.0.0/12	 
IPv6 Private Networks	fc00::/7	 

For more information, see:

- [Grouping Objects](#) on page 159
- [Browsing, Sorting, and Filtering Objects](#) on page 161

Grouping Objects

LICENSE: Any

You can group network, port, VLAN tag, and URL objects. The system allows you to use objects and object groups interchangeably in the web interface. 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.

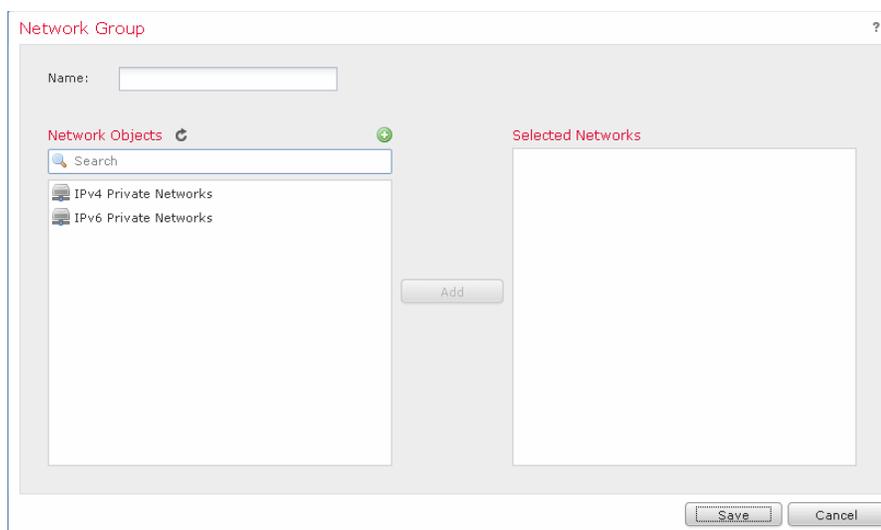
When you edit an object group used in a policy, for example, a network object group used in an access control, network discovery, or intrusion policy, you must reapply the policy for your changes to take effect.

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 VLAN tag group that you are using in a VLAN condition in a saved access control policy.

To group network, port, VLAN tag, or URL objects:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
The Object Management page appears.
2. Under the type of object you want to group, select **Object Groups**.
The page for the type of object you are grouping appears.
3. Click **Add Network Group**, **Add Port Group**, **Add VLAN Tag Group**, or **Add URL Group**.
A pop-up window appears where you can create the group. The following graphic shows the Network Group pop-up window.



4. Type a **Name** for the group. You can use any printable standard ASCII characters except curly braces ({}).
5. Select one or more objects and click **Add**.
 - Use Shift and Ctrl to select 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.

6. Click **Save**.
The group is created.

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.

Displaying 1 - 20 of 21 rows  Page of 2 

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:

ACCESS: Admin/Access Admin/Network Admin

- ▶ Click a column heading. To sort in the opposite direction, click the heading again.

To filter objects or groups:

ACCESS: Admin/Access Admin/Network Admin

- ▶ Type 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

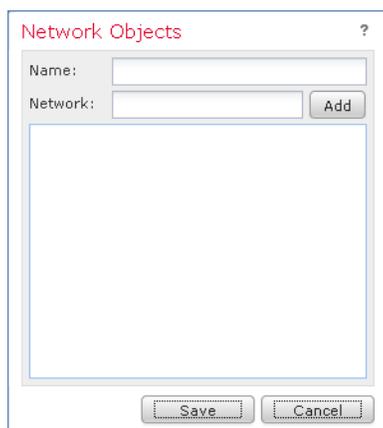
A network object represents one or more IP addresses that you can specify either individually or as address ranges. You can use network objects and groups (see [Grouping Objects](#) on page 159) in various places in the system's web interface, including access control policies, intrusion policy variables, intrusion rules, network discovery rules, event searches, 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, network discovery, or intrusion policy, you must reapply the policy for your changes to take effect.

To create a network object:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
The Object Management page appears.
2. Under **Network**, select **Individual Objects**.
3. Click **Add Network**.
The Network Objects pop-up window appears.



4. Type a **Name** for the network object. You can use any printable standard ASCII characters except curly braces ({}).
5. For each IP address or address range you want to add to the network object, type its value and click **Add**.
6. Click **Save**.
The network object is added.

Working with Security Intelligence Lists and Feeds

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

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.

IMPORTANT! Although they have all other Protection capabilities by default, Series 2 devices cannot perform Security Intelligence filtering.

Comparing Feeds and Lists

A Security Intelligence *feed* is a dynamic collection of IP addresses that the Defense Center 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, Sourcefire provides the *Sourcefire Intelligence Feed*, which represents IP addresses determined by the Sourcefire VRT to have a poor reputation.

When the Defense Center downloads updated feed information, it automatically updates its managed devices. Although it may take a few minutes for a feed update to take effect throughout your deployment, you do not have to reapply access control policies after you create or modify a feed, or after a scheduled feed update.

IMPORTANT! If you want strict control over when the Defense Center downloads a feed from the internet, you can disable automatic updates for that feed. However, Sourcefire 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.

A Security Intelligence *list*, contrasted with a feed, is a simple static list of IP addresses that you manually upload to the Defense Center. 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 an access control policy apply 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 range per line. Comment lines must start with the # character. List source files must use the `.txt` extension.

If the Defense Center 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, the Defense Center updates its managed devices with the addresses it can recognize.

The default health policy includes the Security Intelligence module, which alerts in a few situations involving Security Intelligence filtering, including if the Defense Center cannot update a feed, or if a feed is corrupt or contains no recognizable IP addresses.

Internet Access and High Availability

The system uses port 443/HTTPS to download the Sourcefire 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 Defense Center. If your Defense Center does not have direct access to the feed site, it can use a proxy server (see [Configuring Network Settings](#) on page 2009).

IMPORTANT! The Defense Center 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.

Although Security Intelligence objects are synchronized between Defense Centers in a high availability deployment, only the primary Defense Center downloads feed updates. If the primary Defense Center fails, you must not only make sure that the secondary Defense Center has access to the feed sites, but also must use the web interface on the secondary Defense Center to switch its role to **Active & Primary**. For more information, see [Monitoring and Changing High Availability Status](#) on page 202.

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. (For information on creating and managing network objects and groups, see [Working with Network Objects](#) on page 161.)



The screenshot shows a web interface for managing Security Intelligence objects. At the top, there are three buttons: "Update Feeds", "Add Security Intelligence" (with a green plus icon), and "Filter" (with a magnifying glass icon). Below these is a table with the following columns: "Name", "Type", and an empty column for actions. The table contains five rows of data:

Name	Type	
Global Blacklist	List	 
Global Whitelist	List	 
MyCompany Whitelist <i>Last Updated: 2012-05-17 07:21:53</i>	List	 
Sourcefire Intelligence Feed <i>Last Updated: 2012-05-17 06:58:22</i>	Feed	 
Zeus Blacklist <i>Last Updated: 2012-05-17 07:48:05</i>	Feed	 

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 Sourcefire 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.

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 (add IP addresses using the context menu)	disable or change update frequency	fully modify	upload a modified list only	fully modify
requires access policy control reapply when modified?	yes when deleting (adding IP addresses does not require reapply)	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](#) on page 166
- [Working with the Sourcefire Intelligence Feed](#) on page 167
- [Working with Custom Security Intelligence Feeds](#) on page 168
- [Manually Updating Security Intelligence Feeds](#) on page 170
- [Working with Custom Security Intelligence Lists](#) on page 170
- [Filtering Traffic Based on Security Intelligence Data](#) on page 438

Working with the Global Whitelist and Blacklist

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

In the course of your analysis, you can build a Security Intelligence *global blacklist* by using the IP address context menu in an event view, the Context Explorer, or a dashboard. For example, if you notice a set of routable IP addresses in intrusion events associated with exploit attempts, you can immediately blacklist those IP addresses. You can also build a *global whitelist* in a similar fashion.

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.

When you add an IP address to a global list, the Defense Center automatically updates its managed devices. Although it may take a few minutes for your changes to take effect throughout your deployment, you do not have to reapply access control policies after adding an IP address to a global list. Conversely, after you delete IP addresses from the global whitelist or blacklist, you must apply your access control policies 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 ranges using a /0 netmask in those objects will be ignored and whitelist and blacklist filtering will not occur based on those addresses. Address ranges with a /0 netmask from security intelligence feeds will also be 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.

Because adding an IP address to the global whitelist or blacklist affects access control, you must have one of the following:

- Administrator access
- a combination of default roles: Network Admin or Access Admin, plus Security Analyst and Security Approver
- a custom role with both Modify Access Control Policy and Apply Access Control Policy permissions; see [Using Custom User Roles with Access Control Policies](#) on page 433

To add an IP address to the global whitelist or blacklist using the context menu:

ACCESS: Admin/Custom

1. In an event view, packet view, the Context Explorer, or a dashboard, hover your pointer over an IP address hotspot.

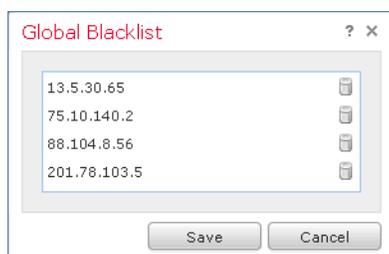
TIP! In an event view or dashboard, hover your pointer over an IP address, not the host icon () to its left.

2. Invoke the context menu:
 - In an event view or dashboard, right-click.
 - In the Context Explorer or packet view, left-click.
3. From the context menu, select either **Whitelist Now** or **Blacklist Now**.
For information on the other options in the context menu, see [Using the Context Menu](#) on page 63.
4. Confirm that you want to whitelist or blacklist the IP address.
A confirmation dialog appears, warning you that it can take several minutes for the addition to take effect.
5. Click **Close**.
After the Defense Center communicates your addition to its managed devices, your deployment begins filtering traffic according to your change.

To remove IP addresses from the global whitelist or blacklist:

ACCESS: Admin/Network Admin

1. On the object manager's Security Intelligence page, next to the global whitelist or blacklist, click the edit icon (✎).
The Global Whitelist or Global Blacklist pop-up window appears.



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 select them, then right-click and select **Delete**.
3. Click **Save**.
Your changes are saved, but you must apply your access control policies for them to take effect.

Working with the Sourcefire Intelligence Feed

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

To help you build blacklists, Sourcefire provides the Sourcefire 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 apply new policies.

Although you cannot delete the Sourcefire 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:

ACCESS: Admin/Network Admin

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

The Sourcefire Security Intelligence pop-up window appears.



2. Edit the **Update Frequency**.

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

3. Click **Save**.

Your changes are saved.

Working with Custom Security Intelligence Feeds

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

Custom or third-party Security Intelligence feeds allow you to augment the Sourcefire Intelligence Feed with other regularly-updated reputable whitelists and blacklists on the internet. You can also set up an internal feed, which is useful if you want to update multiple Defense Centers in your deployment using one source list.

When you configure a feed, you specify its location using a URL; the URL cannot be Punycode-encoded. By default, the Defense Center downloads the entire feed source on the interval you configure, then automatically updates its managed devices.

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

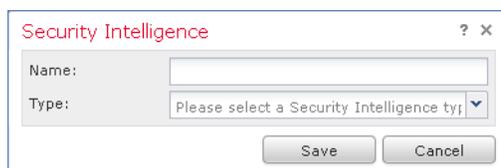
the last time the Defense Center downloaded the feed, the system does not need to re-download it. You may want to use md5 checksums for internal feeds, especially if they are large. The md5 checksum must be stored in a simple text file with only the checksum. Comments are not supported.

To configure a Security Intelligence feed:

ACCESS: Admin/Network Admin

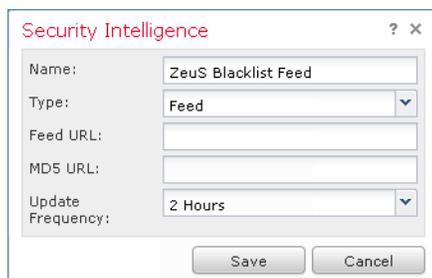
1. On the object manager's Security Intelligence page, click **Add Security Intelligence**.

The Security Intelligence pop-up window appears.



2. Type a **Name** for the feed. You can use any printable standard ASCII characters except curly braces ({}).
3. From the **Type** drop-down list, specify that you want to configure a **Feed**.

The pop-up window updates with new options.



4. Specify a **Feed URL** and optionally, an **MD5 URL**.
5. Select an **Update Frequency**.

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

6. Click **Save**.

The Security Intelligence feed object is created. Unless you disabled feed updates, the Defense Center 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

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

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

To update all Security Intelligence feeds:

ACCESS: Admin/Access Admin/Network Admin

1. On the object manager's Security Intelligence page, click **Update Feeds**.
2. Confirm that you want to update all feeds.

A confirmation dialog appears, warning you that it can take several minutes for the update to take effect.

3. Click **OK**.

After the Defense Center downloads and verifies the feed updates, it communicates any changes to its managed devices. Your deployment begins filtering traffic using the updated feeds.

Working with Custom Security Intelligence Lists

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

A Security Intelligence list is a simple static list of IP addresses and address ranges that you manually upload to the Defense Center. Custom lists are useful if you want to augment and fine-tune feeds or one of the global lists, for a single Defense Center's managed devices.

Note that netmasks for address ranges 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](#) on page 172.

To upload a new Security Intelligence list to the Defense Center:

ACCESS: Admin/Access Admin/Network Admin

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

The Security Intelligence pop-up window appears.



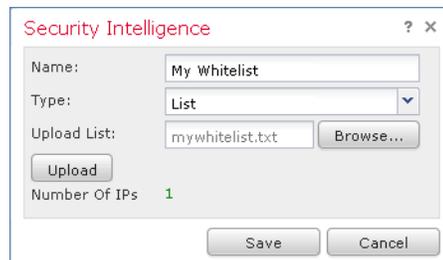
2. Type a **Name** for the list. You can use any printable standard ASCII characters except curly braces ({}).
3. From the **Type** drop-down list, specify that you want to upload a **List**.

The pop-up window updates with new options.



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 ranges that the system found in the list.



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

5. Click **Save**.
The Security Intelligence list object is saved. You can now use it in access control policies.

Updating a Security Intelligence List

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

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 the Defense Center web interface. If you do not have access to the source file, you can download a copy from the Defense Center.

To modify a Security Intelligence list:

ACCESS: Admin/Access Admin/Network Admin

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

The Security Intelligence pop-up window appears.



2. If you need a copy of the list to edit, click **Download**, then follow your browser's prompts to save the list as a text file.
3. Make changes to the list as necessary.
4. On the Security Intelligence pop-up window, click **Browse** to browse to the modified list, then click **Upload**.

The list is uploaded.

5. Click **Save**.

Your changes are saved. If the list is being used by an active access control policy, you must apply the policy for your changes to take effect.

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

You can use port objects and groups (see [Grouping Objects](#) on page 159) in various places in the system's web interface, including access control policies, network discovery rules, and event searches. For example, if your organization uses a custom client that uses a specific range of ports and causes the system to generate excessive and misleading events, you can configure your network discovery policy to exclude monitoring those ports.

You cannot delete a port object that is in use. Additionally, after you edit a port object used in an access control or network discovery policy, you must reapply the policy for your changes to take effect.

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 to the managed device on policy apply. 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:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
The Object Management page appears.
2. Under **Port**, select **Individual Objects**.

3. Click **Add Port**.

The Port Objects pop-up window appears.



4. Type a **Name** for the port object. You can use any printable standard ASCII characters except curly braces ({}).
5. Select a **Protocol**.
You can quickly select **TCP**, **UDP**, **IP**, **ICMP**, or **IPv6-ICMP**, or you can use the **Other** drop-down list to select either a different protocol or **All** protocols.
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.
7. Optionally, restrict a 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.
8. Optionally, select **Other** and a protocol from the drop-down list. If you select **All** protocols, type a port number in the **Port** field.
9. Click **Save**.
The port object is added.

Working with VLAN Tag Objects

LICENSE: Any

Each VLAN tag object you configure represents a VLAN tag or range of tags. You can use VLAN tag objects and groups (see [Grouping Objects](#) on page 159) in various places in the systems's web interface, including access control policies and event searches. For example, you could write an access control rule that applies only to a specific VLAN.

You cannot delete a VLAN tag object that is in use. Additionally, after you edit a VLAN tag object used in an access control policy, you must reapply the policy for your changes to take effect.

To add a VLAN tag object:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
2. Under **VLAN Tag**, select **Individual Objects**.
3. Click **Add VLAN Tag**.

The VLAN Tag pop-up window appears.



4. Type a **Name** for the VLAN tag. You can use any printable standard ASCII characters except curly braces ({}).
5. In the **VLAN Tag** field, type a value for the VLAN tag.
You can specify any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.
6. Click **Save**.
The VLAN tag object is added.

Working with URL Objects

LICENSE: Any

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

Each URL object you configure represents a single URL or IP address. You can use URL objects and groups (see [Grouping Objects](#) on page 159) in various places in the system's web interface, including access control policies and event searches. 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 a URL object used in an access control policy, you must reapply the policy for your changes to take effect.

To add a URL object:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
The Object Management page appears.
2. Under **URL**, select **Individual Objects**.
3. Click **Add URL**.
The URL Objects pop-up window appears.



4. Type a **Name** for the URL object. You can use any printable standard ASCII characters except curly braces ({}).
5. Type the **URL** or IP address for the URL object.
6. Click **Save**.
The URL object is added.

Working with Application Filters

LICENSE: FireSIGHT

SUPPORTED DEVICES: Series 3, virtual

When the Sourcefire 3D System 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 Sourcefire frequently updates and adds additional detectors via system and vulnerability database (VDB) updates. You can also create your own application protocol detectors to enhance the system's detection capabilities.

Application filters group applications according to criteria associated with the applications' risk, business relevance, type, categories, and tags; see the [Application Characteristics table](#) on page 1264. When you create an application protocol detector, you must characterize the application using those criteria as well. 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 [Working with Application Conditions](#) on page 498.

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

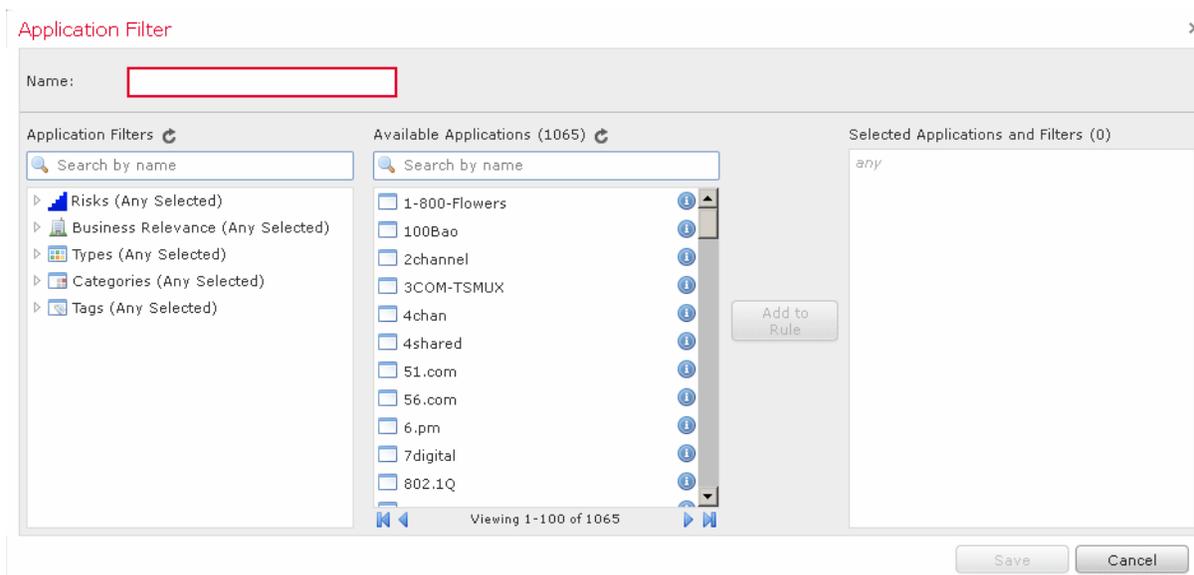
detector, the policy is updated when you update the VDB. Although you must reapply the policy before the system can block the new application, you do not have to update the access control rule that blocks the application.

If the Sourcefire-provided application filters do not group applications according to your needs, you can create your own filters. User-defined filters can group and combine Sourcefire-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 system software or the VDB.

As with Sourcefire-provided application filters, you can use user-defined application filters in access control rules. You can also use user-defined filters in the following additional ways:

- To search for applications using the event viewer; see [Using Objects and Application Filters in Searches](#) on page 1775
- To constrain a table view in a report template; see [Working with Searches in Report Template Sections](#) on page 1749
- To filter application statistics in a Custom Analysis dashboard widget; see [Configuring the Custom Analysis Widget](#) on page 82

You use the object manager (**Objects > 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 Sourcefire-provided application filters that you can select 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 select both filters, the system would display 150 available applications.

The system links different types of filters with an AND operation. For example, if you select 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 pop-up that appears.

After you determine the applications you want to add to the filter, you can add them either individually, or, if you selected 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.

Name	Value 	 
Very High Risk/Very Low Business Relevance	0 application(s) 2 filter(s)	 
Social Networking	1 Filter	 
MyOrganization Custom Apps	5 application(s) 0 filter(s)	 

For information on sorting and filtering the application filters that appear, see [Using the Object Manager](#) on page 159. Note that you cannot delete an application filter that is in use. Additionally, after you edit an application filter used in an access control policy, you must reapply the policy for your changes to take effect.

To create an application filter:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
The Object Management page appears.
2. Click **Application Filters**.
The Application Filters section appears.

3. Click Add Application Filter.

The Application Filter pop-up window appears.

4. Give the filter a Name. You can use any printable standard ASCII characters except curly braces ({}).

5. Optionally, use Sourcefire-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**.

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

6. 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.
- Use Shift and Ctrl keys to select 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 select individual applications and **All apps matching the filter** at the same time.

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**.

8. Click **Save**.
The application filter is saved.

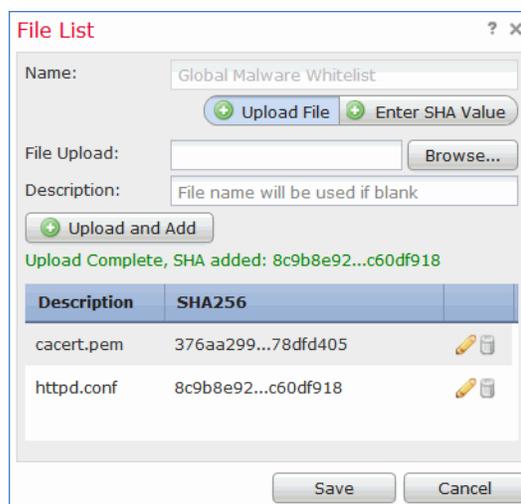
Working with the Global Malware Whitelist

LICENSE: Malware

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

If you use network-based advanced malware protection (AMP), and the Sourcefire cloud incorrectly identifies a file's disposition, you can add the file to the *global malware whitelist*. The system does not perform malware cloud lookups or block malware for files on the list, even files identified as malware by the cloud. Note, however, that the system may still detect and block files on the whitelist using file rules with the Detect and Block Files actions, respectively.



The system's global malware whitelist is included by default in every file policy. You can opt not to use this global list on a per-policy basis.

WARNING! If you enable the global malware whitelist in a file policy, the system ignores files on the list and does not block them as malware, even if the file policy includes Block Malware rules targeting the traffic containing the files. Do **not** include files on this list that are actually malware.

The whitelist is composed of SHA-256 values. You can upload a file and the system calculates the file's SHA-256 value then adds it to the whitelist. You can also calculate a file's SHA-256 value manually and enter it directly. From the event view, you can add a file's SHA-256 value to the whitelist using the context menu.

When you add a file to the global list, you must reapply any access control policies with file policies that use the list for the changes to take effect. Similarly, after you delete or edit file SHA-256 values on the global whitelist, you must reapply your access control policies for your changes to take effect.

Because adding a file to the global malware whitelist affects access control, you must have one of the following to manage all aspects of the global malware whitelist:

- Administrator access
- a combination of Network Admin or Access Admin (to edit the whitelist), Security Approver (to reapply access control policies), and Security Analyst or Security Analyst (RO) (to add a file using the SHA-256 value from the event view) access
- a custom role with Modify Access Control Policy and Object Manager (to edit the whitelist), Apply Access Control Policy (to reapply access control policies), and Modify File Events (to add a file using the SHA-256 value from the event view) permissions; see [Using Custom User Roles with Access Control Policies](#) on page 433

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

- [Adding a File to the Whitelist from the Event View](#) on page 181
- [Adding a File to the Whitelist by Uploading the File](#) on page 182
- [Adding a File to the Whitelist Using the SHA-256 Value](#) on page 184
- [Modifying Files on the Global Malware List](#) on page 185

Adding a File to the Whitelist from the Event View

LICENSE: Malware

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

You can add a file to the global malware whitelist from the file or malware event view, using the context menu on the file's SHA-256 value. If you add a file using the context menu, the system adds the filename as the description.

To add a file to the global whitelist using the context menu:

ACCESS: Admin/Any Security Analyst

1. In the Files or Malware event views, hover your pointer over a SHA-256 hotspot.
2. Right-click and select **Whitelist SHA** from the context menu.
For information on the other options in the context menu, see [Using the Context Menu](#) on page 63.
3. Confirm that you want to add the SHA-256 value representing the file to the whitelist.
A confirmation dialog appears, noting that you must reapply any access control policies with file policies that use the whitelist.
4. Click **Close**.
5. Reapply all access control policies with file policies that use the global malware whitelist.
After the policies apply, the system no longer performs malware cloud lookups on the file or blocks the file as malware.

Adding a File to the Whitelist by Uploading the File

LICENSE: Malware

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

If you have a copy of the file you want to add to the global malware whitelist, you can upload it and the system calculates the file's SHA-256 value and adds the file to the whitelist. The system does not enforce a limit on the size of files for SHA-256 calculation. You cannot upload files that duplicate SHA-256 values already on the list.

To add a file by uploading a file to calculate its SHA-256 value:

ACCESS: Admin/Network Admin

1. Select **Objects > Object Management**, then click **File List**.
2. Next to the global malware whitelist, click the edit icon (✎).
The File List pop-up window appears.

Description	SHA256
-------------	--------

3. Click **Upload File**.
The pop-up window updates to include new fields.

Description	SHA256
-------------	--------

4. Browse to and select the file you want to add to the whitelist.
5. Optionally, type a **Description** for the file. If you leave the description blank, the system uses the file name as the description.
6. Click **Upload and Add**.
The system calculates the SHA-256 value for the file and adds the file and the value to the list.
7. Click **Save**.

8. Reapply all access control policies with file policies that use the global malware whitelist.

After the policies apply, the system no longer performs malware cloud lookups on the file or blocks the file as malware.

Adding a File to the Whitelist Using the SHA-256 Value

LICENSE: Malware

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

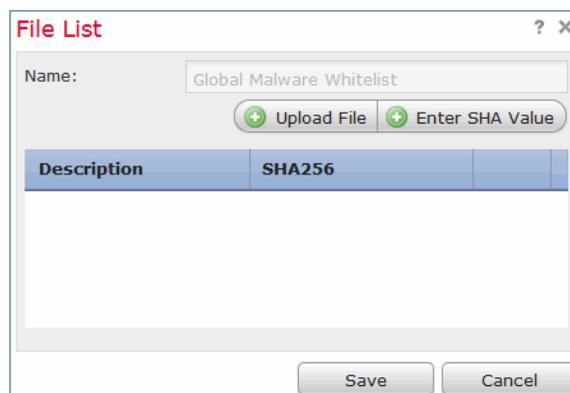
You can submit a file's SHA-256 value to add it to the global malware whitelist. You cannot add duplicate SHA-256 values.

TIP! Use **Show Full Text** in the context menu in the file or malware event views to view the full SHA-256 value for the file. You can copy the value from the full text view.

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

ACCESS: Admin/Network Admin

1. Select **Objects > Object Management**, then click **File List**.
2. Next to the global malware whitelist, click the edit icon (✎).
The File List pop-up window appears.



3. Click **Enter SHA Value**.

The pop-up window updates to include new fields.

Description	SHA256
-------------	--------

4. Type (or paste) the file's **SHA256** value.
5. Optionally, type a description for the file.
The system leaves the description blank if you do not supply the description.
6. Click **Add**.
The system calculates the SHA-256 value for the file and adds the file and the value to the list.
7. Click **Save**.
8. Reapply all access control policies with file policies that use the global malware whitelist.
After the policies apply, the system no longer performs malware cloud lookups on the file or blocks the file as malware.

Modifying Files on the Global Malware List

LICENSE: Malware

SUPPORTED DEVICES: Series 3, virtual

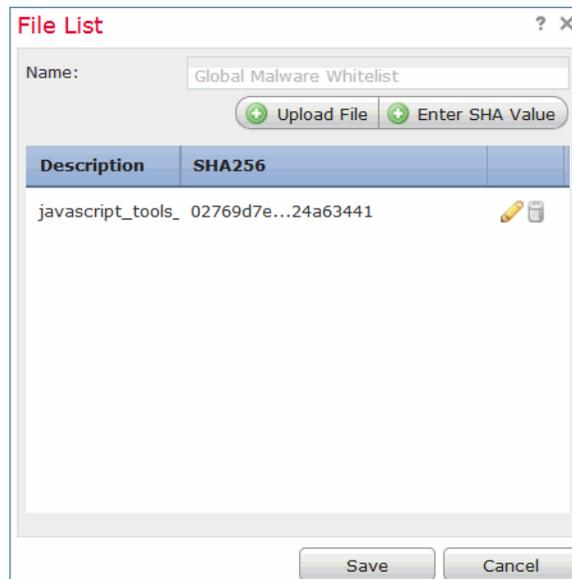
SUPPORTED DEFENSE CENTERS: Any except DC500

You can edit existing file entries on the global malware whitelist.

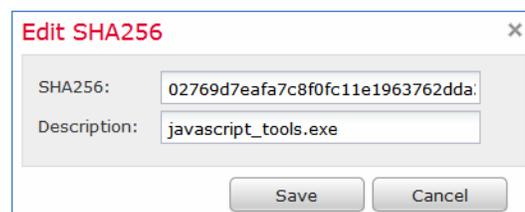
To edit a file on the global malware whitelist:

ACCESS: Admin/Network Admin

1. Select **Objects > Object Management**, then click **File List**.
2. Next to the global malware whitelist, click the edit icon (✎).
The File List pop-up window appears.



3. Next to the file you want to edit, click the edit icon (✎).
The Edit SHA pop-up window appears.



4. Update the **SHA256** value or **Description**.
5. Click **Save**.
The system updates the file entry in the list.

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

6. Click **Save**.

7. Reapply all access control policies with file policies that use the global malware whitelist.

After the policies apply, the system no longer performs malware cloud lookups on or blocks as malware any files you add. It resumes malware protection against any files you deleted from the list.

Working with Security Zones

LICENSE: Any

SUPPORTED DEVICES: Any

A *security zone* is a grouping of one or more inline, passive, switched, or routed interfaces that you can use to manage and classify traffic flow in various policies and configurations. The interfaces in a single zone may span multiple devices; you can also configure multiple zones on a single device. This allows you to divide the network into segments where you can apply various policies. You must assign each interface you configure to a security zone before it can handle traffic, and each interface can belong to only one zone.

In addition to using security zones to group interfaces, you can use zones in various places in the system's web interface, including access control policies, network discovery rules, and event searches. For example, you could write an access control rule that applies only to a specific source or destination zone, or restrict network discovery to traffic to or from a specific zone.

When you update a security zone object, the system saves a new revision of the object. As a result, if you have managed devices in the same security zone that have different revisions of the security zone object, you may log what appear to be duplicate connections. If you notice duplicate connection reporting, you can update all managed devices to use the same revision of the object. In the object manager, edit the security zone, remove all managed devices, save the object, re-add the managed devices, and save the object again. Then, reapply all affected device policies. For more information on applying device policies, see [Applying Changes to Devices](#) on page 211.

You create security zones in one of the following ways:

- The system creates security zones upon device registration, depending on the detection mode you selected for the device during its initial configuration. For example, the system creates a Passive zone in passive deployments, while in inline deployments the system creates External and Internal zones.
- You can create security zones on the fly while configuring interfaces on a managed device.
- You can create security zones using the object manager (**Objects > Object Management**).

The Security Zones page of the object manager lists the zones configured on your managed devices. The page also displays the type of interfaces in each zone, and

you can expand each zone to view which interfaces on which devices belong to each zone.

Name	Type
External	Inline
tamarix	
s1p2	
xiramati	
Internal	Inline

IMPORTANT! All interfaces in a security zone must be of the same type, that is, all inline, passive, switched, or routed. Further, after you create a security zone, you cannot change the type of interfaces it contains.

You cannot delete a security zone that is in use. After you add or remove interfaces from a zone, you must reapply the device configuration to the devices where the interfaces reside. You must also reapply the access control and network discovery policies that use the zone.

To add a security zone:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Objects > Object Management**.
The Object Management page appears.
2. Select **Security Zones**.

3. Click **Add Security Zone**.

The Security Zones pop-up window appears.

The screenshot shows a window titled "Security Zones" with a question mark icon in the top right corner. The window contains the following elements:

- Name:** A text input field.
- Type:** A dropdown menu with the text "Select Security Zone Type".
- Device > Interfaces:** A dropdown menu with the text "Select Device".
- Selected Interfaces:** An empty rectangular box for listing interfaces.
- Add:** A button located between the "Selected Interfaces" box and the "Device > Interfaces" dropdown.
- Save:** A button at the bottom right.
- Cancel:** A button at the bottom right, to the right of the "Save" button.

4. Type a **Name** for the zone. You can use any printable standard ASCII characters except curly braces ({}).
5. Select an interface **Type** for the zone.
After you create a security zone, you cannot change its type.
6. From the **Device > Interfaces** drop-down list, select a device that contains interfaces you want to add to the zone.
7. Select one or more interfaces.
Use the Shift and Ctrl keys to select multiple objects. If you have not yet configured interfaces on your managed devices, you can create an empty zone and add interfaces to it later; skip to step 10.
8. Click **Add**.
The interfaces you selected are added to the zone, grouped by device.
9. Repeat steps 6 through 8 to add interfaces on other devices to the zone.
10. Click **Save**.
The security zone is added.

CHAPTER 5

MANAGING DEVICES

The Sourcefire Defense Center is a key component in the Sourcefire 3D System. You can use the Defense Center to manage the full range of devices that comprise the Sourcefire 3D System, and to aggregate, analyze, and respond to the threats they detect on your network.

By using the Defense Center to manage devices, you can:

- configure policies for all your devices from a single location, making it easier to change configurations
- install various types of software updates on devices
- push health policies to your managed devices and monitor their health status from the Defense Center

The Defense Center aggregates and correlates intrusion events, network discovery information, and device performance data, allowing you to monitor the information that your devices are reporting in relation to one another, and to assess the overall activity occurring on your network.

For more information, see the following sections:

- [Management Concepts](#) on page 191 describes some of the features and limitations involved with managing your devices with a Defense Center.
- [Working in NAT Environments](#) on page 193 describes the principles of setting up the management of your devices in Network Address Translation environments.
- [Configuring High Availability](#) on page 194 describes how to set up two Defense Centers as a high availability pair to help ensure continuity of operations.

- [Working with Devices](#) on page 206 describes how to establish and disable connections between devices and your Defense Center. It also explains how to add, delete, and change the state of managed devices.
- [Configuring Remote Management](#) on page 213 describes how to establish and disable remote management of a managed device.
- [Managing Device Groups](#) on page 217 describes how to create device groups as well as how to add and remove devices from groups.
- [Clustering Devices](#) on page 220 describes how to establish and manage high availability between two managed devices.
- [Editing Device Configuration](#) on page 244 describes the device attributes you can edit and explains how to edit them.
- [Managing Stacked Devices](#) on page 237 describes how to create a stack of managed devices and how to remove devices from a stack.
- [Configuring Interfaces](#) on page 258 explains how to configure interfaces on your managed devices.

Management Concepts

You can use a Defense Center to manage nearly every aspect of a device's behavior. You need only one Defense Center to manage a device, though you can also use a second Defense Center as part of a high availability pair. The sections that follow explain some of the concepts you need to know as you plan your Sourcefire 3D System deployment:

- [What Can Be Managed by a Defense Center?](#) on page 191
- [Beyond Policies and Events](#) on page 192
- [Using Redundant Defense Centers](#) on page 193

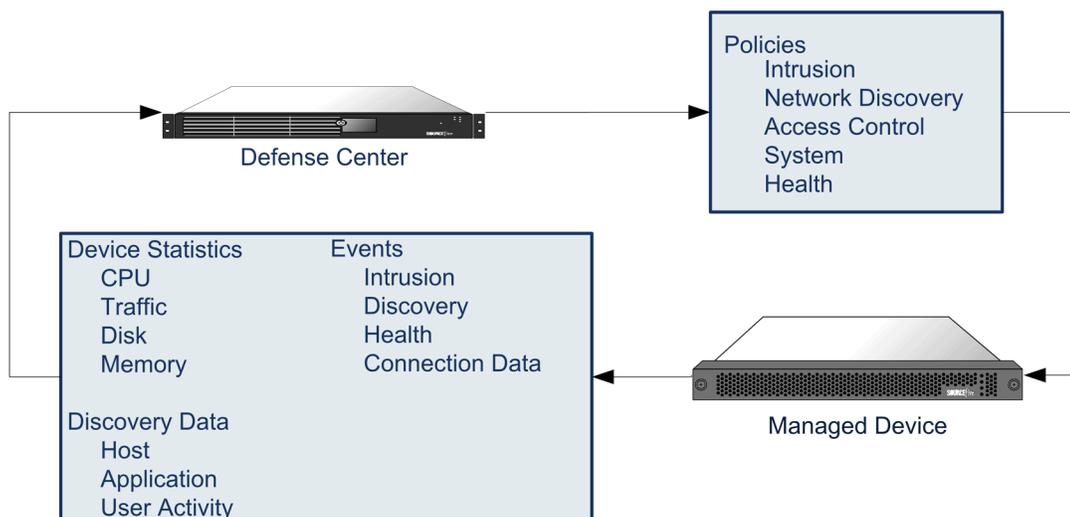
What Can Be Managed by a Defense Center?

You can use your Defense Center as a central management point in a Sourcefire 3D System deployment to manage Sourcefire-branded managed devices, including virtual devices.

IMPORTANT! Sourcefire recommends that you manage no more than three devices with the DC500 model Defense Center. For details on DC500 database limitations see the [Database Event Limits table](#) on page 1979.

When you manage a device, information is transmitted between the Defense Center and the device over a secure, SSL-encrypted TCP tunnel.

The following illustration lists what is transmitted between a Sourcefire Defense Center and its managed devices. Note that the types of events and policies that are sent between the appliances are based on the device type.



Beyond Policies and Events

LICENSE: Any

In addition to applying policies to devices and receiving events from them, you can also perform other device-related tasks on the Defense Center.

Backing Up a Device

If you are storing event data on your device in addition to sending it to the Defense Center, you can use the Defense Center's web interface to back up those events from the device. See [Backing up Your Managed Devices with a Defense Center](#) on page 2209 for more information.

Updating Devices

From time to time, Sourcefire releases updates to the Sourcefire 3D System, including:

- Sourcefire rule updates, which may contain new and updated intrusion rules
- vulnerability database updates
- geolocation updates
- software patches and updates

You can use the Defense Center to install an update on the devices it manages.

Using Redundant Defense Centers

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

You can set up two Defense Centers as a high availability pair. This ensures redundant functionality in case one of the Defense Centers fails. Policies, user accounts, and more are shared between the two Defense Centers. Events are automatically sent to both Defense Centers. See [Configuring High Availability](#) on page 194 for more information.

Working in NAT Environments

LICENSE: Control

Network address translation (NAT) is a method of transmitting and receiving network traffic through a router that involves reassigning the source or destination IP address as the traffic passes through the router. Typical applications using NAT enable multiple hosts on a private network to use a single public IP address to access the public network.

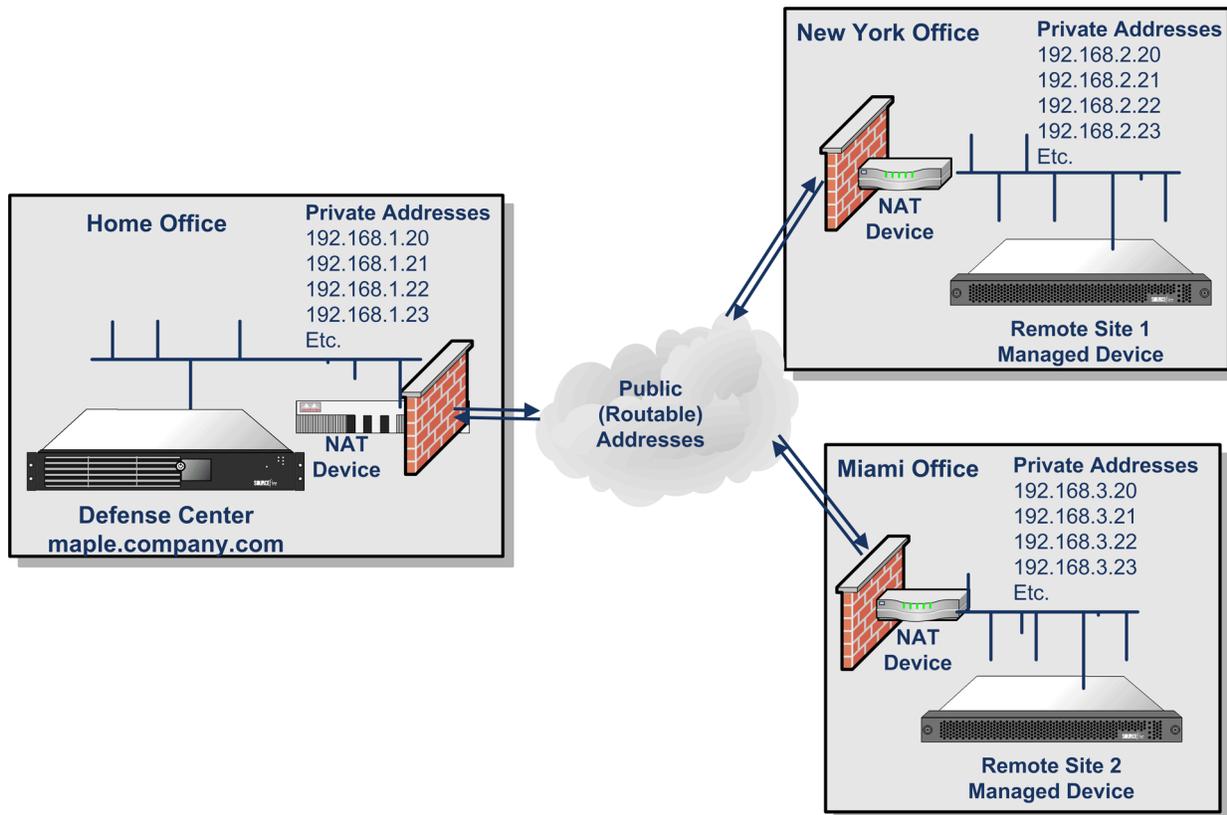
When you add a device to a Defense Center, you establish communications between the appliances. The information you need to establish communications depends on whether the environment uses NAT:

- In an environment without NAT, you need a registration key and the IP addresses or fully qualified domain names of both appliances.
- In an environment with NAT, you need a registration key and a unique NAT ID.

IMPORTANT! The NAT ID **must** be unique among all NAT IDs used to register devices to a Defense Center.

The following diagram shows a Defense Center managing two devices in a NAT environment. You can use the same registration key when adding both devices,

because registration keys do not have to be unique. However, you must use **unique** NAT IDs when adding the devices to the Defense Center.



Configuring High Availability

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

To ensure the continuity of operations, the high availability feature allows you to designate redundant Defense Centers to manage devices. Event data streams from managed devices to both Defense Centers and certain configuration elements are maintained on both Defense Centers. If one Defense Center fails, you can monitor your network without interruption using the other Defense Center.

WARNING! Because the system restricts some functionality to the primary Defense Center, if that appliance fails, you must switch the role of the secondary Defense Center to **Active & Primary**. See [Monitoring and Changing High Availability Status](#) on page 202.

See the following sections for more information about setting up high availability:

- [Using High Availability](#) on page 195 lists the configurations that are and are not shared when you implement high availability.
- [Guidelines for Implementing High Availability](#) on page 199 outlines guidelines you must follow if you want to implement high availability.
- [Setting Up High Availability](#) on page 200 explains how to specify primary and secondary Defense Centers.
- [Monitoring and Changing High Availability Status](#) on page 202 explains how to check the status of your linked Defense Centers and how to change the roles of the Defense Center if the primary Defense Center fails.
- [Disabling High Availability and Unregistering Devices](#) on page 204 explains how to permanently remove the link between linked Defense Centers.
- [Pausing Communication Between Paired Defense Centers](#) on page 205 explains how to pause communications between linked Defense Centers.
- [Restarting Communication Between Paired Defense Centers](#) on page 205 explains how to restart communications between linked Defense Centers.

Using High Availability

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

DC1500s and DC3500s support high availability configurations; DC750s and the virtual Defense Centers do not. Sourcefire **strongly** recommends that both Defense Centers in a high availability pair be the same model. Do **not** attempt to set up high availability between a Defense Center 1500 and a Defense Center 3500.

Although Defense Centers in high availability mode are designated *primary* and *secondary*, you can make policy or other changes to either Defense Center. However, Sourcefire recommends that you change configurations **only** on the primary Defense Center and that you keep your secondary Defense Center as a backup.

Defense Centers periodically update each other on changes to their configurations, and any change you make to one Defense Center should be applied on the other Defense Center within ten minutes. (Each Defense Center has a five-minute synchronization cycle, but the cycles themselves could be out of synchronization by as much as five minutes, so changes appear within two five-minute cycles.) During this ten-minute window, configurations may appear differently on the Defense Centers.

For example, if you create a policy on your primary Defense Center and apply it to a device that is also managed by your secondary Defense Center, the device could contact the secondary Defense Center before the Defense Centers contact each other. Because the device has a policy applied to it that the secondary Defense Center does not recognize, the secondary Defense Center displays a new policy with the name “unknown” until the Defense Centers synchronize.

Also, if you make conflicting policy or other changes to both Defense Centers within the same window between Defense Centers syncs, the last change you make takes precedence, regardless of the designations of the Defense Center as primary and secondary.

Before you establish a high availability pair, note the following prerequisites:

- Make sure both Defense Centers have a user account named `admin` with Administrator privileges. These accounts must use the same password.
- Make sure that other than the `admin` account, the two Defense Centers do not have user accounts with identical user names. Remove or rename one of the duplicate user accounts before you establish high availability.

Note that Defense Centers configured as a high availability pair do not need to be on the same trusted management network, nor do they have to be in the same geographic location.

For more information on which configurations are shared or not shared between members of a high availability pair, see:

- [Shared Configurations](#) on page 196
- [Health and System Policies](#) on page 197
- [Correlation Responses](#) on page 198
- [Licenses](#) on page 198
- [URL Filtering and Security Intelligence](#) on page 198
- [Sourcefire Cloud Connections and Malware Information](#) on page 199
- [User Agents](#) on page 199

Shared Configurations

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

Defense Centers in a high availability pair share the following information:

- user account attributes, authentication configurations, and custom user roles
- authentication objects for user accounts and user awareness, as well as the users and groups that are available to user conditions in access control rules
- custom dashboards
- custom workflows and tables
- device attributes, such as the device's host name, where events generated by the device are stored, and the group in which the device resides
- intrusion policies and their associated rule states
- file policies
- access control policies and their associated rules
- local rules

- custom intrusion rule classifications
- variable values and user-defined variables
- network discovery policies
- user-defined application protocol detectors and the applications they detect
- activated custom fingerprints
- host attributes
- network discovery user feedback, including notes and host criticality; the deletion of hosts, applications, and networks from the network map; and the deactivation or modification of vulnerabilities
- correlation policies and rules, compliance white lists, and traffic profiles
- change reconciliation snapshots and report settings
- intrusion rule updates
- geolocation updates
- vulnerability database updates

Health and System Policies

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

Health and system policies for Defense Centers and managed devices are shared in high availability pairs. Allow enough time to ensure that information about health policies, modules, blacklists, is synchronized on a newly activated Defense Center.

IMPORTANT! Although system policies are shared by Defense Centers in a high availability pair, they are not automatically applied. If you want identical system policies on both Defense Centers, apply the policy after it synchronizes.

Defense Centers in a high availability pair share the following system and health policy information:

- system policies
- system policy configurations (what policy is applied where)
- health policies
- health monitoring configurations (what policy is applied where)
- which appliances are blacklisted from health monitoring
- which appliances have individual health monitoring policies blacklisted

Correlation Responses

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

Although Defense Centers share correlation policies, rules, and responses, Defense Centers do not share the *associations* between correlation rules and their responses. This is to avoid launching duplicate responses when correlation policies are violated.

You must upload and install any custom remediation modules and configure remediation instances on your secondary Defense Center before remediations are available to associate with correlation policies. If the primary Defense Center fails, not only should you quickly associate your correlation policies with the appropriate responses and remediations on the secondary Defense Center, but you must also use the web interface on the secondary Defense Center to switch its role to **Active & Primary** to maintain continuity of operations. For more information on switching roles, see [Monitoring and Changing High Availability Status](#) on page 202. For more information about correlation responses, see [Creating Correlation Policies](#) on page 1513 and [Creating Remediations](#) on page 1607.

When you restore your primary Defense Center after a failure, if you created associations between rules or white lists and their responses and remediations on the secondary Defense Center, make sure you remove the associations so responses and remediations will only be generated by the primary Defense Center.

Licenses

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

Defense Centers in a high availability pair do **not** share licenses. You must add equivalent licenses to each member of the pair. For more information, see [Understanding Licensing](#) on page 2036.

URL Filtering and Security Intelligence

LICENSE: URL Filtering or Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

URL filtering and Security Intelligence configurations and information are synchronized between Defense Centers in a high availability deployment. However, only the primary Defense Center downloads URL category and reputation data and for updates to Security Intelligence feeds.

If the primary Defense Center fails, not only must you make sure that the secondary Defense Center can access the URL filtering cloud and any configured feed sites, but you must also use the web interface on the secondary Defense Center to switch its role to **Active & Primary**. For information on switching roles, see [Monitoring and Changing High Availability Status](#) on page 202.

Sourcefire Cloud Connections and Malware Information

LICENSE: Any or Malware

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

Although they share file policies and related configurations, Defense Centers in a high availability pair share neither Sourcefire cloud connections nor malware dispositions. To ensure continuity of operations, and to ensure that detected files' malware dispositions are the same on both Defense Centers, both primary and secondary Defense Centers must have access to the cloud. For more information, see [Working with Malware Protection and File Control](#) on page 1197.

User Agents

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

User Agents can connect to up to five Defense Centers at a time. You should connect agents to the primary Defense Center. If the primary Defense Center fails, you must make sure that any agents can communicate with the secondary Defense Center. See [Introduction to Network Discovery](#) on page 1250 for more information.

Guidelines for Implementing High Availability

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

To take advantage of high availability, you must follow the guidelines in the following sections.

Primary and Secondary Defense Center Requirements

You must designate one Defense Center as the primary Defense Center and one as the secondary.

Regardless of their designations as primary and secondary, both Defense Centers can be configured with policies, rules, managed devices, and so on before you set up high availability.

To avoid confusion, start with the secondary Defense Center in its original state. That is, you have not created or modified any policies, nor created any new rules, nor have you previously managed any devices with it. To make sure the secondary Defense Center is in its original state, restore it to factory defaults. Note that this also deletes event and configuration data from the Defense Center. For more information, see the *Sourcefire 3D System Installation Guide*.

IMPORTANT! You cannot configure a recurring task schedule on the inactive Defense Center. You must recreate the recurring task schedule on a newly activated Defense Center when it changes from inactive to active.

Version Requirements

Both Defense Centers must be running the same software and rule update version. Additionally, this software version must be the same or newer than the software version of managed devices.

Communication Requirements

By default, paired Defense Centers use port 8305/tcp for communications. You can change the port as described in [Changing the Management Port](#) on page 217.

The two Defense Centers do not need to be on the same network segment, but each of the Defense Centers must be able to communicate with the other and with the devices they share. That is, the primary Defense Center must be able to contact the secondary Defense Center at the IP address on the secondary Defense Center's own management interface, and vice versa. In addition, each Defense Center must be able to contact the devices it manages or the devices must be able to contact the Defense Center.

Setting Up High Availability

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

To use high availability, you must designate one Defense Center as the primary and another Defense Center of the same model as the secondary. For information about editing the remote management communications between the two appliances, see [Editing Remote Management](#) on page 216.

WARNING! Sourcefire recommends that you change configurations **only** on the primary Defense Center and that you use your secondary Defense Center as a backup.

Before you configure high availability, make sure you synchronize time settings between the Defense Centers you want to link. For details on setting time, see [Synchronizing Time](#) on page 1991.

Depending upon the number of policies and custom standard text rules they have, it may take up to 10 minutes before all the rules and policies appear on both Defense Centers. You can view the High Availability page to check the status of the link between the two Defense Centers. You can also monitor the Task Status to see when the process completes. See [Monitoring and Changing High Availability Status](#) on page 202.

If one of the Defense Centers in the high availability pair must be reimaged, disable the high availability link first. After you reimage the Defense Center, re-establish the high availability pair and the data will synchronize from the existing Defense Center to the newly added Defense Center. If a Defense Center cannot be reimaged (for example, the appliance has failed), contact Sourcefire Support.

To set up high availability for two Defense Centers:

ACCESS: Admin

1. Log into the Defense Center that you want to designate as the secondary Defense Center.
2. Select **System > Local > Registration**.
The Registration page appears.
3. Click **High Availability**.
The High Availability page appears.
4. Click the **secondary Defense Center** option.
The Secondary Defense Center Setup page appears.



The screenshot shows a form with three text input fields. The first field is labeled 'Primary DC Host *', the second is 'Registration Key *', and the third is 'Unique NAT ID'. Below the fields is a button labeled 'Register'.

5. Type the hostname or IP address of the primary Defense Center in the **Primary DC Host** text box.

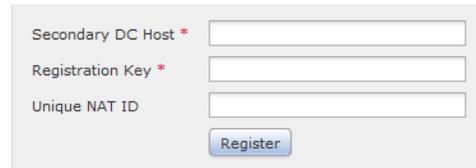
WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

Note that you can leave the **Primary DC Host** field empty if the management host does not have a routable address. In that case, use both the **Registration Key** and the **Unique NAT ID** fields.

6. Type a one-time-use registration key in the **Registration Key** text box
7. Optionally, in the **Unique NAT ID** field, type a unique alphanumeric registration ID that you want to use to identify the primary Defense Center. See [Working in NAT Environments](#) on page 193 for more information.
8. Click **Register**.
A success message appears, and the Peer Manager page appears, showing the current state of the secondary Defense Center.
9. Using an account with Admin access, log into the Defense Center that you want to designate as the primary.
10. Select **System > Local > Registration**.
The Registration page appears.
11. Click **High Availability**.
The High Availability page appears.

12. Click the **primary Defense Center** option.

The Primary Defense Center Setup page appears.



13. Type the hostname or IP address of the secondary Defense Center in the **Secondary DC Host** text box.

WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

14. Type the same one-time-use registration key in the **Registration Key** text box you used in step 6.
15. If you used a unique NAT ID on the secondary Defense Center, type the same registration ID that you used in step 7 in the **Unique NAT ID** text box.
16. Click **Register**.

A success message appears, and the Peer Manager page appears, showing the current state of the primary Defense Center.

Monitoring and Changing High Availability Status

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

After you have identified your primary and secondary Defense Centers, you can use one of them to view status information about the other, including:

- IP address
- product model
- operating system
- operation system version
- the local role (Active & Primary or Inactive & Secondary)
- time the Defense Centers last synchronized

You can also use the High Availability page to change the roles of the Defense Centers if the primary Defense Center fails. Because the system restricts the

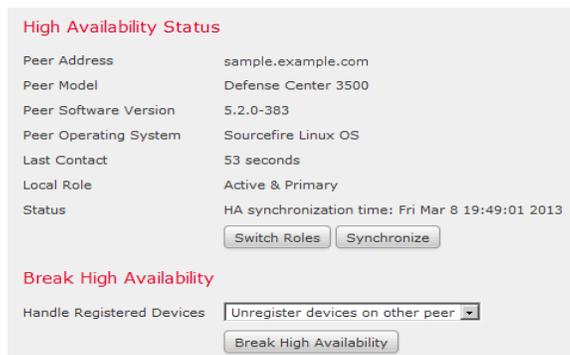
following functionality to the primary Defense Center, if that appliance fails, you must switch the role of the secondary Defense Center to **Active & Primary**:

- Updates to URL category and reputation data; see [URL Filtering and Security Intelligence](#) for more information.
- Updates to Security Intelligence feeds; see [URL Filtering and Security Intelligence](#) for more information.
- Associations between correlation rules and responses; see [Correlation Responses](#) for more information.

To check high availability status:

ACCESS: Admin

1. Log into one of the Defense Centers that you linked using high availability.
2. Select **System > Local > Registration**.
The Registration page appears.
3. Click **High Availability**.
The High Availability page appears.



The screenshot shows the 'High Availability Status' page. It contains a table of system information for the peer device, followed by control buttons and a section for breaking high availability.

High Availability Status	
Peer Address	sample.example.com
Peer Model	Defense Center 3500
Peer Software Version	5.2.0-383
Peer Operating System	Sourcefire Linux OS
Last Contact	53 seconds
Local Role	Active & Primary
Status	HA synchronization time: Fri Mar 8 19:49:01 2013

Below the table are two buttons: **Switch Roles** and **Synchronize**.

Under the heading **Break High Availability**, there is a dropdown menu for 'Handle Registered Devices' with the selected option 'Unregister devices on other peer', and a **Break High Availability** button.

4. Under **High Availability Status**, you can view the following information about the other Defense Center in the high availability pair:
 - the IP address
 - the model name
 - the software version
 - the operating system
 - the length of time since the last contact between the two Defense Centers
 - the local role (Active & Primary or Inactive & Secondary)
 - the time the Defense Centers last synchronized
 - the option to switch roles between the two Defense Centers

5. The two Defense Centers automatically synchronize within ten minutes (five minutes for each Defense Center) after any action that affects a shared feature. For example, if you create a new policy on one Defense Center, it is automatically shared with the other Defense Center within 5 minutes. However, if you want to synchronize the policy immediately, click **Synchronize**.

IMPORTANT! If you delete a device from a Defense Center configured in a high availability pair and intend to re-add it, Sourcefire recommends that you wait at least five minutes before adding the device back. This interval ensures that the high availability pair resynchronizes first. If you do not wait five minutes, it may take more than one synchronization cycle to add the device to both Defense Centers.

6. Click **Switch Roles** to change the local role from Active & Primary to Inactive & Secondary, or from Inactive & Secondary to Active & Primary.
The roles are switched between the two peers.
7. Click **Peer Manager** in the toolbar.
The Peer Manager page appears.

Host	Last Modified	Status	State	
10.10.10.10	2011-11-29 12:56:58	Registered	<input checked="" type="checkbox"/>	

You can view the following information:

- the IP address of the other Defense Center in the high availability pair
- the status, registered or unregistered, of the communications link
- the state, enabled or disabled, of the high availability pair

For information about editing the remote management communications between the two appliances, see [Editing Remote Management](#) on page 216.

Disabling High Availability and Unregistering Devices

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

If you want to remove one of the Defense Centers from a high availability pair, you must first disable the high availability link between them.

To disable a high availability pair:

ACCESS: Admin

1. Log into one of the Defense Centers in the high availability pair.
2. Select **System > Local > Registration**.
The Registration page appears.
3. Click **High Availability**.
The High Availability page appears.

4. Select one of the following options from the **Handle Registered Devices** drop-down list:
 - To control all the managed devices with the Defense Center where you are accessing this page, select **Unregister devices on the other peer**.
 - To control all the managed devices with the other Defense Center, select **Unregister devices on this peer**.
 - To stop managing the devices altogether, select **Unregister devices on both peers**.
5. Click **Break High Availability**.

After you answer the prompt **Do you really want to Break High Availability?** by selecting **OK**, high availability is disabled and any managed devices are deleted from the Defense Centers according to your selection.

You can enable high availability with a different Defense Center as described in [Setting Up High Availability](#) on page 200.

Pausing Communication Between Paired Defense Centers

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

If you want to temporarily disable high availability, you can disable the communications channel between the Defense Centers.

To disable the communications channel for a high availability pair:

ACCESS: Admin

1. Click **Peer Manager**.

The Peer Manager page appears.

Host	Last Modified	Status	State
10.10.10.10	2011-11-29 12:56:58	Registered	<input checked="" type="checkbox"/>

2. Click the slider to disable the communications channel between the two Defense Centers.

For information about editing the remote management communications between the two appliances, see [Editing Remote Management](#) on page 216.

Restarting Communication Between Paired Defense Centers

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

If you temporarily disabled high availability, you can enable the communications channel between the Defense Centers to restart high availability.

To enable the communications channel for a high availability pair:

ACCESS: Admin

1. Click **Peer Manager**.

The Peer Manager page appears.

Host	Last Modified	Status	State
10.10.10.10	2011-11-29 15:06:44	Registered	<input type="checkbox"/> 

2. Click the slider to enable the communications channel between the two Defense Centers.

For information about editing the remote management communications between the two appliances, see [Editing Remote Management](#) on page 216.

Working with Devices

LICENSE: Any

You can use the Defense Center to manage the full range of devices that are a part of the Sourcefire 3D System. When you manage a device, you set up a two-way, SSL-encrypted communication channel between the Defense Center and the device. The Defense Center uses this channel to send information to the device about how you want to analyze and manage your network traffic.

As the device evaluates the traffic, it generates events and sends them to the Defense Center using the same channel.

See the following sections for more information about managing devices:

- [Understanding the Device Management Page](#) on page 206
- [Adding Devices to the Defense Center](#) on page 208
- [Configuring Remote Management](#) on page 213
- [Managing Device Groups](#) on page 217
- [Clustering Devices](#) on page 220
- [Editing Device Configuration](#) on page 244
- [Configuring Interfaces](#) on page 258

Understanding the Device Management Page

LICENSE: Any

The Device Management page provides you with a range of information and options that you can use to manage your registered devices, device clusters, and

device groups. The page displays a list of all the devices currently registered on the Defense Center.

Name	License Type	Health Policy	System Policy	Access Control Policy
By Group [v] Add... [v]				
Ungrouped (4)				
birch 10.10.10.10 - 3D Sensor 3500 - v5.2.0	Protection, Control	katsura_health_policy	katsura_system_policy	Default Access Control

You can use the **sort-by** drop-down list to sort the appliance list according to your needs. Devices are displayed in the appliance list grouped by the category you select. You can sort by:

- Group (that is, device group); see [Managing Device Groups](#) on page 217 for more information
- Type (that is, the type of licenses applied to the device); see [Licensing the Sourcefire 3D System](#) on page 2036 for more information
- Model (that is, the model of the device being managed by the Defense Center)
- Health Policy; see [Using Health Monitoring](#) on page 2109 for more information
- System Policy; see [Managing System Policies](#) on page 1961 for more information
- Access Control Policy; see [Managing Access Control Policies](#) on page 456 for more information

For device groups, you can expand and collapse the list of devices in the group. The list appears collapsed by default.

See the Appliance List Fields table for more information about the appliance list.

Appliance List Fields

FIELD	DESCRIPTION
Name	A list of the hostname, IP address, device model, and software version for each device. The status icon to the left of the appliance indicates its current health status.
License Type	The licenses that are enabled on the managed device.
Health Policy	The currently applied health policy for the device. You can click the name of the health policy to view a read-only version of the policy. See Editing Health Policies on page 2147 for information about modifying an existing health policy.

Appliance List Fields (Continued)

FIELD	DESCRIPTION
System Policy	The currently applied system policy for the device. You can click the name of the system policy to view a read-only version of the policy. See Managing System Policies on page 1961 for more information.
Access Control Policy	A link to the currently applied access control policy. See Managing Access Control Policies on page 456.

See the following sections for more features accessible from the Device Management page:

- [Adding Devices to the Defense Center](#) on page 208
- [Managing Device Groups](#) on page 217
- [Clustering Devices](#) on page 220

Adding Devices to the Defense Center

LICENSE: Any

When you manage a device, you set up a two-way, SSL-encrypted communication channel between the Defense Center and the device. The Defense Center uses this channel to send information about how you want to analyze your network traffic to the device. As the device evaluates the traffic, it generates events and sends them to the Defense Center using the same channel.

Before you manage a device with a Defense Center, you must make sure that the network settings are configured correctly on the device. This is usually completed as part of the installation process. See [Configuring Network Settings](#) on page 2009 for more information.

Note that if you registered a Defense Center and a device using IPv4 and want to convert them to IPv6, you must delete and re-register the device.

You can select an access control policy to apply to a device as you register it to the Defense Center. If you attempt to apply a policy that requires a license that is not enabled on the device, the policy apply fails. Also, if the access control policy is associated with a network discovery policy, that policy also fails to apply as network discovery requires access control to properly apply to a managed device.

When you register a device cluster or device stack, although you can select licenses, these licenses cannot be applied upon device registration. This ensures that the cluster or stack is running the proper licenses to prevent it from entering a degraded state with mismatched licenses. After registration, you can evaluate the licenses in either the general properties (cluster) or stack properties (stack) of the Device Management page. For more information, see [Establishing Device Clusters](#) on page 223 or [Establishing Device Stacks](#) on page 239.

When you register a Series 2 device, although you can select licenses, any licenses you select are not applied upon device registration. Series 2 devices automatically have Protection capabilities, with the exception of Security Intelligence filtering. You cannot disable these capabilities, nor can you apply other licenses to a Series 2 device.

TIP! To modify the detailed configuration of a device, click the edit icon (✎) next to the device. See [Editing Device Configuration](#) on page 244 and [Configuring Interfaces](#) on page 258 for more information.

To add a device to a Defense Center:

ACCESS: Admin/Network Admin

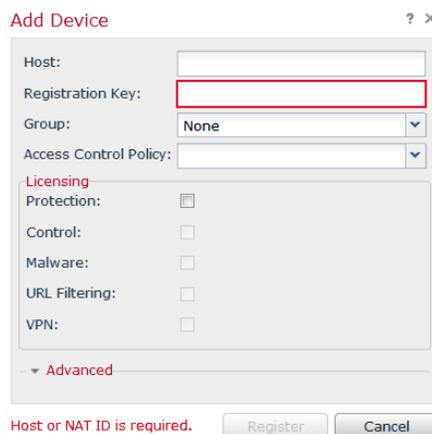
1. Configure the device to be managed by the Defense Center.

Use the procedure in [Configuring Remote Management](#) on page 213. After the device confirms communication with the Defense Center, the Pending Registration status appears.

Host	Last Modified	Status	State	
katsura	2011-12-06 16:51:05	Pending Registration	<input checked="" type="checkbox"/>	 

IMPORTANT! In some high availability deployments where network address translation (NAT) is used, you may also need to add the secondary Defense Center as a manager. For more information, contact Sourcefire Support.

2. Select **Devices > Device Management**.
The Device Management page appears.
3. From the **Add** drop-down menu, select **Add Device**.
The Add Device pop-up window appears.



4. In the **Host** field, type the IP address or the hostname of the device you want to add.

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

Note that in a NAT environment, you may not need to specify the IP address or host name of the device, if you already specified the IP address or host name of the Defense Center when you configured the device to be managed by the Defense Center. For more information, see [Working in NAT Environments](#) on page 193.

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

5. In the **Registration Key** field, type the same registration key that you used when you configured the device to be managed by the Defense Center.
6. Optionally, add the device to a device group by selecting the group from the **Group** drop-down list.

For more information about device groups, see [Managing Device Groups](#) on page 217.

7. From the **Access Control Policy** drop-down list, select an initial policy to apply to the device:
 - The **Default Access Control** policy blocks all traffic from entering your network.
 - The **Default Intrusion Prevention** policy allows all traffic that is also passed by the Balanced Security and Connectivity intrusion policy.
 - The **Default Network Discovery** policy allows all traffic, which is inspected by network discovery only.
 - You can select any existing user-defined access control policy.

For more information, see [Using Access Control Policies](#) on page 424.

8. Select licenses to apply to the device. Note that:
 - Control and Malware licenses require a Protection license; a URL Filtering license requires both Protection and Control licenses.
 - You cannot enable a Malware license on a pre-Version 5.1.1 device.
 - Although you can enable a Control license on a virtual device, virtual devices do **not** support fast-path rules, switching, routing, or clustering.
 - You cannot change the license settings on clustered devices.
 - For stacked devices, you enable or disable the licenses for the stack on the Stack page of the appliance editor.

- When you register a Series 2 device, any licenses you select are not applied upon device registration. Series 2 devices automatically have Protection capabilities, with the exception of Security Intelligence filtering. You cannot disable these capabilities, nor can you apply other licenses to a Series 2 device.

For more information, see [Licensing the Sourcefire 3D System](#) on page 2036.

9. If you used a NAT ID to identify the device when you configured it to be managed by the Defense Center, expand the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field.

10. To allow the device to transfer packets to the Defense Center, select the **Transfer Packets** check box.

This option is enabled by default. If you disable it, you completely prohibit packet transfer to the Defense Center.

11. Click **Register**.

The device is added to the Defense Center. Note that it may take up to two minutes for the Defense Center to verify the device's heartbeat and establish communication.

Applying Changes to Devices

LICENSE: Any

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

TIP! You can apply device changes from the Device Management page or from the Interfaces tab of the appliance editor.

To apply changes to the device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to apply changes, click the apply icon (👍).

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 browser window. For more information, see [Using the Device Management Revision Comparison Report](#) on page 212.

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:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the appliance where you want to apply changes, click the apply icon (✓).
The Apply Device Changes pop-up window appears. Note that the appliance must have unapplied changes or the apply icon remains disabled.
3. Click **View Changes**.
The Device Management Revision Comparison Report page appears in a new window.
4. Click **Previous** and **Next** to scroll through the differences between the current appliance configuration and the proposed appliance configuration.
5. Optionally, click **Comparison Report** to produce a PDF version of the report.

Deleting Devices

LICENSE: Any

If you no longer want to manage a device, you can delete it from the Defense Center. Deleting a device severs all communication between the Defense Center

and the device. To manage the device again at a later date, you must re-add it to the Defense Center.

IMPORTANT! If you delete a device from a Defense Center configured in a high availability pair and want to re-add it, Sourcefire recommends that you wait at least five minutes before re-adding it. This interval ensures that the high availability pair resynchronizes so that both Defense Centers recognize the deletion. If you do not wait five minutes, it may take more than one synchronization cycle to add the device to both Defense Centers.

To delete a device from the Defense Center:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management.**

The Device Management page appears.

2. Next to the device you want to delete, click the delete icon ().

When prompted, confirm that you want to delete the device. Communication between the device and the Defense Center is discontinued and the device is deleted from the Device Management page. If the device has a system policy that causes it to receive time from the Defense Center via NTP, the device reverts to local time management.

Configuring Remote Management

LICENSE: Any

Before you can manage one Sourcefire 3D 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 Defense Center. After you configure remote management, you can use the managing appliance's web interface to add the managed appliance to your deployment.

To enable communications between two appliances, you must provide a way for the appliances to recognize each other. There are three criteria the Sourcefire 3D 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 required registration key that identifies the connection
- an optional unique alphanumeric NAT ID that can help the Sourcefire 3D System establish communications in a NAT environment

The NAT ID **must** be unique among all NAT IDs used to register managed appliances. For more information, see [Working in NAT Environments](#) on page 193.

When you register a managed device to a Defense Center, the access control policy you select applies to the device. In addition, the network discovery policy on the Defense Center automatically 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, causing the network discovery policy apply to fail as well. If, for example, you select an access control policy with an intrusion policy as the default action, and do not enable the Protection license, both the access control policy and the network discovery policy apply fail.

To configure remote management of the local appliance:

ACCESS: Admin

1. On the web interface for the appliance you want to manage, select **System > Local > Registration**.

The Remote Management page appears.

WARNING! Sourcefire **strongly** recommends that you not change the value for the management port. If you change it, you must also change it for all appliances in your deployment that need to communicate with each other. For more information, see [Changing the Management Port](#) on page 217.

2. Click **Add Manager**.

The Add Remote Management page appears.

Host	Last Modified	Status	State
oak.example.com	2011-10-07 10:10:45	Registered	Enabled (Disable)

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 Sourcefire 3D System uses the NAT ID you will provide later to identify the remote manager on the managed appliance's web interface.

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

4. In the **Registration Key** field, type the registration key that you want to use to set up communications between appliances.
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.
6. Click **Save**.
After the appliances confirm that they can communicate with each other, the Pending Registration status appears.
7. Use the managing appliance's web interface to add this appliance to your deployment.

For more information, see [Adding Devices to the Defense Center](#) on page 208.

IMPORTANT! When enabling remote management of a device, in some high availability deployments that use NAT, you may also need to add the secondary Defense Center as a manager. For more information, contact Sourcefire 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 Sourcefire 3D System. Although you can use the hostname as the display name of the appliance, entering a different display name does not change the hostname.

TIP! You can click the slider to block the communication channel between the managed device and the Defense Center. Click the slider to enable the communication channel.

To edit remote management:

ACCESS: Admin

1. Select **System > Local > Registration**.

The Remote Management page appears.

Host	Last Modified	Status	State
oak.example.com	2011-10-07 10:10:45	Registered	Enabled (Disable)

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

The Edit Remote Management page appears.

Name * oak.example.com
Host * oak
Save Cancel

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

5. Click **Save**.
Your changes are saved.

Changing the Management Port

LICENSE: Any

Sourcefire 3D System appliances communicate using a two-way, SSL-encrypted communication channel, which by default is on port 8305.

Although Sourcefire **strongly** recommends that you keep the default setting, if the management port conflicts with other communications on your network, you can choose a different port. Usually, changes to the management port are made during installation of the Sourcefire 3D System.

WARNING! If you change the management port, you must change it for all appliances in your deployment that need to communicate with each other.

To change the management port:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.
2. Click **Network**.
The Network Settings page appears.
3. In the **Remote Management Port** field, enter the port number that you want to use.
4. Click **Save**.
The management port is changed.
5. Repeat this procedure for every appliance in your deployment that must communicate with this appliance.

Managing Device Groups

LICENSE: Any

The Defense Center allows you to group devices so you can easily apply policies and install updates on multiple devices. You can expand and collapse the list of devices in the group. The list appears collapsed by default.

See the following sections for more information:

- [Adding Device Groups](#) on page 218
- [Editing Device Groups](#) on page 218
- [Deleting Device Groups](#) on page 219

Adding Device Groups

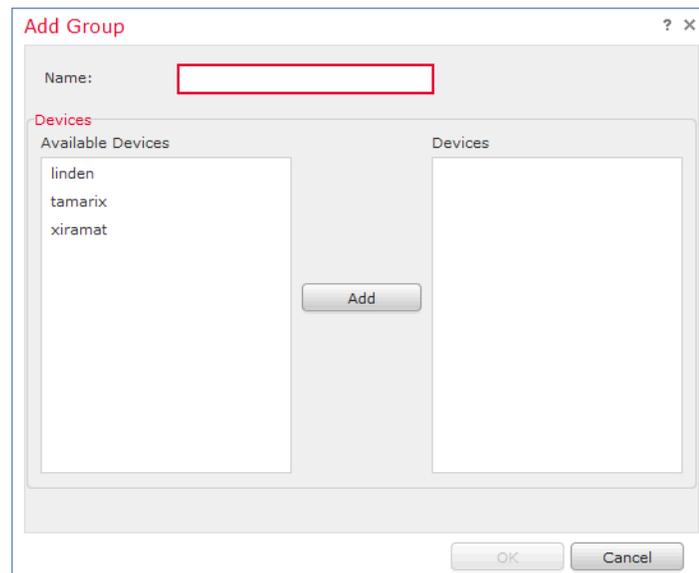
LICENSE: Any

The following procedure explains how to add a device group so you can easily apply policies and install updates on multiple devices.

To create a device group and add devices to it:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. From the **Add** drop-down menu, select **Add Group**.
The Add Group pop-up window appears.



3. In the **Name** field, type the name of the group.
4. Under **Available Devices**, select one or more appliances to add to the device group. Use Ctrl or Shift while clicking to select multiple appliances.
5. Click **Add** to include the selected appliances in the device group.
6. Click **OK**.
The device group is added.

Editing Device Groups

LICENSE: Any

You can change the set of devices that reside in any device group. You must remove an appliance from its current group before you can add it to a new group.

Moving an appliance to a new group does not change its policy to the policy previously applied to the group. To change the device's policy, you must apply a new policy to the device or device group.

To edit a device group:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device group you want to edit, click the edit icon ().
The Edit Group pop-up window appears.
3. Optionally, in the **Name** field, type a new name for the group
4. Under **Available Devices**, select one or more appliances to add to the device group. Use Ctrl or Shift while clicking to select multiple appliances.
5. Click **Add** to include the selected appliances in the device group.
6. To remove selected appliances from the device group, click the delete icon ().
7. Click **OK**.
The changes to the device group are saved.

Deleting Device Groups

LICENSE: Any

If you delete a device group that contains devices, the devices are moved to the Ungrouped category on the Device Management page. They are not deleted from the Defense Center.

To delete a device group:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device group you want to delete, click the delete icon ().
3. When prompted, confirm that you want to delete the device group.
The device group is deleted.

Clustering Devices

LICENSE: Control

SUPPORTED DEVICES: Series 3

With device clustering (also called device high availability), you can establish redundancy of networking functionality and configuration data between two peer devices or two peer device stacks. See [Managing Stacked Devices](#) on page 237 for more information about stacking devices.

You achieve configuration redundancy by clustering two peer devices or two peer device stacks as a single logical system for policy applies, system updates, and registration. The system automatically synchronizes other configuration data.

Clustering Requirements

Before you can configure a device cluster, both devices or device stack primary members must be the same model and have identical copper or fiber interfaces. Both devices or device stacks must also be running the same software and have the same licenses. Device stacks must have identical hardware configurations; for example, you can cluster a 3D8290 with a 3D8290. If the devices are targeted by NAT policies, both peers must have the same NAT policy. After you cluster the devices, you cannot change the license options for individual clustered devices, but you can change the license for the entire cluster. See [Establishing Device Clusters](#) on page 223 for more information.

Clustering Failover and Maintenance Mode

With a device cluster, the system fails over either manually or automatically. You manually trigger failover by placing one of the clustered devices or stacks in maintenance mode. For more information about maintenance mode, see [Placing a Clustered Device into Maintenance Mode](#) on page 229.

Automatic failover occurs when the health of the active device or stack becomes compromised or during a system update. If the health of the backup device or stack becomes similarly compromised, the system does not fail over and enters a degraded state. The system also does not fail over when one of the devices or device stacks is in maintenance mode. Note that disconnecting the stacking cable from an active stack sends that stack into maintenance mode. Shutting down the secondary device in an active stack also sends that stack into maintenance mode.

Applying Policies and Updates

When you apply policies, you apply them to the device cluster instead of the individual devices or stacks. If the policy fails, the system does not apply it to either device or stack. The policy first applies to the active device or stack and then the backup, so that the cluster always has one peer handling network traffic.

Clustered devices receive updates as a single entity rather than individual devices or stacks. When the update is started, the system first applies it to the backup device or stack, which goes into maintenance mode until any necessary

processes restart and the device begins processing traffic again. The system then applies the update to the active device or stack, which follows the same process.

Achieving Redundancy Without Clustering Devices

In most cases, you can achieve Layer 3 redundancy without clustering devices by using the Sourcefire Redundancy Protocol (SFRP). SFRP allows devices to act as redundant gateways for specified IP addresses. With network redundancy, you configure two devices or stacks to provide identical network connections, ensuring connectivity for other hosts on the network. For more information about SFRP, see [Configuring SFRP](#) on page 304.

You determine how to configure device high availability depending on your Sourcefire 3D System deployment: passive, inline, routed, or switched. You can also deploy your system in multiple roles at once. Of the four deployment types, only passive deployments require that you cluster devices or stacks to provide redundancy. You can establish network redundancy for the other deployment types with or without device clusters. The following sections provide a brief overview of high availability in each deployment type.

Passive Deployment Redundancy

Passive interfaces are generally connected to tap ports on central switches, which allows them to analyze all of the traffic flowing across the switch. If multiple devices are connected to the same tap feed, the system generates events from each of the devices. When clustered, devices act as either active or backup, which allows the system to analyze traffic even in the event of a system failure while also preventing duplicate events.

Inline Deployment Redundancy

Because an inline set has no control over the routing of the packets being passed through it, it must always be active in a deployment. Therefore, redundancy relies on external systems to route traffic correctly. You can configure redundant inline sets with or without device clusters.

To deploy redundant inline sets, you configure the network topology so that it allows traffic to pass through only one of the inline sets while preventing circular routing. If one of the inline sets fails, the surrounding network infrastructure detects the loss of connectivity to the gateway address and adjusts the routes to send traffic through the redundant set.

Routed Deployment Redundancy

Hosts in an IP network must use a well-known gateway address to send traffic to different networks. Establishing redundancy in a routed deployment requires that routed interfaces share the gateway addresses so that only one interface handles traffic for that address at any given time. To accomplish this, you must maintain an equal number of IP addresses on a virtual router. One interface advertises the address. If that interface goes down, the backup interface begins advertising the address.

In non-clustered devices, you use SFRP to establish redundancy by configuring gateway IP addresses shared between multiple routed interfaces. You can configure SFRP with or without device clusters. You can also establish redundancy using dynamic routing such as OSPF or RIP.

Switched Deployment Redundancy

You establish redundancy in a switched deployment using the Spanning Tree Protocol (STP). STP is a protocol that manages the topology of bridged networks. It is specifically designed to allow redundant links to provide automatic backup for switched interfaces without configuring backup links. Devices in a switched deployment rely on STP to manage traffic between redundant interfaces. Two devices connected to the same broadcast network receive traffic based on the topology calculated by STP. See [Configuring Advanced Virtual Switch Settings](#) on page 291 for more information about enabling STP.

IMPORTANT! Sourcefire strongly recommends that you enable STP when configuring a virtual switch that you plan to deploy in a device cluster.

See the following sections for more information about clustering devices and stacks:

- [Establishing Device Clusters](#) on page 223
- [Editing Device Clusters](#) on page 224
- [Configuring Individual Devices in a Cluster](#) on page 225
- [Configuring Individual Device Stacks in a Cluster](#) on page 226
- [Configuring Interfaces on a Clustered Device](#) on page 227
- [Switching the Active Peer in a Cluster](#) on page 228
- [Placing a Clustered Device into Maintenance Mode](#) on page 229
- [Replacing a Device in a Clustered Stack](#) on page 229
- [Establishing Clustered State Sharing](#) on page 230
- [Troubleshooting Clustered State Sharing](#) on page 233
- [Separating Clustered Devices](#) on page 236
- [Configuring SFRP](#) on page 304
- [Configuring HA Link Interfaces](#) on page 262

Establishing Device Clusters

LICENSE: Control

SUPPORTED DEVICES: Series 3

Before you establish a device cluster, you must meet the following prerequisites:

- Configure interfaces on each device or each primary device in a stack.
- Each device or device stack primary member that you include in the cluster must be the same model and have identical copper or fiber interfaces.
- Both devices or device stacks must have normal health status, run the same software, and have the same licenses. See [Using the Health Monitor](#) on page 2163 for more information. In particular, the devices cannot have hardware failures that would cause them to enter maintenance mode and trigger a failover.
- You cannot mismatch devices and stacks in a cluster. You must cluster single devices with single devices or device stacks with device stacks that have identical hardware configurations. For example, you can cluster a 3D8290 with a 3D8290.
- If the devices are targeted by NAT policies, both peers must have the same NAT policy.

When establishing a device cluster, you designate one of the devices or stacks as active and the other as backup. The system applies a merged configuration to the clustered devices. If there is a conflict, the system applies the configuration from the device or stack you designated as active.

After you cluster the devices, you cannot change the license options for individual clustered devices, but you can change the license for the entire cluster. See [Editing Device Clusters](#) on page 224 for more information. If there are interface attributes that need to be set on switched interfaces or routed interfaces, the system establishes the cluster, but sets it to a pending status. After you configure the necessary attributes, the system completes the device cluster and sets it to a normal status.

After you establish clustered pair, the system treats the peer devices or stacks as a single device on the Device Management page. Device clusters display the cluster icon () in the appliance list. Any configuration changes you make are synchronized between the clustered devices. The Device Management page displays which device or stack in the cluster is active, which changes after manual or automatic failover. See [Placing a Clustered Device into Maintenance Mode](#) on page 229 for more information about manual failover.

Removing registration of a device cluster from a Defense Center removes registration from both devices or stacks. You remove a device cluster from the Defense Center as you would an individual managed device. See [Deleting Devices](#) on page 212 for more information.

You can then register the cluster on another Defense Center. To register clustered single devices, you add remote management to the active device in the cluster

and then add that device to the Defense Center, which adds the entire cluster. To register clustered stacked devices, you add remote management to the primary device of the either stack and then add that device to the Defense Center, which adds the entire cluster. See [Adding Devices to the Defense Center](#) on page 208 for more information.

After you establish a device cluster, you can configure a high availability link interface, as explained in [Configuring HA Link Interfaces](#) on page 262.

To cluster devices or device stacks:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

2. From the **Add** drop-down menu, select **Add Cluster**.

The Add Cluster pop-up window appears.



3. In the **Name** field, type the name of the cluster.

You may enter alphanumeric characters and special characters, with the exception of the following characters, which are invalid: +, (,), {, }, #, &, \, <, >, ?, ', and "

4. Select the **Active** device or stack for the cluster.
5. Select the **Backup** device or stack for the cluster.
6. Click **Cluster**.

The device cluster is added. This process takes a few minutes as the process synchronizes system data.

Editing Device Clusters

LICENSE: Control

SUPPORTED DEVICES: Series 3

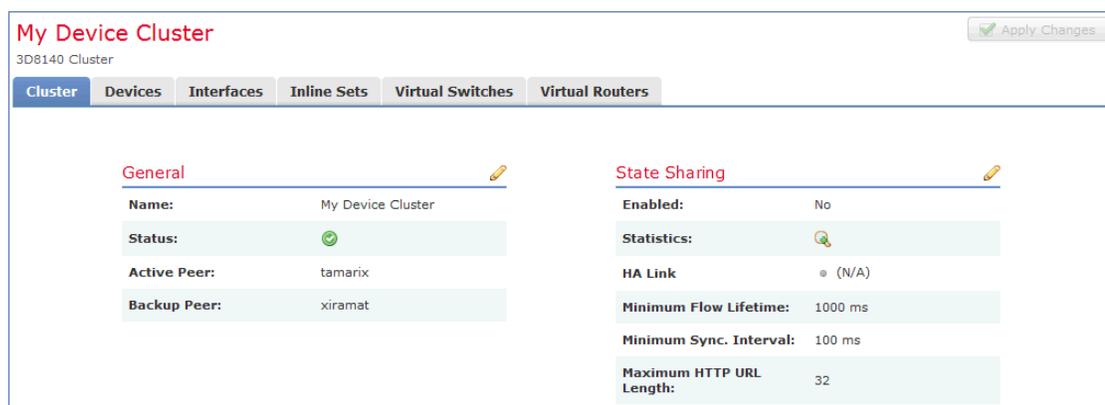
After you establish a device cluster, most changes you make to the device configuration also change the configuration of the entire cluster.

You can view the status of the cluster by hovering your pointer over the status icon in the General section. You can also view which device or stack is the active peer and backup peer in the cluster.

To edit a device cluster:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device cluster where you want to edit the configuration, click the edit icon (✎).
The Cluster page appears.



3. Use the sections on the Cluster page to make changes to the clustered configuration as you would a single device configuration.

See the following sections for more information:

- [Editing Assigned Device Names](#) on page 244
- [Enabling and Disabling Device Licenses](#) on page 246
- [Establishing Clustered State Sharing](#) on page 230
- [Editing Advanced Device Settings](#) on page 251

Configuring Individual Devices in a Cluster

LICENSE: Control

SUPPORTED DEVICES: Series 3

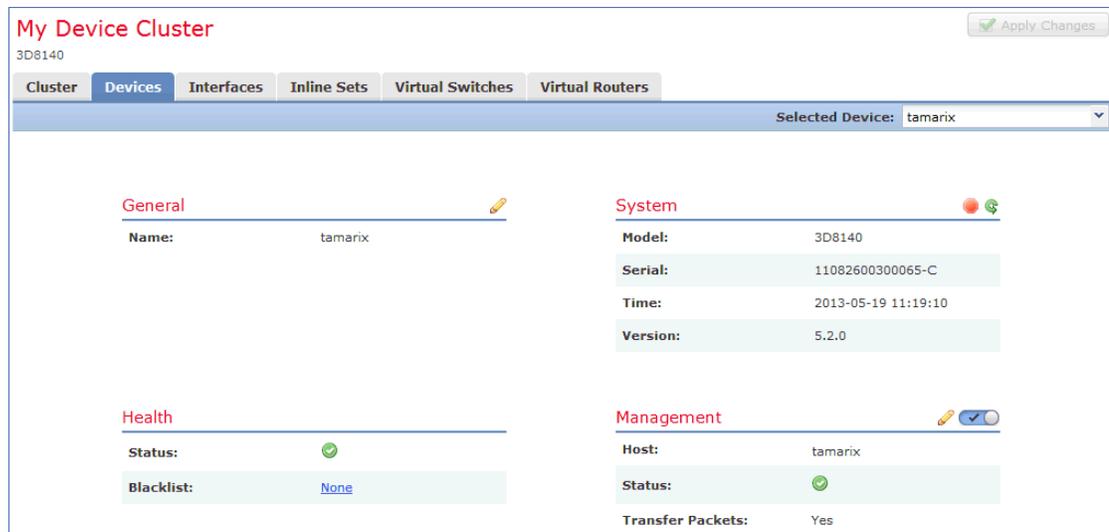
After you establish a device cluster, you can still configure some attributes for each device within the cluster. You can make changes to a clustered device just as you would to a single device.

To configure an individual device in a cluster:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.

2. Next to the device cluster where you want to edit the configuration, click the edit icon (✎).
The Cluster page appears.
3. Click **Devices**.
The Devices page appears.



4. From the **Selected Device** drop-down list, select the device you want to modify.
5. Use the sections on the Devices page to make changes to the individual clustered device as you would a single device.
See the following sections for more information:
 - [Editing Assigned Device Names](#) on page 244
 - [Editing Device System Settings](#) on page 247
 - [Viewing the Health of a Device](#) on page 248
 - [Editing Device Management Settings](#) on page 249

Configuring Individual Device Stacks in a Cluster

LICENSE: Control

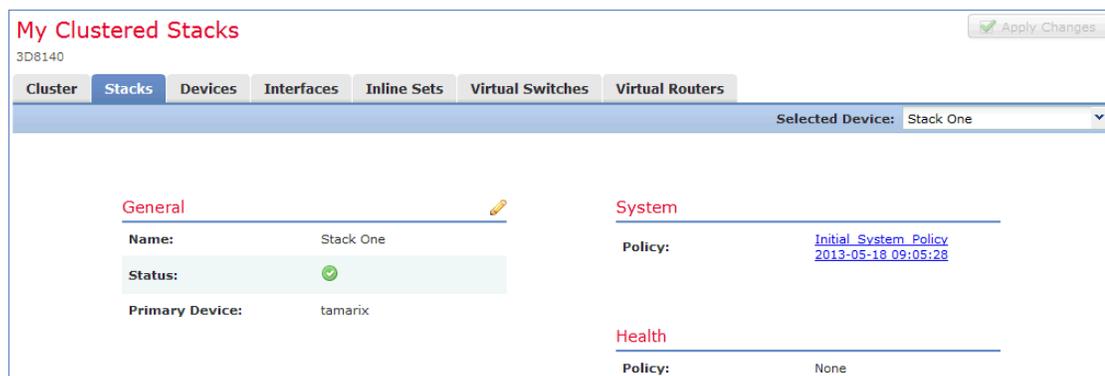
SUPPORTED DEVICES: Series 3

After you cluster a pair of stacked devices, the system limits the stack attributes that you can edit. You can edit the name of a stack in a clustered stack. In addition, you can edit the network configuration of the stack, as described in [Configuring Interfaces on a Clustered Device](#) on page 227.

To edit the name of a stack in a cluster:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device cluster where you want to edit the configuration, click the edit icon (✎).
The Cluster page appears.
3. Click **Stacks**.
The Stacks page appears.
From the **Selected Device** drop-down list, select the stack you want to modify.



4. Next to the General section, click the edit icon (✎).
The General pop-up window appears.
5. In the **Name** field, type a new assigned name for the stack.
You may enter alphanumeric characters and special characters, with the exception of the following characters, which are invalid: +, (,), {, }, #, &, \, <, >, ?, ', and ".
?
6. Click **Save**.
The new name is saved. Note that your changes do not take effect until you apply the stack configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring Interfaces on a Clustered Device

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can configure interfaces on individual devices in a cluster. However, you must also configure an equivalent interface on the peer device in the cluster. For clustered stacks, you configure identical interfaces on the primary devices of the

stacks. When you configure virtual routers, you select the stack where you want to configure the routers. See [Configuring Virtual Routers](#) on page 306 for more information.

The Interfaces page of a clustered device includes the hardware and interfaces views that you find on an individual device. See [Configuring Interfaces](#) on page 258 for more information.

To configure interfaces on a clustered device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device cluster where you want to configure interfaces, click the edit icon (✎).
The Cluster page appears.
3. Click **Interfaces**.
The Interfaces page appears.



4. From the **Selected Device** drop-down list, select the device you want to modify.
5. Configure interfaces as you would on an individual device. See [Configuring Interfaces](#) on page 258 for more information.

Switching the Active Peer in a Cluster

LICENSE: Control

SUPPORTED DEVICES: Series 3

After you establish a device cluster, you can manually switch the active and backup peer devices or stacks.

To switch the active peer in a cluster:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.

2. Next to the device cluster that you want to change the active peer, click the switch active peer icon ()
The Switch Active Peer pop-up window appears.
3. Click **Yes** to immediately make the backup device the active device in the cluster. Click **No** to cancel and return to the Device Management page.

Placing a Clustered Device into Maintenance Mode

LICENSE: Control

SUPPORTED DEVICES: Series 3

After you establish a cluster, you can manually trigger failover by placing one of the clustered devices or stacks into maintenance mode to perform maintenance on the devices. In maintenance mode, the system administratively takes down all interfaces except for the management interface. After maintenance is completed, you can re-enable the device to resume normal operation.

IMPORTANT! You cannot place both members of a cluster into maintenance mode at the same time.

To place a clustered device into maintenance mode:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the clustered device you want to place in maintenance mode, click the toggle maintenance mode icon ()
The Confirm Maintenance Mode pop-up window appears.
3. Click **Yes** to confirm maintenance mode or click **No** to cancel.
4. Click the toggle maintenance mode icon () again to bring the device out of maintenance mode.

Replacing a Device in a Clustered Stack

LICENSE: Control

SUPPORTED DEVICES: Series 3

After you place a stack that is a cluster member into maintenance mode, you can replace a secondary device in the stack for another device. You can only select devices that are not currently stacked or clustered. The new device must follow the same guidelines for establishing a device stack. See [Establishing Device Stacks](#) on page 239.

To replace a device in a clustered stack:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the stack member you want to place into maintenance mode, click the toggle maintenance mode icon ()
The Confirm Maintenance Mode pop-up window appears.
3. Click **Yes** to confirm maintenance mode or click **No** to cancel.
4. Click the replace device icon ()
The Replace Device pop-up window appears.
5. Select the **Replacement Device** from the drop-down list.
6. Click **Replace** to replace the device or click **Cancel** to keep the current device and return to the Device Management page.
7. Click the toggle maintenance mode icon () again to bring the stack immediately out of maintenance mode.
You do not need to reapply the device configuration.

Establishing Clustered State Sharing

LICENSE: Control

SUPPORTED DEVICES: Series 3

Clustered state sharing allows clustered devices or clustered stacks to synchronize as much state as necessary, so that if either device or stack fails, the other peer can take over with no interruption to traffic flow. Without state sharing, the following features may not fail over properly:

- Strict TCP enforcement
- Unidirectional access control rules
- Blocking persistence

Note, however, that enabling state sharing slows system performance.

You must configure and enable HA link interfaces on both devices or the primary stacked devices in the cluster before you can configure clustered state sharing. 3D8250 devices require a 10G HA link, while other model devices require a 1G HA link. See [Configuring HA Link Interfaces](#) on page 262 for more information.

Strict TCP Enforcement

SUPPORTED DEVICES: Series 3

When you enable strict TCP enforcement for a domain, the system drops any packets that are out of order on TCP sessions. For example, the system drops

non-SYN packets received on an un-established connection. With state sharing, devices in the cluster allow TCP sessions to continue after failover without having to reestablish the connection, even if strict TCP enforcement is enabled. You can enable strict TCP enforcement on inline sets, virtual routers, and virtual switches.

Unidirectional Access Control Rules

If you have configured unidirectional access control rules, network traffic may match a different access control rule than intended when the system reevaluates a connection midstream after failover. For example, consider if you have a policy containing the following two access control rules:

Rule 1: Allow from 192.168.1.0/24 to 192.168.2.0/24

Rule 2: Block all

Without state sharing, if an allowed connection from 192.168.1.1 to 192.168.2.1 is still active following a failover and the next packet is seen as a response packet, the system denies the connection. With state sharing, a midstream pickup would match the existing connection and continue to be allowed.

Blocking Persistence

While many connections are blocked on the first packet based on access control rules or other factors, there are cases where the system allows some number of packets through before determining that the connection should be blocked. With state sharing, the system immediately blocks the connection on the peer device or stack as well.

When establishing clustered state sharing, you can configure the following options:

Enabled

Click the check box to enable state sharing. Clear the check box to disable state sharing.

Minimum Flow Lifetime

Specify the minimum time (in milliseconds) for a session before the system sends any synchronization messages for it. You can use any integer from 0 to 65535. The system does not synchronize any sessions that have not met the minimum flow lifetime, and the system synchronizes only when a packet is received for the connection.

Minimum Sync. Interval

Specify the minimum time (in milliseconds) between update messages for a session. You can use any integer from 0 to 65535. The minimum synchronization interval prevents synchronization messages for a given connection from being sent more frequently than the configured value after the connection reaches the minimum lifetime.

Maximum HTTP URL Length

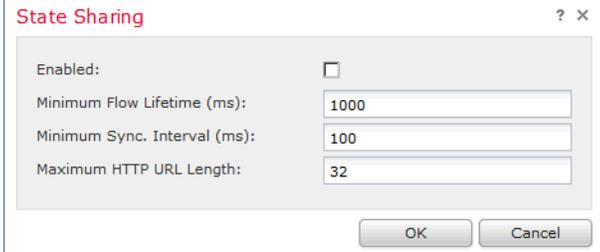
Specify the maximum characters for the URL the system synchronizes between the clustered devices. You may use any integer from 0 to 225.

IMPORTANT! Sourcefire recommends that you use the default values, unless your deployment presents a good reason to change them. Decreasing the values allows increased clustered peer readiness, while increasing the values allows better performance.

To establish clustered state sharing:

ACCESS: Admin/Network Admin

1. Configure HA link interfaces for each device in the cluster.
See [Configuring HA Link Interfaces](#) on page 262 for more information.
2. Select **Devices > Device Management**.
The Device Management page appears.
3. Next to the device cluster you want to edit, click the edit icon (✎).
The Cluster page appears.
4. Next to the State Sharing section, click the edit icon (✎).
The State Sharing pop-up window appears.



Enabled:	<input type="checkbox"/>
Minimum Flow Lifetime (ms):	1000
Minimum Sync. Interval (ms):	100
Maximum HTTP URL Length:	32

5. Configure the state sharing, as described earlier in this section.
6. Click **OK**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Troubleshooting Clustered State Sharing

LICENSE: Control

SUPPORTED DEVICES: Series 3

After you enable state sharing, you can view the following information about the configuration in the State Sharing section of the Cluster page:

- The HA link interface that is being used and its current link state
- Detailed synchronization statistics for troubleshooting issues

The state sharing statistics are primarily counters for different aspects of the clustered synchronization traffic sent and received, along with some other error counters. In addition, you can view the latest system logs for each device in the cluster.

See the following sections for more information about the statistics you can view for each device and how you can use them to troubleshoot your clustered state sharing configuration.

Messages Received (Unicast)

Messages received are the number of cluster synchronization messages received from the clustered peer.

The value should be close to the number of messages sent by the peer. During active use, the values may not match, but should be close. If traffic stops, the values should become stable and the messages received will match the messages sent.

For troubleshooting, you should view both the messages received and the messages sent, compare the rate of increase, and make sure the values are close. The sent value on each peer should be incrementing at approximately the same rate as the received value on the opposite peer.

Contact Support if the received messages stop incrementing or increment slower than the messages sent by the peer.

Packets Received

The system batches multiple messages into single packets in order to decrease overhead. The Packets Received counter displays the total number of these data packets, as well as other control packets that have been received by a device.

The value should be close to the number of packets sent by the peer device. During active use, the values may not match, but should be close. Because the number of messages received should be close and incrementing at the same rate as the number of messages sent by the peer, the number of packets received should have the same behavior.

For troubleshooting, you should view both the packets received and the messages sent, compare the rate of increase, and make sure the values are increasing at the same rate. If the sent value on the clustered peer is incrementing, the received value on the device should also increase at the same rate.

Contact Support if the received packets stop incrementing or increment slower than the messages sent by the peer.

Total Bytes Received

Total bytes received are the number of bytes that make up the packets received by the peer.

The value should be close to the number of bytes sent by the other peer. During active use, the values may not match, but should be close.

For troubleshooting, you should view both the total bytes received and the messages sent, compare the rate of increase, and make sure the values are increasing at the same rate. If the sent value on the clustered peer is incrementing, the received value on the device should also increase at the same rate.

Contact Support if the received bytes stop incrementing or increment slower than the messages sent by the peer.

Protocol Bytes Received

Protocol bytes received are the number of bytes of protocol overhead received, which includes everything but the payload of session state synchronization messages.

The value should be close to the number of bytes sent by the peer. During active use, the values may not match, but should be close.

For troubleshooting, you should view the total bytes received to discover how much actual state data is being shared in comparison to protocol data. If the protocol data is a large percentage of the data being sent, you can adjust the minimum sync interval.

Contact Support if the protocol bytes received increment at a similar rate to the total bytes received. Protocol bytes received should be minimal in relation to the total bytes received.

Messages Sent

Messages sent are the number of cluster synchronization messages sent to the clustered peer.

This data is useful in comparison to the number of messages received. During active use, the values may not match, but should be close.

For troubleshooting, you should view both the messages received and the messages sent, compare the rate of increase, and make sure the values are close.

Contact Support if the messages sent increment at a similar rate to the total bytes received.

Bytes Sent

Bytes sent are the total number of bytes sent that make up the cluster synchronization messages sent to the peer.

This data are useful in comparison to the number of messages received. During active use, the values may not match, but should be close. The number of bytes received on the peer should be close to, but not more than this value.

Contact Support if the total bytes received is not incrementing at about the same rate as the bytes sent.

Tx Errors

Tx errors are the number of memory allocation failures the system encounters when trying to allocate space for messages to be sent to the clustered peer.

This value should be zero at all times on both peers. Contact Support if this number is not zero or if the number steadily increases, which indicates the system has encountered an error where it cannot allocate memory.

Tx Overruns

Tx overruns are the number of times the system attempts and fails to place a message into the transit queue.

This value should be zero at all times on both peers. When the value is not zero or is steadily increasing, it indicates that the system is sharing too much data across the HA link that cannot be sent quickly enough.

You should increase the HA link MTU if it was previously set below the default value of 9918. You can change the minimum flow lifetime and minimum synchronization interval settings to reduce the amount of data shared across the HA link to prevent the number from incrementing.

Contact Support if this value persists or continues to increase.

Recent Logs

The system log displays the most recent clustered synchronization messages. The log should not display any ERROR or WARN messages. It should remain comparable between the peers, such as the same number of sockets being connected.

However, the data displayed may be opposite in some instances, for example, one peer reports that it received a connection from the other peer and references different IP addresses. The log provides a comprehensive view of the clustered state sharing connection, and any errors within the connection.

Contact Support if the log displays an ERROR or WARN message, or any message that does not appear to be purely informational.

To view clustered state sharing statistics:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device cluster you want to edit, click the edit icon (✎).
The Cluster page for the device cluster appears.
3. In the State Sharing section, click the view statistics icon (📊).
The State Sharing Statistics pop-up window appears.

	Active Peer	Backup Peer
Device	xiramat	tamarix
Messages Received (Unicast)	0	0
Packets Received	1307	1308
Total Bytes Received	10456	10464
Protocol Bytes Received	10456	10464
Messages Sent	0	0
Packets Sent	1442	1307
Bytes Sent	11536	10456
TX Errors	0	0
TX Overruns	0	0
Recent Logs	View	View

4. Optionally, select a **Device** to view if your cluster is composed of device stacks.
5. Optionally, click **Refresh** to update the statistics.
6. Optionally, click **View** to view the latest data log for each clustered device.

Separating Clustered Devices

LICENSE: Control

SUPPORTED DEVICES: Series 3

When you break device clustering, the active device or stack retains full deployment functionality. The backup device or stack loses its interface

configurations and fails over to the active device or stack, unless you choose to leave the interface configurations active, in which case the backup device or stack resumes normal operation. Breaking a cluster always removes the configuration of passive interfaces on the backup devices. Any devices in maintenance mode resume normal operation upon breaking the cluster.

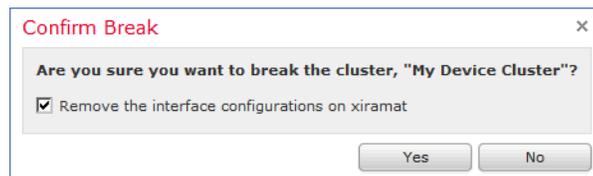
To separate a clustered device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

2. Next to the device cluster you want to break, click the break cluster icon (). The Confirm Break pop-up window appears.



3. Optionally, select the check box to remove the interface configurations on the backup device or stack, which means all interfaces except for the management interface are administratively taken down.
4. Click **Yes**.
The device cluster is separated.

Managing Stacked Devices

LICENSE: Any

SUPPORTED DEVICES: 3D8140, 3D8200 family, 3D9900

You can increase the amount of traffic inspected on a network segment by using devices in the following stacked configurations:

- two 3D8140s
- two 3D9900s
- up to four 3D8250s
- a 3D8260 (a primary device and a secondary device)
- a 3D8270 (a primary device with 40G capacity and two secondary devices)
- a 3D8290 (a primary device with 40G capacity and three secondary devices)

For more information about stacked configurations, see the *Sourcefire 3D System Installation Guide*.

When you establish a stacked configuration, you combine the resources of each stacked device into a single, shared configuration.

You designate one device as the *primary* device, where you configure the interfaces for the entire stack. You designate the other devices as *secondary*. Secondary devices must not be currently sensing any traffic and must not have link on any interface.

Connect the primary device to the network segment you want to analyze in the same way you would configure a single device. See [Configuring Interfaces](#) on page 258 for more information. Connect the secondary devices to the primary device using the stacked device cabling instructions found in the *Sourcefire 3D System Installation Guide*.

All devices in the stacked configuration must have the same hardware, run the same software version, and have the same licenses. If the devices are targeted by NAT policies, both the primary and secondary device must have the same NAT policy. See [Managing NAT Policies](#) on page 391 for more information. You must apply updates to the entire stack from the Defense Center. If an update fails on one or more devices in the stack, the stack enters a mixed-version state. You cannot apply policies to or update a stack in a mixed-version state. To correct this state, you can break the stack or remove individual devices with different versions, update the individual devices, then reestablish the stacked configuration. After you stack the devices, you can change the licenses only for the entire stack at once.

After you establish the stacked configuration, the devices act like a single, shared configuration. If the primary device fails, no traffic is passed to the secondary devices. Health alerts are generated indicating that the stacking heartbeat has failed on the secondary devices. See [Using Health Monitoring](#) on page 2109 for more information.

If a secondary device fails, the primary device continues to sense traffic, generate alerts, and send traffic to all secondary devices. On failed secondary devices, traffic is dropped. A health alert is generated indicating loss of link.

You can use a device stack as you would a single device in your deployment, with a few exceptions. If you have clustered devices, you cannot stack a device cluster or a device in a clustered pair. See [Clustering Devices](#) on page 220 for more information. You also cannot configure NAT on a device stack.

IMPORTANT! If you use eStreamer to stream event data from stacked devices to an external client application, collect the data from each device and ensure that you configure each device identically. The eStreamer settings are not automatically synchronized between stacked devices.

See the following sections for more information:

- [Establishing Device Stacks](#) on page 239
- [Editing Device Stacks](#) on page 241
- [Configuring Individual Devices in a Stack](#) on page 242
- [Separating Stacked Devices](#) on page 243

Establishing Device Stacks

LICENSE: Any

SUPPORTED DEVICES: 3D8140, 3D8200 family, 3D9900

You can increase the amount of traffic inspected on a network segment by stacking two fiber-based 3D9900s, two 3D8140 devices, up to four 3D8250s, a 3D8260, a 3D8270, or a 3D8290, and using their combined resources in a single, shared, configuration. Before you begin, you must:

- decide which unit will be the primary device.
- cable the units properly before designating the primary/secondary relationship. For information about cabling, see the *Sourcefire 3D System Installation Guide*.

IMPORTANT! If you have clustered devices, you cannot stack a device cluster or a device in a clustered pair. However, you can cluster a device stack. See [Clustering Devices](#) on page 220 for more information.

After you establish a device stack, the system treats the devices as a single device on the Device Management page. Device stacks display the stack icon

() in the appliance list.

Removing registration of a device stack from a Defense Center also removes registration from both devices. You delete stacked devices from the Defense Center as you would a single managed device; you can then register the stack on another Defense Center. You only need to register one of the stacked devices on the new Defense Center for the entire stack to appear. See [Deleting Devices](#) on page 212 and [Adding Devices to the Defense Center](#) on page 208 for more information.

After you establish the device stack, you cannot change which devices are primary or secondary unless you break and reestablish the stack. However, you can add secondary devices to an existing stack of two or three 3D8250s, a 3D8260, or a 3D8270 up to the limit of four 3D8250s in a stack. For additional devices, the primary device in the stack must have the necessary stacking NetMods for additional cabled devices. For example, if you have a 3D8260 where the primary only has a single stacking NetMod, you cannot add another secondary device to this stack. You add secondary devices to an existing stack in the same manner that you initially establish a stacked device configuration.

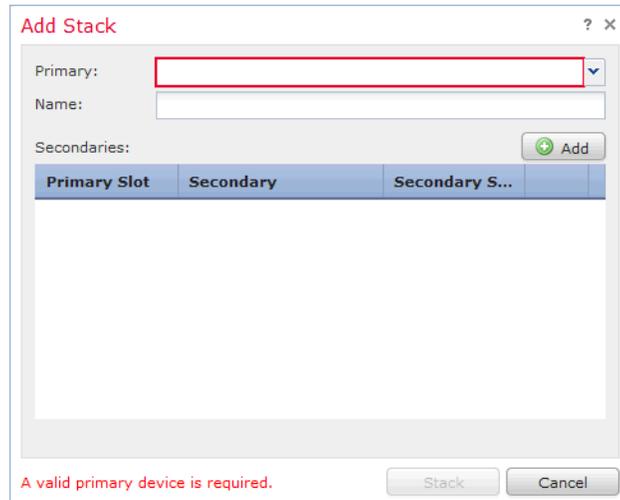
To establish a stacked device configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

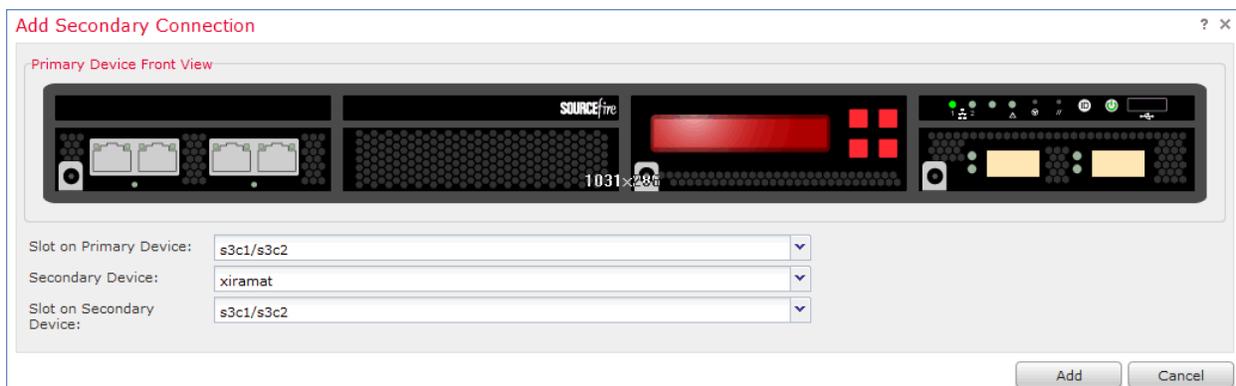
2. From the **Add** drop-down menu, select **Add Stack**.
The Add Stack pop-up window appears.



3. From the **Primary** drop-down list, select the device that you cabled for primary operation.

IMPORTANT! If you edit a device that is not cabled as the primary device, you cannot perform the next series of steps.

4. In the **Name** field, type the name of the stack. You may enter alphanumeric characters and special characters, with the exception of the following characters, which are invalid: +, (,), {, }, #, &, \, <, >, ?, ', and ".
5. Click **Add** to select the devices you want to form a stack with.
The Add Secondary Connection pop-up window appears. The following graphic displays the primary device front view for a 3D8140.



6. From the **Slot on Primary Device** drop-down list, select the stacking network module that connects the primary device to the secondary device.

7. From the **Secondary Device** drop-down list, select the device you cabled for secondary operation.

IMPORTANT! All devices in a stack must be of the same hardware model: 3D9900 with 3D9900, 3D8140 with 3D8140, or 3D8250 with 3D8250. You can add secondary 3D8250s to a properly configured and cabled 3D8260 or 3D8270.

8. From the **Slot on Secondary Device** drop-down list, select the stacking network module that connects the secondary device to the primary device.
9. Click **Add**.
The Add Stack window reappears with the new secondary device included.
10. Optionally, repeat steps 5 through 9 if you are adding secondary devices to an existing stack of 3D8250s, a 3D8260, or a 3D8270.
11. Click **Stack**.
The device stack is established or the additional secondary devices are added. Note that this process takes a few minutes as the process synchronizes system data.

Editing Device Stacks

LICENSE: Any

SUPPORTED DEVICES: 3D8140, 3D8200 family, 3D9900

After you establish a device stack, most changes you make to the device configuration also change the configuration of the entire stack. On the Stack page of the appliance editor, you can make changes to the stack configuration as on the Device page of a single device.

You can change the display name of the stack, enable and disable licenses, view system and health policies, configure automatic application bypass, and set up fast-path rules.

See the following sections for more information:

- [Editing Assigned Device Names](#) on page 244
- [Enabling and Disabling Device Licenses](#) on page 246
- [Editing Advanced Device Settings](#) on page 251

To edit a stacked configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.

2. Next to the stacked device where you want to edit the configuration, click the edit icon (✎).
The Stack page for that device appears.
3. Use the sections on the Stack page to make changes to the stacked configuration as you would a single device configuration.
See the following sections for more information:
 - [Editing Assigned Device Names](#) on page 244
 - [Enabling and Disabling Device Licenses](#) on page 246
 - [Editing Advanced Device Settings](#) on page 251

Configuring Individual Devices in a Stack

LICENSE: Any

SUPPORTED DEVICES: 3D8140, 3D8200 family, 3D9900

After you establish a device stack, you can still configure some attributes for only one device within the stack. On the Devices page of the appliance editor, you can make changes to a device configured in a stack as on the Device page of a single device.

You can change the display name of a device, view system settings, shut down or restart a device, view health information, and edit device management settings.

See the following sections for more information:

- [Editing Assigned Device Names](#) on page 244
- [Editing Device System Settings](#) on page 247
- [Viewing the Health of a Device](#) on page 248
- [Editing Device Management Settings](#) on page 249

To configure an individual device in a stack:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the stacked device where you want to edit the configuration, click the edit icon (✎).
The Stack page for that device appears.
3. Click **Devices**.
The Devices page appears.
4. From the **Selected Device** drop-down list, select the device you want to modify.

5. Use the sections on the Devices page to make changes to the individual stacked device as you would a single device.

See the following sections for more information:

- [Editing Assigned Device Names](#) on page 244
- [Editing Device System Settings](#) on page 247
- [Viewing the Health of a Device](#) on page 248
- [Editing Device Management Settings](#) on page 249

Configuring Interfaces on a Stacked Device

LICENSE: Any

SUPPORTED DEVICES: 3D8140, 3D8200 family, 3D9900

Except for the management interface, you configure stacked device interfaces on the Interfaces page of the primary device in the stack. You can select any device in the stack to configure the management interface. See [Configuring the Management Interface](#) on page 260 for more information.

The Interfaces page of a Series 3 stacked device includes the hardware and interfaces views that you find on an individual device. The interfaces page of a 3D9900 does not include these views. See [Configuring Interfaces](#) on page 258 for more information.

To configure interfaces on a stacked device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the stacked device where you want to configure interfaces, click the edit icon (✎).
The Stack page for that device appears.
3. Click **Interfaces**.
The Interfaces page appears.
4. From the **Selected Device** drop-down list, select the device you want to modify.
5. Configure interfaces as you would on an individual device. See [Configuring Interfaces](#) on page 258 for more information.

Separating Stacked Devices

LICENSE: Any

SUPPORTED DEVICES: 3D8140, 3D8200 family, 3D9900

If you no longer need to use a stacked configuration for your devices, you can break the stack and separate the devices.

To separate stacked devices:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device stack you want to break, click the break stack icon ().
The Confirm Break pop-up window appears.

TIP! To remove a secondary device from a stack of three or more 3D8250 devices without breaking the stack, click the remove from stack icon (). Removing the secondary device causes a brief disruption of traffic inspection, traffic flow, or link state as the system reconfigures the stack for operation without the extra device.

3. Click **Yes**.
The device stack is separated.

Editing Device Configuration

LICENSE: Any

The Device page of the appliance editor displays detailed device configuration and information. It also allows you to make changes to some parts of device configuration, such as enabling and disabling licenses, shutting down and restarting a device, modifying management, and setting up fast-path rules.

See the following sections for more information:

- [Editing Assigned Device Names](#) on page 244
- [Enabling and Disabling Device Licenses](#) on page 246
- [Editing Device System Settings](#) on page 247
- [Viewing the Health of a Device](#) on page 248
- [Editing Device Management Settings](#) on page 249
- [Understanding Advanced Device Settings](#) on page 250

Editing Assigned Device Names

LICENSE: Any

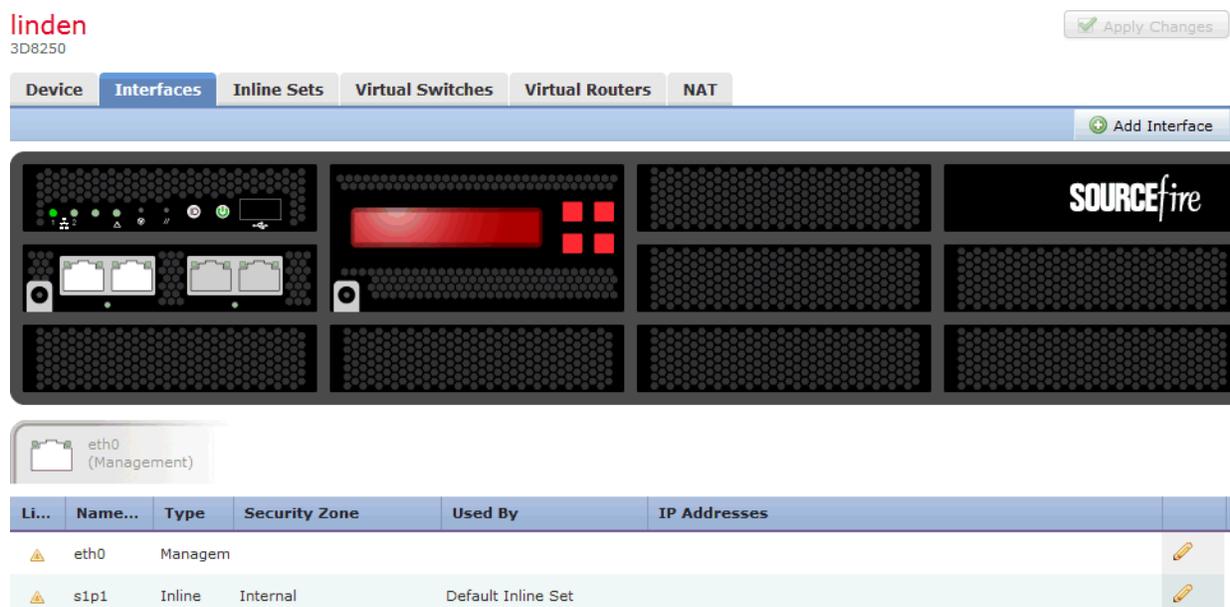
The General section of the Device tab shows the device name, which you can change.

To edit the assigned name of a device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to edit the assigned name, click the edit icon (✎).

The Interfaces page for that device appears. The following graphic shows the Interfaces page for a Series 3 device.



3. Click **Device**.
The Device page appears.

TIP! For stacked devices, you edit the assigned device name for the stack on the Stack page of the appliance editor. You can edit the assigned device name for an individual device on the Devices page of the appliance editor.

4. Next to the General section, click the edit icon (✎).
The General pop-up window appears.



5. In the **Name** field, type a new assigned name for the device. You may enter alphanumeric characters and special characters, with the exception of the following characters, which are invalid: +, (,), {, }, #, &, \, <, >, ?, ', and ".
6. Click **Save**.
The new name is saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Enabling and Disabling Device Licenses

LICENSE: Any

SUPPORTED DEVICES: Series 3, virtual

You can enable licenses on your device if you have available licenses on your Defense Center. Note that:

- Control and Malware licenses require a Protection license; URL Filtering VPN licenses require both Protection and Control licenses.
- You cannot enable a Malware license on a pre-Version 5.1.1 device.
- You cannot enable a VPN license on a virtual device.
- Although you can enable a Control license on a virtual device, virtual devices do **not** support fast-path rules, switching, routing, or clustering.
- You cannot change the license settings on clustered devices.
- Because Series 2 devices automatically have Protection capabilities, with the exception of Security Intelligence filtering, you cannot disable these capabilities, nor can you apply other licenses to a Series 2 device.

For more information, see [Licensing the Sourcefire 3D System](#) on page 2036.

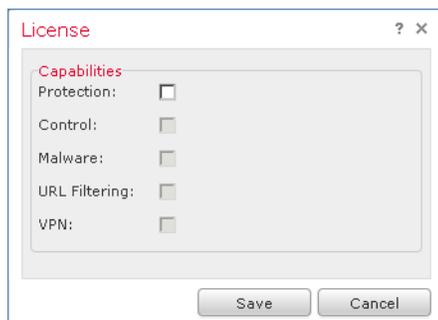
To enable or disable device licenses:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to enable or disable licenses, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Device**.
The Devices tab appears.

TIP! For stacked devices, you enable or disable the licenses for the stack on the Stack page of the appliance editor.

- Next to the **License** section, click the edit icon (✎).
The License pop-up window appears.



- You have the following options:
 - To enable a license, select the check box next to the license name.
 - To disable a license, clear the check box next to the license name.
- Click **Save**.
The changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Editing 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.

System Section Table Fields

FIELD	DESCRIPTION
Model	The model name and number for the managed device.
Serial	The serial number of the chassis of the managed device.
Time	The current system time of the device.
Version	The version of the software currently installed on the managed device.
Policy	A link to the system policy currently applied to the managed device.

You can also shut down or restart the device.

To shut down and restart a managed device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device that you want to restart, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Device**.
The Devices tab appears.

TIP! For stacked devices, you shut down or restart an individual device on the Devices page of the appliance editor.

4. To shut down the device, click the shut down device icon (🔴).
5. When prompted, confirm that you want to shut down the device.
You are returned to the Device Management page.
6. To restart the device, click the restart device icon (🔄).
7. When prompted, confirm that you want to restart the device.
The device is restarted.

Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Viewing the Health of a Device

LICENSE: Any

The Health section of the Device tab displays health-related information. You can view an icon showing the current health status of the managed device. You can also click the icon to navigate to the Health Monitor page for that device. See [Interpreting Health Monitor Status](#) on page 2165 for more information.

You can click the **Policy** link to view a read-only version of the currently applied health policy. See [Editing Health Policies](#) on page 2147 for more information.

You can also click the **Blacklist** link to go to the Health Blacklist page, where you can enable and disable health blacklist modules. See [Blacklisting a Health Policy Module](#) on page 2158 for more information.

Editing Device Management Settings

LICENSE: Any

The Management section of the Device tab displays a list of remote management information, as described in the following table:

Management Section Table Fields

FIELD	DESCRIPTION
Host	The current management hostname or IP address of the device.
Status	Whether local event storage is enabled for the device.
Transfer Packets	Whether packet data is transferred to the Defense Center.

You can use the Management section to specify the management host and regenerate the virtual IP. You can also specify management options regarding local event storage and whether packet data should be transferred to the Defense Center.

TIP! You can click the slider to block the communication channel between the Defense Center and the managed device. Click the slider to enable the communication channel.

To modify management options on a device:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to modify management options, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Device**.
The Devices tab appears.

TIP! For stacked devices, you modify management options on an individual device on the Devices page of the appliance editor.

- Next to the **Management** section, click the edit icon (✎).
The Management pop-up window appears.

- In the **Host** field, enter the name or IP address of the management host.
- Clear the **Enabled** check box to disable management of the device. Select the check box to enable management.
- Select the **Transfer Packets** check box to allow packet data to be stored on the Defense Center with events. Clear the check box to prevent the managed device from sending packet data with the events.
- Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Understanding Advanced Device Settings

LICENSE: Any

SUPPORTED DEVICES: feature dependent

The Advanced section of the Device tab displays a table of advanced configuration settings, as described in the Advanced Section Table Fields table.

Advanced Section Table Fields

FIELD	DESCRIPTION	SUPPORTED DEVICES
Application Bypass	The state of Automatic Application Bypass on the device.	Any
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.	Any
Inspect Local Router Traffic	Whether the device inspects traffic received on routed interfaces that is destined for itself, such as ICMP, DHCP, and OSPF traffic.	Series 3
Fast-Path Rules	The number of fast-path rules that have been created on the device.	8000 Series, 3D9900

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

- [Automatic Application Bypass](#) on page 251
- [Editing Advanced Device Settings](#) on page 251
- [Configuring Fast-Path Rules](#) on page 253

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.

The AAB option is off by default. 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.

Typically, you use Rule Latency Thresholding in the intrusion policy to fast-path packets after the latency threshold value is exceeded. Rule Latency Thresholding does not shut down the engine or generate troubleshoot data. For more information, see [Understanding Rule Latency Thresholding](#) on page 1021.

IMPORTANT! 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.

If detection is bypassed, the device generates a health monitoring alert. For more information on that health monitoring alert, see [Using the Health Monitor](#) on page 2163.

For more information about enabling Automatic Application Bypass and setting the bypass threshold, see [Editing Advanced Device Settings](#) on page 251.

Editing Advanced Device Settings

LICENSE: Any

SUPPORTED DEVICES: feature dependent

You can use the Advanced section of the Devices tab to modify the Automatic Application Bypass and Inspect Local Router Traffic settings. You can also configure fast-path rules, as explained in [Configuring Fast-Path Rules](#) on page 253.

Note the following:

- you can configure fast-path rules only on 8000 Series and 3D9900 devices.
- you can configure **Inspect Local Router Traffic** only on Series 3 devices

To modify advanced device settings:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

2. Next to the device where you want to edit advanced device settings, click the edit icon (✎).

The Interfaces tab for that device appears.

3. Click **Device**.

The Devices tab appears.

TIP! For stacked devices, you edit the advanced device settings for the stack on the Stack page of the appliance editor.

4. Next to the **Advanced** section, click the edit icon (✎).

The Advanced pop-up window appears.

The screenshot shows a window titled "Advanced" with a search icon and a close icon in the top right corner. The window contains the following elements:

- Automatic Application Bypass:
- Bypass Threshold (ms):
- Inspect Local Router Traffic:
- Fast-Path Rules: A table with columns: #, Domain, Initiator, Initiator P..., Responder, Responder P..., Protocol, VLAN, Bidirectio... There are two buttons: "New IPv4 Rule" and "New IPv6 Rule".
- Save and Cancel buttons at the bottom right.

5. 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](#) on page 251.
6. 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.

7. Optionally, select the **Inspect Local Router Traffic** check box to inspect exception traffic when deployed as a router.
8. Optionally, configure fast-path rules. For more information, see [Configuring Fast-Path Rules](#) on page 253.
9. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring Fast-Path Rules

LICENSE: Any

SUPPORTED DEVICES: 8000 Series, 3D9900

You can create fast-path rules to send traffic directly through a device with no further inspection. Fast-path rules divert traffic that does not need to be analyzed to bypass the device. Fast-path rules either send traffic to the fast-path (out of the interface) or allow it to continue into the device for further analysis. Their advantage is the speed at which they determine the correct path for the traffic. Because the fast-path rules function at the hardware level, they only determine limited information about the packet.

See the following sections for more information:

- [Adding IPv4 Fast-Path Rules](#) on page 253
- [Adding IPv6 Fast-Path Rules](#) on page 255
- [Deleting Fast-Path Rules](#) on page 257

Adding IPv4 Fast-Path Rules

LICENSE: Any

SUPPORTED DEVICES: 8000 Series, 3D9900

Fast-path rules send traffic to the fast-path (out of the interface) or into the device for further analysis. You can use the following criteria to select the IPv4 traffic you want to divert to the fast-path and not inspect:

- initiator or responder IP address or CIDR block
- protocol
- initiator or responder port, for TCP or UDP protocols
- VLAN ID
- bidirectional option

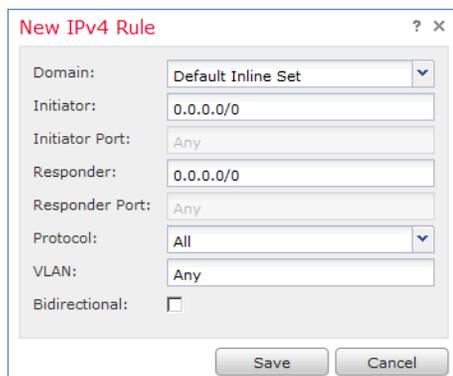
Note that the outermost ID is used for fast-path rules.

TIP! To edit an existing fast-path rule, click the edit icon () next to the rule.

To build or edit IPv4 fast-path rules:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add a fast-path rule, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Device**.
The Devices tab appears.
4. Next to the **Advanced** section, click the edit icon (✎).
The Advanced pop-up window appears.
5. Click **New IPv4 Rule** to add a fast-path rule.
The New IPv4 Rule pop-up window appears.



6. From the **Domain** drop-down list, select an inline set or passive security zone. See [Setting Up an IPS Device](#) on page 266 for more information.
7. Use CIDR notation in the **Initiator** and the **Responder** fields to designate the IP address ranges of initiators or responders whose packets should bypass further analysis.
Your rule matches packets from the designated initiators or packets to the designated responders. For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
8. Optionally, from the **Protocol** drop-down list, select the protocol on which you want the rule to act or select **All** to match traffic from any protocol on the list.

9. Optionally, if you chose the TCP or UDP protocol in step 8, enter initiator and responder ports in the **Initiator Port** and the **Responder Port** fields to designate ports.

TIP! You can enter a comma-separated list of port numbers in each rule. You cannot use port ranges in IPv4 fast-path rules.

If you also select the **Bidirectional** option, your filter criteria are narrowed to packets from those initiator ports or packets to those responder ports.

10. Optionally, enter a VLAN ID in the **VLAN** field.
Your rule matches only traffic for that VLAN.
11. Optionally, select the **Bidirectional** option to filter all traffic traveling between the specified initiator and responder IP addresses. Clear the option to filter only traffic from the specified initiator IP address to the specified responder IP address.
12. Click **Save**.

The rule is added under Fast-Path Rules in the Advanced pop-up window. Although the rule is added, you must click **Save** again to save the rule. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Adding IPv6 Fast-Path Rules

LICENSE: Any

SUPPORTED DEVICES: 8000 Series, 3D9900

Fast-path rules send traffic to the fast-path (out of the interface) or into the device for further analysis. You can use the following criteria to select the IPv6 traffic you want to divert to the fast-path and not inspect:

- protocol
- initiator or responder port, for TCP or UDP protocols
- VLAN ID
- bidirectional option

Note that the outermost VLAN ID is used for fast-path rules.

TIP! To edit an existing fast-path rule, click the edit icon () next to the rule.

To add an IPv6 fast-path rule:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.

2. Next to the device where you want to add a fast-path rule, click the edit icon (✎).

The Interfaces tab for that device appears.

3. Click **Device**.

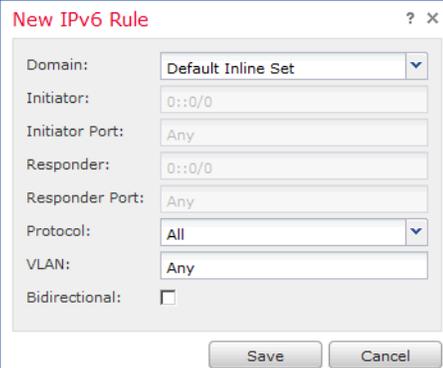
The Devices tab appears.

4. Next to the **Advanced** section, click the edit icon.

The Advanced pop-up window appears.

5. Click **New IPv6 Rule** to add a fast-path rule.

The New IPv6 Rule pop-up window appears. Note that the initiator and responder fields are fixed and indicate that the filter applies to IPv6 packets from any initiator or responder.



6. From the **Domain** drop-down list, select an inline set or passive security zone. See [Setting Up an IPS Device](#) on page 266 for more information.
7. Optionally, from the **Protocol** drop-down list, select the protocol on which you want the rule to act or select **All** to match traffic from any protocol on the list. Your fast-path rule matches only the selected protocol's packets.
8. Optionally, if you chose the TCP or UDP protocol in step 7, enter initiator and responder ports in the **Initiator Port** and the **Responder Port** fields to designate ports.

TIP! You can enter a comma-separated list of port numbers in each rule. You cannot use port ranges in IPv6 fast-path rules.

9. Optionally, enter a VLAN ID in the **VLAN** field. Your rule matches only traffic for that VLAN.

10. Optionally, select **Bidirectional** to filter all traffic traveling between the specified initiator and responder ports. Clear the option to specify that your rule matches only packets from those initiator ports or packets to those responder ports.

IMPORTANT! Because IPv6 fast-path rules do not specify initiator and responder IP addresses, the bidirectional option only functions with initiator and responder ports.

11. Click **Save**.
The rule is added under Fast-Path Rules in the Advanced pop-up window.
12. In the Advanced pop-up window, click **Save**.
The rule is saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Deleting Fast-Path Rules

LICENSE: Any
SUPPORTED DEVICES: 8000 Series, 3D9900

The following procedure explains how to delete any IPv4 or IPv6 fast-path rule.

To delete any fast-path rule:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to delete a fast-path rule, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Device**.
The Devices tab appears.
4. Next to the **Advanced** section, click the edit icon (✎).
The Advanced pop-up window appears.
5. Next to the fast-path rule you want to delete, click the delete icon (🗑️).
6. When prompted, confirm that you want to delete the rule.
The rule is removed from the Advanced pop-up window.

7. Click **Save**.

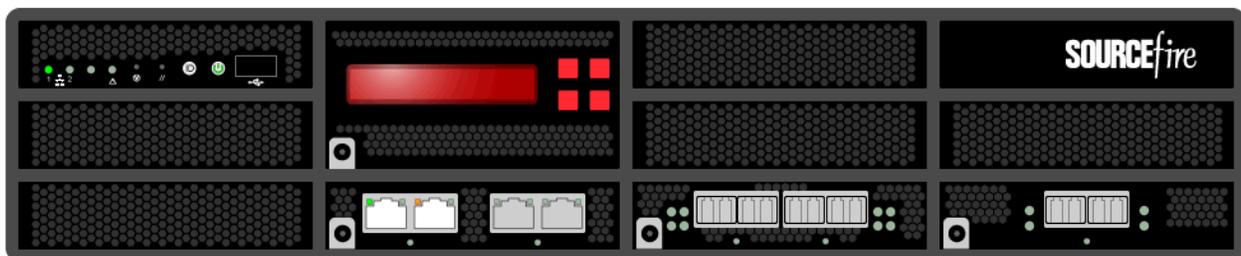
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring Interfaces

LICENSE: Any

You can configure the interfaces of a managed device, according to your Sourcefire 3D System deployment, from the Interfaces page of the appliance editor.

The top of the Interfaces page displays a physical hardware view of a managed Series 3 device. Series 2 and virtual devices do not have physical hardware views. The following graphic shows the hardware view for a 3D8250.



The Using the Hardware View table explains how to use the physical hardware view.

Using the Hardware View

To...	YOU CAN...
view a network module's type, part number, and serial number	hover your cursor over the dark circle in the lower left corner of the network module.
select an interface in the interfaces table view	click the interface.
open an interface editor	double-click the interface.

Using the Hardware View (Continued)

To...	YOU CAN...
view the name of the interface, the type of interface, whether the interface has link, the interface's speed setting, and whether the interface is currently in bypass mode	hover your cursor over the interface.
view the details about an error or warning	hover your cursor over the affected port on the network module.

The interfaces table view, which is below the Series 3 hardware view, lists all the interfaces you have configured on a device. The table includes summarized information about each interface, as described in the following table. Note that only 8000 Series devices display the MAC Address and IP Address columns. See the Interfaces Table View Fields table for more information.

Interfaces Table View Fields

FIELD	DESCRIPTION
Link	The current link state of the interface. Logical interfaces have the same link state as their parent physical interface.
Name	The name of the interface. Interface names are auto-generated with the exception of hybrid interfaces, which are user-defined. Physical interfaces display the name of the physical interface. Logical interfaces display the name of the physical interface and the assigned VLAN tag.
Type	The configuration of the interface: <ul style="list-style-type: none"> • None • Management • Passive • Inline • Switched • Routed • Hybrid • HA link
Security Zone	The security zone where the interface is assigned.
Used by	The inline set, virtual switch, or virtual router where the interface is assigned.

Interfaces Table View Fields (Continued)

FIELD	DESCRIPTION
MAC Address	The MAC address displayed for the interface when it is enabled for switched and routed features. For virtual devices, the MAC address is displayed so that you can match the network adapters configured on your device to the interfaces that appear on the Interfaces page.
IP Addresses	IP addresses assigned to the interface. Hover your pointer over an IP address to view whether it is active or inactive. Inactive IP addresses are grayed out.

Note that you can only configure a total of 1024 interfaces on a managed device. See the following sections for details on the different ways you can configure interfaces on a device:

- [Configuring the Management Interface](#) on page 260
- [Configuring HA Link Interfaces](#) on page 262
- [Configuring the Interface MTU](#) on page 263
- [Disabling Interfaces](#) on page 264
- [Preventing Duplicate Connection Logging](#) on page 265
- [Setting Up an IPS Device](#) on page 266
- [Setting Up Virtual Switches](#) on page 281
- [Setting Up Virtual Routers](#) on page 295
- [Setting Up Hybrid Interfaces](#) on page 352

Configuring the Management Interface

LICENSE: Any

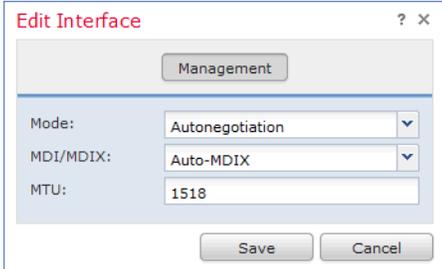
You can configure the link mode and MDI/MDIX settings for the management interface from the appliance editor. You must configure all other management interface settings, such as IPv4 address, IPv6 address, and DNS settings locally on the device.

WARNING! Changing the maximum transmission unit (MTU) interrupts traffic on the device. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To configure the management interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to configure the management interface, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Next to the management interface, click the edit icon (✎).
The Edit Interface pop-up window appears.



The screenshot shows a dialog box titled "Edit Interface" with a "Management" tab. It contains three configuration fields: "Mode" set to "Autonegotiation", "MDI/MDIX" set to "Auto-MDIX", and "MTU" set to "1518". At the bottom, there are "Save" and "Cancel" buttons.

4. From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to autonegotiate speed and duplex settings.
5. From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.
Normally, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.
6. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.
The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.
7. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring HA Link Interfaces

LICENSE: Any

SUPPORTED DEVICES: Series 3

After you establish a device cluster, you can configure a physical interface as a high availability (HA) link interface. This link acts as a redundant communications channel for sharing health information between the clustered devices. When you configure an HA link interface on one device, you automatically configure an interface on the second device. You must configure both HA links on the same broadcast domain. See [Clustering Devices](#) on page 220 for more information.

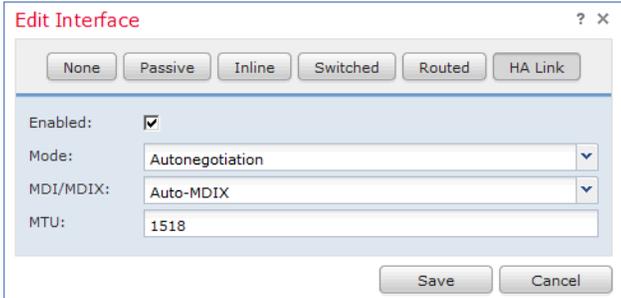
Dynamic NAT relies on dynamically allocating IP addresses and ports to map to other IP addresses and ports. Without an HA link, these mappings are lost in a failover, causing all translated connections to fail as they are routed through the now active device in the cluster.

WARNING! Changing the maximum transmission unit (MTU) interrupts traffic on the device. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, the device model, and the interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To configure an HA link interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the clustered device where you want to configure the HA link interface, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Next to the interface you want to configure as a HA link interface, click the edit icon (✎).
The Edit Interface pop-up window appears.
4. Click **HA Link** to display the HA link options.



The screenshot shows the 'Edit Interface' dialog box with the following configuration:

- Mode: HA Link (selected)
- Enabled:
- Mode: Autonegotiation
- MDI/MDIX: Auto-MDIX
- MTU: 1518

Buttons: Save, Cancel

5. Select the **Enabled** check box to allow the HA link interface to provide link.
If you clear the check box, the interface becomes disabled and administratively taken down.
6. From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to autonegotiate speed and duplex settings.
7. From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX.
Normally, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.
8. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.
The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, the device model, and the interface type. See [Configuring the Interface MTU](#) on page 263 for more information.
9. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring the Interface MTU

LICENSE: Any

When you change the maximum transmission unit (MTU) on an interface, the following capabilities may be affected:

- traffic inspection, including application awareness and control, URL filtering, intrusion detection and prevention, and connection logging
- traffic flow, including switching, routing, and related functionality
- link state

The manner and duration of network traffic interruption depends on the interface type and how your devices are configured and deployed.

Consider the following when setting the interface MTU:

- Sourcefire 3D System software versions prior to 5.2 do not support Series 2 devices.
- The minimum IPv6 MTU setting is 1280.

The following table lists MTU configuration ranges for managed devices running different software versions.

MTU Range by Software Version and Device

FOR THIS SOFTWARE VERSION...	ON THIS MODEL DEVICE...	THE MTU RANGE (INTERFACE) IS...
5.1 to 5.1.1	Series 3, virtual	1518-10228 (hybrid) 1518-9018 (all others)
5.1.1 to pre-5.2	Series 3, virtual	1518-9234 (management) 1518-10228 (hybrid) 1518-10172 (all others)
5.2 or greater	Series 2 except 3D6500, 3D9900	594-1518 (all interfaces)
5.2 or greater	3D6500, 3D9900, virtual	594-9018 (all interfaces)
5.2 or greater	Series 3	594-9234 (management) 594-10172 (inline, passive) 594-9922 (all others)

Disabling Interfaces

You can disable an interface by setting the interface type to **None**. Disabled interfaces appear grayed out in the interface list.

LICENSE: Any

To disable an interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to disable the interface, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Next to the interface you want to disable, click the edit icon (✎).
The Edit Interface pop-up window appears.
4. Click **None**.

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 Devices](#) on page 211 for more information.

Preventing Duplicate Connection Logging

When you update a security zone object, the system saves a new revision of the object. As a result, if you have managed devices in the same security zone that have different revisions of the security zone object configured in the interfaces, you may log what appear to be duplicate connections.

If you notice duplicate connection reporting, you can update all managed devices to use the same revision of the object.

To synchronize security zone object revisions across devices:

1. Select **Devices > Device Management**.

The Device Management page appears.

WARNING! You must not reapply managed device changes to any device until you have edited the zone setting for interfaces on all devices you want to sync.

2. Next to the device where you want to update the security zone selection, click the edit icon ().
- The Interfaces tab for that device appears.
3. For each interface logging duplicate connection events, change the **Security Zone** to another zone, click **Save**, then change it back to the desired zone, and click **Save** again.
4. Repeat steps 2 through 3 for each device logging duplicate events.
5. After all interfaces on all devices have been edited, apply device changes to all managed devices at once.

CHAPTER 6

SETTING UP AN IPS DEVICE

You can configure your device in either a passive or inline IPS deployment. In a passive deployment, you deploy the system out of band from the flow of network traffic. In an inline deployment, you configure the system transparently on a network segment by binding two ports together.

The following sections describe configuring your device for passive and inline deployments of the Sourcefire 3D System:

- [Understanding Passive IPS Deployments](#) on page 266
- [Configuring Passive Interfaces](#) on page 267
- [Understanding Inline IPS Deployments](#) on page 269
- [Configuring Inline Interfaces](#) on page 269
- [Configuring Inline Sets](#) on page 271

Understanding Passive IPS Deployments

LICENSE: Protection

In a passive IPS deployment, the Sourcefire 3D System monitors traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This provides the system visibility within the network without being in the flow of network traffic. When configured in a passive deployment, the system cannot take certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally and no traffic received on these interfaces is retransmitted.

Configuring Passive Interfaces

LICENSE: Protection

You can configure one or more physical ports on a managed device as passive interfaces.

WARNING! Changing the maximum transmission unit (MTU) interrupts traffic on the device. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To configure a passive interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

Name	License Type	Health Policy	System Policy	Access Control Pol...
Ungrouped (1)				
linden 10.10.10.10 - 3D8250	Protect & Control w/URL Filtering	Blacklisted Power Supply	katsura system policy	My Access Control Policy

2. Next to the device where you want to configure the passive interface, click the edit icon (✎).

The Interfaces tab appears.

linden
3D8250

Apply Changes

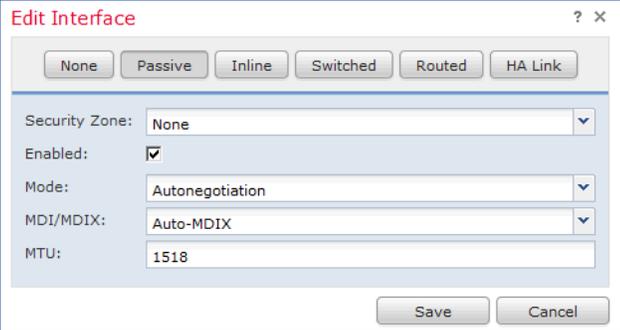
Device | **Interfaces** | Inline Sets | Virtual Switches | Virtual Routers | NAT

Add Interface

Li...	Name...	Type	Security Zone	Used By	IP Addresses
eth0	eth0	Managem			
s1p1	s1p1	Inline	Internal	Default Inline Set	

3. Next to the interface you want to configure as a passive interface, click the edit icon (✎).

The Edit Interface pop-up window appears.



4. Click **Passive** to display the passive interface options.
5. Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.
6. Select the **Enabled** check box to allow the passive interface to monitor traffic. If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.
7. From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to automatically negotiate speed and duplex settings. Note that mode settings are available only for copper interfaces.
8. From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX. Note that MDI/MDIX settings are available only for copper interfaces.
By default, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.
9. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.
The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.
10. Click **Save**.

The passive interface is configured. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Understanding Inline IPS Deployments

LICENSE: Protection

In an inline IPS deployment, you configure the Sourcefire 3D System transparently on a network segment by binding two ports together. This allows the system to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

Configuring Inline Interfaces

LICENSE: Protection

You can configure one or more physical ports on a managed device as inline interfaces. You must assign a pair of inline interfaces to an inline set before they can handle traffic in an inline deployment.

IMPORTANT! If you configure an interface as an inline interface, the adjacent port on its NetMod automatically becomes an inline interface as well to complete the pair.

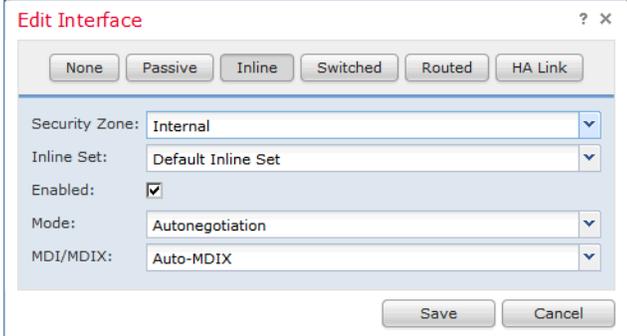
To configure inline interfaces on a virtual device, you must create the inline pair using adjacent interfaces.

To configure an inline interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to configure the inline interface, click the edit icon (✎).
The Interfaces tab appears.
3. Next to the interface you want to configure as an inline interface, click the edit icon (✎).
The Edit Interface pop-up window appears.

4. Click **Inline** to display the inline interface options.



The screenshot shows a dialog box titled "Edit Interface" with a close button (X) and a help button (?). At the top, there are six tabs: "None", "Passive", "Inline", "Switched", "Routed", and "HA Link". The "Inline" tab is selected. Below the tabs, there are several configuration fields: "Security Zone" with a dropdown menu showing "Internal"; "Inline Set" with a dropdown menu showing "Default Inline Set"; "Enabled" with a checked checkbox; "Mode" with a dropdown menu showing "Autonegotiation"; and "MDI/MDIX" with a dropdown menu showing "Auto-MDIX". At the bottom of the dialog, there are "Save" and "Cancel" buttons.

5. Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.
6. From the **Inline Set** drop-down list, select an existing inline set or select **New** to add a new inline set.

Note that if you add a new inline set, you must configure it on the Device Management page (**Devices > Device Management > Inline Sets**) after you set up the inline interface. For more information, see [Adding Inline Sets](#) on page 271.
7. Select the **Enabled** check box to allow the inline interface to handle traffic.

If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.
8. From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to automatically negotiate speed and duplex settings. Note that Mode settings are available only for copper interfaces.
9. From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX. Note that MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to **Auto-MDIX**, which automatically handles switching between MDI and MDIX to attain link.
10. Click **Save**.

The inline interface is configured. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring Inline Sets

LICENSE: Protection

Before you can use inline interfaces in an inline deployment, you must configure inline sets and assign inline interface pairs to them. An inline set is a grouping of one or more inline interface pairs on a device; an inline interface pair can belong to only one inline set at a time.

See the following sections for more information:

- [Viewing Inline Sets](#) on page 271
- [Adding Inline Sets](#) on page 271
- [Configuring Advanced Inline Set Options](#) on page 275
- [Deleting Inline Sets](#) on page 279

Viewing Inline Sets

LICENSE: Protection

The Inline Sets tab of the Device Management page displays a list of all inline sets you have configured on a device. Note that you cannot configure inline sets to go into bypass mode on a virtual device. The Inline Sets Table View Fields table includes summary information about each set.

Inline Sets Table View Fields

FIELD	DESCRIPTION
Name	The name of the inline set.
Interface Pairs	A list of all pairs of inline interfaces assigned to the inline set. A pair is not available when you disable either interface in the pair from the Interfaces tab.
Bypass	The configured bypass mode of the inline set.

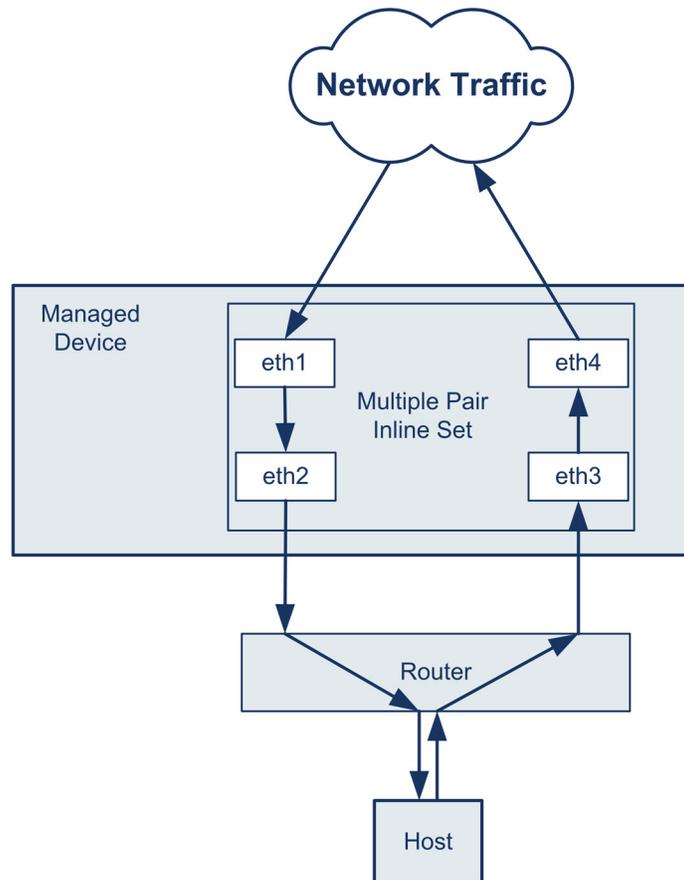
Adding Inline Sets

LICENSE: Protection

You can add inline sets from the Inline Sets tab of the Device Management page. You can also add inline sets as you configure inline interfaces.

You can assign only inline interface pairs to an inline set. If you want to create an inline set before you configure the inline interfaces on your managed devices, you can create an empty inline set and add interfaces to it later.

You can also add multiple interface pairs when your network uses asynchronous routing, as shown in the following diagram.



Your network may be set up to route traffic between a host on your network and external hosts through different inline interface pairs, depending on whether the traffic is inbound or outbound. If you include only one interface pair in an inline set, the device may not correctly analyze your network traffic because it might see only half of the traffic.

For devices with inline sets, a software bridge is automatically set up to transport packets after the device restarts. If the device is restarting, there is no software bridge running anywhere. If you enable bypass mode on the inline set, it goes into hardware bypass while the device is restarting. In that case, you may lose a few seconds of packets as the system goes down and comes back up, due to

renegotiation of link with the device. However, the system will pass traffic while Snort is restarting.

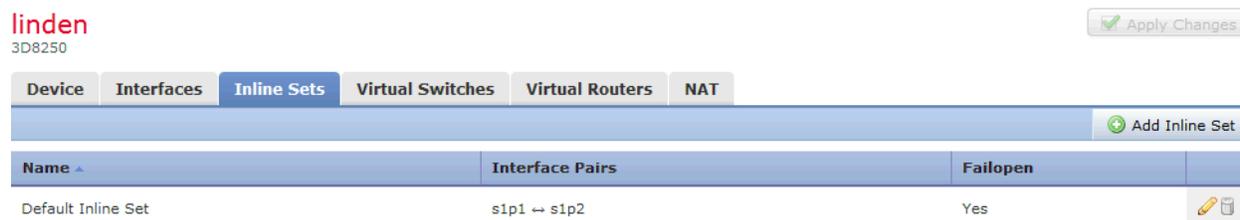
WARNING! Changes you make to an existing inline set may interrupt traffic on the device. Changing the maximum transmission unit (MTU) interrupts traffic on the device; some packets are transmitted without inspection and dropped. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To edit an existing inline set, click the edit icon (✎) next to the set.

To add an inline set:

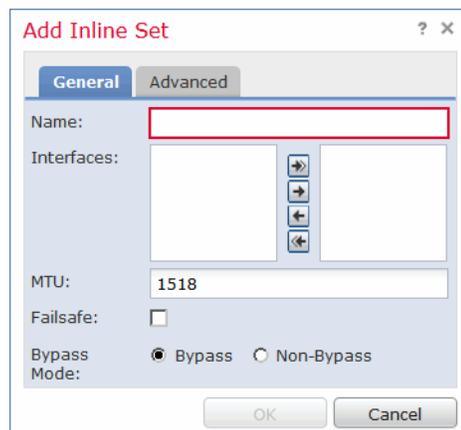
ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the inline set, click the edit icon (✎).
The Interfaces tab appears.
3. Click **Inline Sets**.
The Inline Sets tab appears.



4. Click **Add Inline Set**.

The Add Inline Set pop-up window appears.



5. In the **Name** field, type a name for the inline set. You can use alphanumeric characters and spaces.

6. You have two options for selecting inline interface pairs to add to the inline set:

- Next to **Interfaces**, select one or more inline interface pairs, then click the add selected icon (➔). Use Ctrl or Shift to select multiple inline interface pairs.
- To add all interface pairs to the inline set, click the add all icon (➔➔).

TIP! To remove inline interfaces from the inline set, select one or more inline interface pairs and click the remove selected icon (⬅). To remove all interface pairs from the inline set, click the remove all icon (⬅➔). Disabling either interface in a pair from the Interfaces tab also removes the pair.

7. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.

The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

8. Optionally, select **Failsafe** to specify that traffic is allowed to bypass detection and continue through the device. Managed devices monitor internal traffic buffers and bypass detection if those buffers are full.

Note that only Series 3 and 3D9900 devices support this option.

9. Select the bypass mode to configure how the relays in the inline interfaces respond when an interface fails:
 - Select **Bypass** to allow traffic to continue to pass through the interfaces.
 - Select **Non-Bypass** to block traffic.

IMPORTANT! In bypass mode, you may lose a few packets when you reboot the appliance. Also note that you cannot configure bypass mode for inline sets on a virtual device, for non-bypass NetMods on 8000 Series devices, or for SFP modules on 3D7115 or 3D7125 devices.

10. Click **OK**.

The inline set is added. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

TIP! To configure advanced settings for the inline set, such as tap mode, link state propagation, and transparent inline mode, see [Configuring Advanced Inline Set Options](#) on page 275.

Configuring Advanced Inline Set Options

LICENSE: Protection

SUPPORTED DEVICES: feature dependent

There are a number of options you may consider as you configure inline sets. See the sections below for more information about each option.

Tap Mode

SUPPORTED DEVICES: Series 3, 3D9900

Tap mode is available on 3D9900 and Series 3 devices when you create an inline or inline with fail-open interface set.

With tap mode, the device is deployed inline, but instead of the packet flow passing through the device, a copy of each packet is sent to the device and the network traffic flow is undisturbed. Because you are working with copies of packets rather than the packets themselves, rules that you set to drop and rules that use the replace keyword do not affect the packet stream. However, rules of these types do generate intrusion events when they are triggered, and the table view of intrusion events indicates that the triggering packets would have dropped in an inline deployment.

There are benefits to using tap mode with devices that are deployed inline. For example, you can set up the cabling between the device and the network as if the device were inline and analyze the kinds of intrusion events the device generates. Based on the results, you can modify your intrusion policy and add the drop rules

that best protect your network without impacting its efficiency. When you are ready to deploy the device inline, you can disable tap mode and begin dropping suspicious traffic without having to reconfigure the cabling between the device and the network.

Note that you cannot enable this option and strict TCP enforcement on the same inline set.

Propagate Link State

SUPPORTED DEVICES: Series 2, Series 3

Link state propagation is a feature for inline sets configured in bypass mode so both pairs of an inline set track state.

Link state propagation automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the link state of the other interface changes automatically to match it. Link state propagation is available for both copper and fiber configurable bypass interfaces.

IMPORTANT! When link state propagation triggers, fiber inline sets configured as fail-open on Series 2 devices (except those on 3D9900s) activate hardware bypass mode. In this case, the interface cards involved do not come out of bypass automatically; you must bring them out of bypass mode manually. For more information about fiber interfaces in inline sets and hardware bypass, see [Removing Bypass Mode on Fiber Inline Sets Configured to Fail Open](#) on page 279.

Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

You cannot disable link state propagation for inline sets configured on clustered devices.

Note that virtual devices do not support this option.

Transparent Inline Mode

Transparent Inline Mode option allows the device to act as a “bump in the wire” and means that the device forwards all the network traffic it sees, regardless of its source and destination. Note that you cannot disable this option on Series 3 or 3D9900 devices.

Strict TCP Enforcement

SUPPORTED DEVICES: Series 3

To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed. Strict enforcement also blocks:

- non-SYN TCP packets for connections where the three-way handshake was not completed
- non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK
- non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

Note that Series 2 and virtual devices do not support this option. In addition, you cannot enable this option and tap mode on the same inline set.

WARNING! Managed devices must be running at least Version 5.1 of the Sourcefire 3D System before you can enable strict TCP enforcement. Although you can apply a device configuration with strict TCP enforcement to pre-5.1 devices, those devices do not perform strict enforcement.

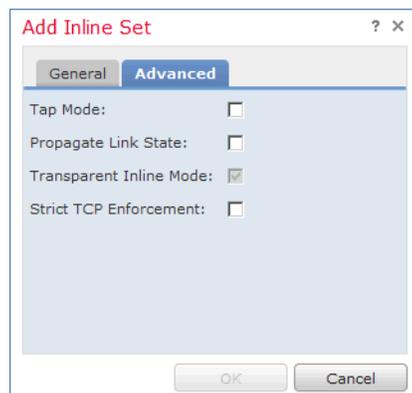
To configure advanced inline set options:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to edit the inline set, click the edit icon (✎).
The Interfaces tab appears.
3. Click **Inline Sets**.
The Inline Sets tab appears.
4. Next to the inline set you want to edit, click the edit icon (✎).
The Edit Inline Set pop-up window appears.

5. Click **Advanced**.

The Advanced tab appears.



6. Optionally, select **Tap Mode** to enable tap mode on the inline interfaces of Series 3 and 3D9900 devices.
Note that virtual devices and Series 2 devices other than 3D9900 do not support this option. In addition, you cannot enable Tap Mode and Strict TCP Enforcement on the same inline set.
7. Optionally, select **Propagate Link State** on Series 2 or Series 3 devices. This option is especially useful if the routers on your network are able to reroute traffic around a network device that is down.
You cannot disable link state propagation for inline sets configured on clustered devices.
Note that virtual devices do not support this option.
8. Optionally, select **Strict TCP Enforcement** to enable strict TCP enforcement on Series 3 devices.
Note that Series 2 and virtual devices do not support this option. In addition, you cannot enable Strict TCP Enforcement and Tap Mode on the same inline set.
9. Optionally, select **Transparent Inline Mode**.
Note that you cannot disable this option on Series 3 or 3D9900 devices.
10. Click **OK**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Removing Bypass Mode on Fiber Inline Sets Configured to Fail Open

LICENSE: Protection

SUPPORTED DEVICES: Series 2 except 3D9900

When link state propagation is enabled on a Series 2 device with a fiber inline set configured to fail open and the device goes into bypass mode, all network traffic passes through the inline set without being analyzed. When the links restore, most fiber inline sets configured to fail open do not return from bypass automatically. You can use a command line tool to force the inline set out of bypass mode.

This tool works on inline sets with fiber inline interfaces configured to fail open. It is not necessary to use this tool on inline sets with copper inline interfaces set to fail open.

IMPORTANT! Contact Technical Support if you are having issues with inline sets configured to fail open on your device.

To force a fiber inline set configured to fail open out of bypass mode on a device:

ACCESS: Admin/Network Admin

1. Open a terminal window on your device and sign in as the admin user.
2. Enter the following at the command line:

```
sudo /var/sf/bin/unbypass_cards.sh
```

You are prompted for your password.
3. When the interfaces switch out of bypass mode, a message in syslog indicates the device is analyzing traffic. For example:

```
Fiber pair has been reset by un_bypass
```

Deleting Inline Sets

LICENSE: Protection

When you delete an inline set, any inline interfaces assigned to the set become available for inclusion in another set. The interfaces are not deleted.

To delete an inline set:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to delete the inline set, click the edit icon ().
The Interfaces tab appears.

3. Click **Inline Sets**.
The Inline Sets tab appears.
4. Next to the inline set you want to delete, click the delete icon ().
5. When prompted, confirm that you want to delete the inline set.
The inline set is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

CHAPTER 7

SETTING UP VIRTUAL SWITCHES

You can configure a managed device in a Layer 2 deployment so that it provides packet switching between two or more networks. In a Layer 2 deployment, you can configure virtual switches on managed devices to operate as standalone broadcast domains, dividing your network into logical segments. A virtual switch uses the media access control (MAC) address from a host to determine where to send packets.

When you configure a virtual switch, the switch initially broadcasts packets through every available port on the switch. Over time, the switch uses tagged return traffic to learn which hosts reside on the networks connected to each port.

IMPORTANT! In a Layer 2 deployment, you cannot block egress traffic based on destination network or destination security zone. You must instead write access control rules that block ingress traffic based on blocking source network or source security zone. For more information on adding zones and networks to access control rules, see [Adding Zone Conditions](#) on page 492 and [Adding Network Conditions](#) on page 493.

A virtual switch must contain two or more switched interfaces to handle traffic. For each virtual switch, traffic becomes limited to the set of ports configured as switched interfaces. For example, if you configure a virtual switch with four switched interfaces, packets sent in through one port for broadcast can only be sent out of the remaining three ports on the switch.

When you configure a physical switched interface, you must assign it to a virtual switch. You can also define additional logical switched interfaces on a physical port as needed.

Note that you cannot configure virtual switches, physical switched interfaces, or logical switched interfaces on a virtual device.

WARNING! If a Layer 2 deployment fails for any reason, the device no longer passes traffic.

See the following sections for more information about configuring a Layer 2 deployment:

- [Configuring Switched Interfaces](#) on page 282
- [Configuring Virtual Switches](#) on page 288

Configuring Switched Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can set up switched interfaces to have either physical or logical configurations. You can configure physical switched interfaces for handling untagged VLAN traffic. You can also create logical switched interfaces for handling traffic with designated VLAN tags.

In a Layer 2 deployment, the system drops any traffic received on an external physical interface that does not have a switched interface waiting for it. If the system receives a packet with no VLAN tag and you have not configured a physical switched interface for that port, it drops the packet. If the system receives a VLAN-tagged packet and you have not configured a logical switched interface, it also drops the packet.

The system handles traffic that has been received with VLAN tags on switched interfaces by stripping the outermost VLAN tag on ingress before any rules evaluation or forwarding decisions. Packets leaving the device through a VLAN-tagged logical switched interface are encapsulated with the associated VLAN tag on egress.

Note that if you change the parent physical interface to inline or passive, the system deletes all the associated logical interfaces.

See the following sections for more information:

- [Configuring Physical Switched Interfaces](#) on page 283
- [Adding Logical Switched Interfaces](#) on page 285
- [Deleting Logical Switched Interfaces](#) on page 287

Configuring Physical Switched Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can configure one or more physical ports on a managed device as switched interfaces. You must assign a physical switched interface to a virtual switch before it can handle traffic.

WARNING! Changing the maximum transmission unit (MTU) interrupts traffic on the device and packets are dropped. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To configure a physical switched interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

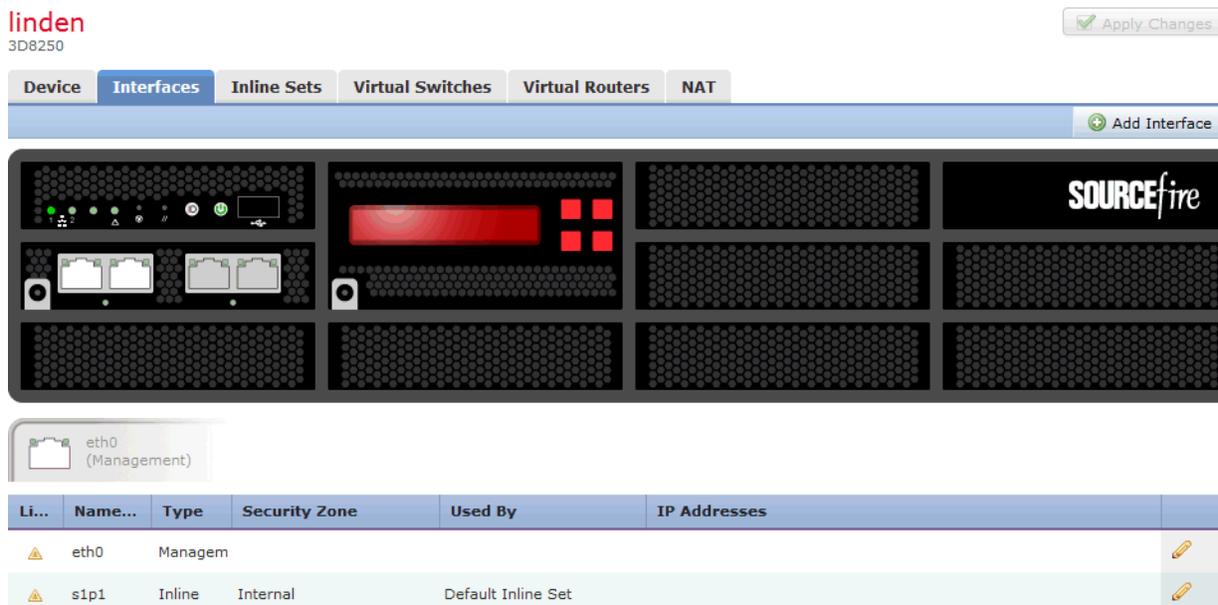
The Device Management page appears.



Name	License Type	Health Policy	System Policy	Access Control Policy
📁 Ungrouped (4)				

- Next to the device where you want to configure the switched interface, click the edit icon (✎).

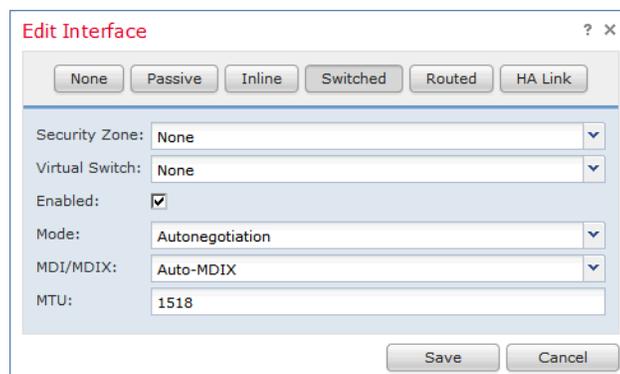
The Interfaces tab appears.



- Next to the interface you want to configure as a switched interface, click the edit icon (✎).

The Edit Interface pop-up window appears.

- Click **Switched** to display the switched interface options.



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

6. Optionally, from the **Virtual Switch** drop-down list, select an existing virtual switch or select **New** to add a new virtual switch.

Note that if you add a new virtual switch, you must configure it on the Virtual Switches tab of the Device Management page (**Devices > Device Management > Virtual Switches**) after you set up the switched interface. See [Adding Virtual Switches](#) on page 289.

7. Select the **Enabled** check box to allow the switched interface to handle traffic. If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.

8. From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to auto negotiate speed and duplex settings. Note that mode settings are available only for copper interfaces.

9. From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX. Note that MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

10. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.

The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

11. Click **Save**.

The physical switched interface is configured. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Adding Logical Switched Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

For each physical switched interface, you can add multiple logical switched interfaces. You must associate each logical interface with a VLAN tag to handle

traffic received by the physical interface with that specific tag. You must assign a logical switched interface to a virtual switch to handle traffic.

WARNING! Any changes you make to the maximum transmission unit (MTU) interrupts switched traffic on the device and packets are dropped. Managed devices must be running at least Version 5.1.1 of the Sourcefire 3D System before you can set their MTU to an upper limit of 10172. Pre-5.1.1 devices have an MTU limit of 9018; these devices will drop packets that exceed this limit.

To edit an existing logical switched interface, click the edit icon (✎) next to the interface.

To add a logical switched interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the switched interface, click the edit icon (✎).
The Interfaces tab appears.
3. Click **Add Interface**.
The Add Interface pop-up window appears.
4. Click **Switched** to display the switched interface options.

The screenshot shows the 'Add Interface' dialog box with the 'Switched' tab selected. The 'Interface' dropdown is set to 's1p3', 'VLAN Tag' is '1', 'Security Zone' is 'None', 'Virtual Switch' is 'None', 'Enabled' is checked, and 'MTU' is '1518'. The 'Save' and 'Cancel' buttons are at the bottom right.

5. From the **Interface** drop-down list, select the physical interface that will receive the VLAN-tagged traffic.
6. In the **VLAN Tag** field, type a tag value that gets assigned to inbound and outbound traffic on this interface. The value can be any integer from 1 to 4094.
7. Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.

8. Optionally, from the Virtual Switch drop-down list, select an existing virtual switch or select **New** to add a new virtual switch.
Note that if you add a new virtual switch, you must configure it on the Device Management page (**Devices > Device Management > Virtual Switches**) after you set up the switched interface. See [Adding Virtual Switches](#) on page 289.
9. Select the **Enabled** check box to allow the switched interface to handle traffic.
If you clear the check box, the interface becomes disabled and administratively taken down. If you disable a physical interface, you also disable all of the logical interfaces associated with it.
10. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.
In Version 5.1.1 or greater, the value can be any integer from 1518 to 10172. In pre-5.1.1 devices, the value can be 1518 to 9018.
11. Click **Save**.
The logical switched interface is added. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Deleting Logical Switched Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

When you delete a logical switched interface, you remove it from the physical interface where it resides, as well as the virtual switch and security zone it is associated with.

To delete a switched interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Select the managed device that contains the switched interface you want to delete and click the edit icon (✎) for that device.
The Interfaces tab for that device appears.
3. Next to the logical switched interface you want to delete, click the delete icon (🗑).
4. When prompted, confirm that you want to delete the interface.
The interface is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Configuring Virtual Switches

LICENSE: Control

SUPPORTED DEVICES: Series 3

Before you can use switched interfaces in a Layer 2 deployment, you must configure virtual switches and assign switched interfaces to them. A virtual switch is a group of switched interfaces that process inbound and outbound traffic through your network.

See the following sections for more information about configuring virtual switches:

- [Viewing Virtual Switches](#) on page 288
- [Adding Virtual Switches](#) on page 289
- [Configuring Advanced Virtual Switch Settings](#) on page 291
- [Deleting Virtual Switches](#) on page 294

Viewing Virtual Switches

LICENSE: Control

SUPPORTED DEVICES: Series 3

The Virtual Switches tab of the Device Management page displays a list of all the virtual switches you have configured on a device. The page includes summary information about each switch, as described in the following table.

Virtual Switches Table View Fields

FIELD	DESCRIPTION
Name	The name of the virtual switch.
Interfaces	All switched interfaces that are assigned to the virtual switch. Interfaces that you have disabled from the Interfaces tab are not available.
Hybrid Interface	The optionally configured hybrid interface that ties the virtual switch to a virtual router.

Virtual Switches Table View Fields (Continued)

FIELD	DESCRIPTION
Unicast Packets	Unicast packet statistics for the virtual switch, including: <ul style="list-style-type: none">• Unicast packets received• Unicast packets forwarded (excludes drops by host)• Unicast packets unintentionally dropped
Broadcast Packets	Broadcast packet statistics for the virtual switch, including: <ul style="list-style-type: none">• Broadcast packets received• Broadcast packets forwarded• Broadcast packets unintentionally dropped

Adding Virtual Switches

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can add virtual switches from the Virtual Switches tab of the Device Management page. You can also add switches as you configure switched interfaces.

You can assign only switched interfaces to a virtual switch. If you want to create a virtual switch before you configure the switched interfaces on your managed devices, you can create an empty virtual switch and add interfaces to it later.

TIP! To edit an existing virtual switch, click the edit icon () next to the switch.

Note that any changes you make to an existing virtual switch may interrupt traffic on the device.

To add a virtual switch:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the virtual switch, click the edit icon ().
The Interfaces tab appears.

3. Click **Virtual Switches**.

The Virtual Switches tab appears.

Name	Interface	Hybrid Interface	Unicast Packets	Broadcast Packets	
Test Virtual Switch	s1p3		Received: 0 Forwarded: 0 Dropped: 0	Received: 0 Forwarded: 0 Dropped: 0	

4. Click **Add Virtual Switch**.

The Add Virtual Switch pop-up window appears.

Add Virtual Switch ? X

General Advanced

Name:

Available Selected

s1p3

Add

Hybrid Interface: None

Save Cancel

5. In the **Name** field, type a name for the virtual switch. You can use alphanumeric characters and spaces.
6. Under **Available**, select one or more switched interfaces to add to the virtual switch.

TIP! Interfaces that you have disabled from the Interfaces tab are not available; disabling an interface after you add it removes it from the configuration.

7. Click **Add**.
8. Optionally, from the **Hybrid Interface** drop-down list, select a hybrid interface that ties the virtual switch to a virtual router. For more information, see [Setting Up Hybrid Interfaces](#) on page 352.

9. Click **Save**.

The virtual switch is added. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

TIP! To configure advanced settings for the switch, such as static MAC entries and spanning tree protocol, see [Configuring Advanced Virtual Switch Settings](#) on page 291.

Configuring Advanced Virtual Switch Settings

LICENSE: Control

SUPPORTED DEVICES: Series 3

When adding or editing a virtual switch, you can add static MAC entries, enable Spanning Tree Protocol (STP), drop Bridge Protocol Data Units (BPDU), and enable strict TCP enforcement.

Over time, a virtual switch learns MAC addresses by tagging return traffic from the network. Optionally, you can manually add a static MAC entry, which designates that a MAC address resides on a specific port. Regardless of whether you ever receive traffic from that port, the MAC address remains static in the table. You can specify one or more static MAC addresses for each virtual switch.

STP is a network protocol used to prevent network loops. BPDUs are exchanged through the network, carrying information about network bridges. The protocol uses BPDUs to identify and select the fastest network links, if there are redundant links in the network. If a network link fails, Spanning Tree fails over to an existing alternate link.

If your virtual switch routes traffic between VLANs, similar to a router on a stick, BPDUs enter and exit the device through different logical switched interfaces, but the same physical switched interface. As a result, STP identifies the device as a redundant network loop, which can cause issues in certain Layer 2 deployments. To prevent this, you can configure the virtual switch at the domain level to have the device drop BPDUs when monitoring traffic.

IMPORTANT! Sourcefire strongly recommends that you enable STP when configuring a virtual switch that you plan to deploy in a device cluster.

To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed. Strict enforcement also blocks:

- non-SYN TCP packets for connections where the three-way handshake was not completed
- non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK

- non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

Note that if you associate the virtual switch with a logical hybrid interface, the switch uses the same strict TCP enforcement setting as the virtual router associated with the logical hybrid interface. You cannot specify strict TCP enforcement on the switch in this case.

WARNING! Managed devices must be running at least Version 5.1 of the Sourcefire 3D System before you can enable strict TCP enforcement. Although you can apply a device configuration with strict TCP enforcement to pre-5.1 devices, those devices do not perform strict enforcement.

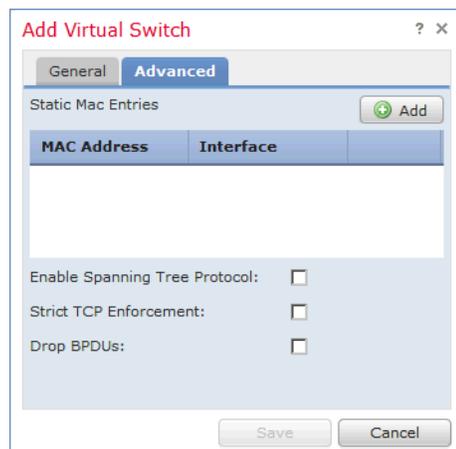
To configure advanced virtual switch settings:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device that contains the virtual switch you want to edit, click the edit icon (✎).
The Interfaces tab appears.
3. Click **Virtual Switches**.
The Virtual Switches tab appears.
4. Next to the virtual switch that you want to edit, click the edit icon (✎).
The Edit Virtual Switch pop-up window appears.

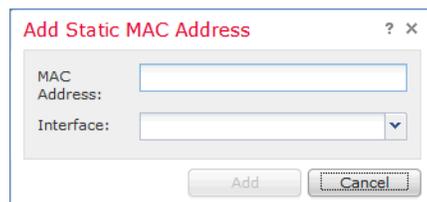
5. Click **Advanced**.

The Advanced tab appears.



6. To add a static MAC entry, click **Add**.

The Add Static MAC Address pop-up window appears.



7. In the **MAC Address** field, type the address using the standard format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).
8. From the **Interface** drop-down list, select the interface where you want to assign the MAC address.
9. Click **Add**.

The MAC address is added to the Static MAC Entries table.

To edit a MAC address, click the edit icon (✎). To delete a MAC address, click the delete icon (🗑).

10. Optionally, to enable the Spanning Tree Protocol, select **Enable Spanning Tree Protocol**. Select **Enable Spanning Tree Protocol** only if your virtual switch switches traffic between multiple network interfaces.
You cannot select **Drop BPDUs** unless you clear **Enable Spanning Tree Protocol**.
11. Optionally, select **Strict TCP Enforcement** to enable strict TCP enforcement.

If you associate the virtual switch with a logical hybrid interface, this option does not appear and the switch uses the same setting as the virtual router associated with the logical hybrid interface.

12. Optionally, select **Drop BPDUs** to drop BPDUs at the domain level. Select **Drop BPDUs** only if your virtual switch routes traffic between VLANs on a single physical interface.
You cannot select **Enable Spanning Tree Protocol** unless you clear **Drop BPDUs**.
13. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

Deleting Virtual Switches

LICENSE: Control

SUPPORTED DEVICES: Series 3

When you delete a virtual switch, any switched interfaces assigned to the switch become available for inclusion in another switch.

To delete a virtual switch:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Select the managed device that contains the virtual switch you want to delete and click the edit icon () for that device.
The Interfaces tab for that device appears.
3. Click **Virtual Switches**.
The Virtual Switches tab appears.
4. Next to the virtual switch that you want to delete, click the delete icon ().
5. When prompted, confirm that you want to delete the virtual switch.
The virtual switch is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211 for more information.

CHAPTER 8

SETTING UP VIRTUAL ROUTERS

You can configure a managed device in a Layer 3 deployment so that it routes traffic between two or more interfaces. You must assign an IP address to each interface and assign the interfaces to a virtual router to route traffic.

You can configure the system to route packets by making packet forwarding decisions according to the destination address. Interfaces configured as routed interfaces receive and forward the Layer 3 traffic. Routers obtain the destination from the outgoing interface based on the forwarding criteria, and access control rules designate the security policies to be applied.

In Layer 3 deployments, you can define static routes. In addition, you can configure Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) dynamic routing protocols. You can also configure a combination of static routes and RIP or static routes and OSPF.

Note that you cannot configure virtual routers, physical routed interfaces, or logical routed interfaces on a virtual device.

WARNING! If a Layer 3 deployment fails for any reason, the device no longer passes traffic.

See the following sections for more information about configuring a Layer 3 deployment:

- [Configuring Routed Interfaces](#) on page 296
- [Configuring Virtual Routers](#) on page 306
- [Configuring NAT on a Virtual Router](#) on page 341

Configuring Routed Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can set up routed interfaces with either physical or logical configurations. You can configure physical routed interfaces for handling untagged VLAN traffic. You can also create logical routed interfaces for handling traffic with designated VLAN tags.

In a Layer 3 deployment, the system drops any traffic received on an external physical interface that does not have a routed interface waiting for it. If the system receives a packet with no VLAN tag and you have not configured a physical routed interface for that port, it drops the packet. If the system receives a VLAN-tagged packet and you have not configured a logical routed interface, it also drops the packet.

The system handles traffic that has been received with VLAN tags on switched interfaces by stripping the outermost VLAN tag on ingress prior to any rules evaluation or forwarding decisions. Packets leaving the device through a VLAN-tagged logical routed interface are encapsulated with the associated VLAN tag on egress. The system drops any traffic received with a VLAN tag after the stripping process completes.

Note that if you change the parent physical interface to inline or passive, the system deletes all the associated logical interfaces.

See the following sections for more information:

- [Configuring Physical Routed Interfaces](#) on page 296
- [Adding Logical Routed Interfaces](#) on page 300
- [Deleting Logical Routed Interfaces](#) on page 304

Configuring Physical Routed Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can configure one or more physical ports on a managed device as routed interfaces. You must assign a physical routed interface to a virtual router before it can route traffic.

You can add static Address Resolution Protocol (ARP) entries to a routed interface. If an external host needs to know the MAC address of the destination IP address it needs to send traffic to on your local network, it sends an ARP request. When you configure static ARP entries, the virtual router responds with an IP address and associated MAC address.

Note that disabling the **ICMP Enable Responses** option for routed interfaces does not prevent ICMP responses in all scenarios. You can add rules to an access control policy to drop packets where the destination IP is the routed interface's IP and the protocol is ICMP. For more information about creating access control

rules, see [Understanding and Writing Access Control Rules](#) on page 472. If you have enabled the **Inspect Local Router Traffic** option on the managed device, it drops the packets before they reach the host, thereby preventing any response. For more information about inspecting local router traffic, see [Understanding Advanced Device Settings](#) on page 250.

WARNING! Changing the maximum transmission unit (MTU) interrupts traffic on the device and packets are dropped. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To configure a physical routed interface:

ACCESS: Admin/Network Admin

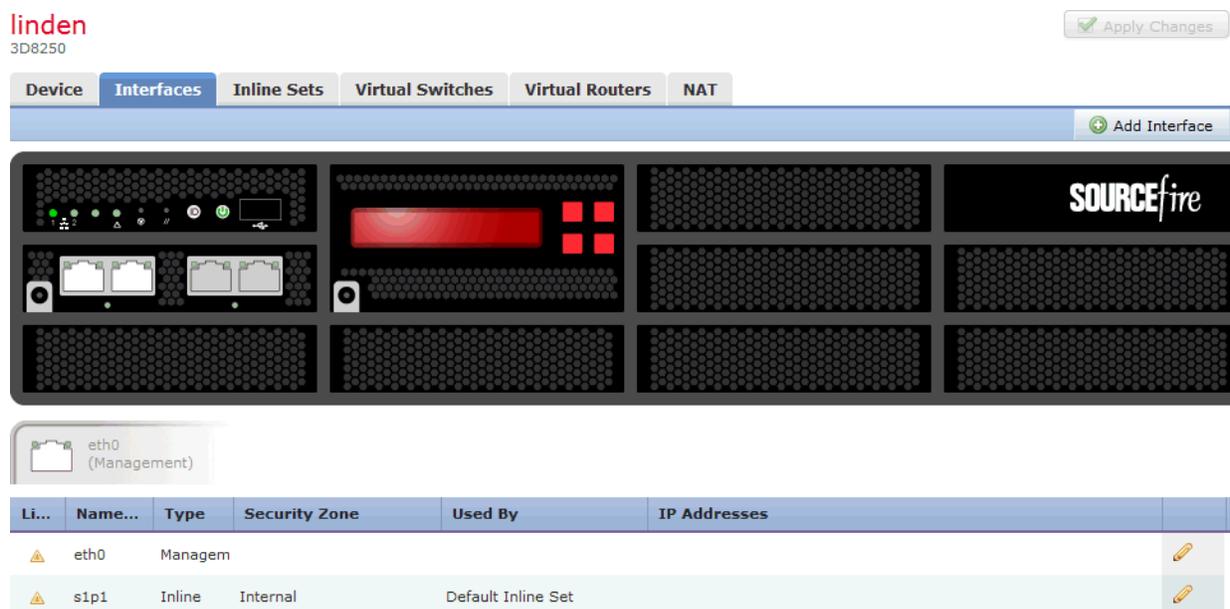
1. Select **Devices > Device Management**.

The Device Management page appears.



2. Next to the device where you want to configure the routed interface, click the edit icon (✎).

The Interfaces tab for that device appears.



3. Next to the interface you want to configure as a routed interface, click the edit icon (✎).

The Edit Interface pop-up window appears.

4. Click **Routed** to display the routed interface options.

Edit Interface ? x

None Passive Inline Switched **Routed** HA Link

Security Zone: None

Virtual Router: None

Enabled:

Mode: Autonegotiation

MDI/MDIX: Auto-MDIX

MTU: 1518

ICMP: Enable Responses

IPv6 NDP: Enable Router Advertisement

IP Addresses:

Address	Type
---------	------

Static ARP Entries:

IP Address	MAC Address
------------	-------------

Save Cancel

5. Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.
6. Optionally, from the **Virtual Router** drop-down list, select an existing virtual router or select **New** to add a new virtual router.

Note that if you add a new virtual router, you must configure it on the Virtual Routers tab of the Device Management page (**Devices > Device Management > Virtual Routers**) after you set up the routed interface. See [Adding Virtual Routers](#) on page 307.

7. Select the **Enabled** check box to allow the routed interface to handle traffic. If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.
8. From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to auto negotiate speed and duplex settings. Note that mode settings are available only for copper interfaces.

9. From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX. Note that MDI/MDIX settings are available only for copper interfaces.

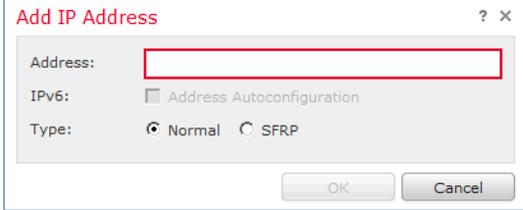
Normally, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

10. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed. Note that the MTU is the Layer 2 MTU/MRU and not the Layer 3 MTU.

The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

11. Next to **ICMP**, select the **Enable Responses** check box to allow the interface to respond to ICMP traffic such as pings and traceroute.
12. Next to **IPv6 NDP**, select the **Enable Router Advertisement** check box to enable the interface to broadcast router advertisements.
13. To add an IP address, click **Add**.

The Add IP Address pop-up window appears.



14. In the **Address** field, type the routed interface's IP address and subnet mask using CIDR notation.
15. Optionally, if your organization uses IPv6 addresses, next to the **IPv6** field, select the **Address Autoconfiguration** check box to set the IP address of the interface automatically.
16. For **Type**, select either Normal or SFRP.
For SFRP options, see [Configuring SFRP](#) on page 304 for more information.

17. Click **OK**.

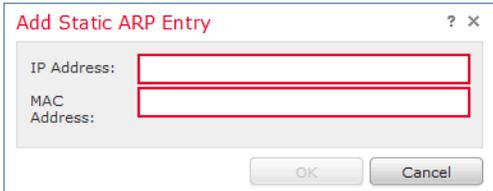
The IP address is added.

To edit an IP address, click the edit icon (✎). To delete an IP address, click the delete icon (🗑).

IMPORTANT! When adding an IP address to a routed interface of a clustered device, you must add a corresponding IP address to the routed interface on the cluster peer.

18. To add a static ARP entry, click **Add**.

The Add Static ARP Entry pop-up window appears.



19. In the **IP Address** field, type an IP address for the static ARP entry.
20. In the **MAC Address** field, type a MAC address to associate with the IP address. Enter the address using the standard format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).
21. Click **OK**.
- The static ARP entry is added.

TIP! To edit a static ARP entry, click the edit icon (✎). To delete a static ARP entry, click the delete icon (🗑).

22. Click **Save**.

The physical routed interface is configured. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Logical Routed Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

For each physical routed interface, you can add multiple logical routed interfaces. You must associate each logical interface with a VLAN tag to handle traffic received by the physical interface with that specific tag. You must assign a logical routed interface to a virtual router to route traffic.

Note that disabling the **ICMP Enable Responses** option for routed interfaces does not prevent ICMP responses in all scenarios. You can add rules to an access control policy to drop packets where the destination IP is the routed interface's IP and the protocol is ICMP. For more information about creating access control rules, see [Understanding and Writing Access Control Rules](#) on page 472. If you have enabled the **Inspect Local Router Traffic** option on the managed device, it drops the packets before they reach the host, thereby preventing any response. For more information about inspecting local router traffic, see [Understanding Advanced Device Settings](#) on page 250.

WARNING! Changing the maximum transmission unit (MTU) interrupts traffic on the device and packets are dropped. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To edit an existing routed interface, click the edit icon () next to the interface.

To add a logical routed interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the routed interface, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **Add Interface**.
The Add Interface pop-up window appears.

4. Click **Routed** to display the routed interface options.

The screenshot shows the 'Add Interface' configuration window with the 'Routed' tab selected. The configuration fields are as follows:

- Interface: s1p3
- VLAN Tag: 1
- Security Zone: None
- Virtual Router: None
- Enabled:
- MTU: 1518
- ICMP: Enable Responses
- IPv6 NDP: Enable Router Advertisement

Below these fields are two empty tables for configuration:

Address	Type
---------	------

IP Address	MAC Address
------------	-------------

Buttons for 'Add', 'Save', and 'Cancel' are visible at the bottom of the window.

5. From the **Interface** drop-down list, select the physical interface where you want to add the logical interface.
6. In the **VLAN Tag** field, type a tag value that gets assigned to inbound and outbound traffic on this interface. The value can be any integer from 1 to 4094.
7. Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.
8. Optionally, from the **Virtual Router** drop-down list, select an existing virtual router or select **New** to add a new virtual router.

Note that if you add a new virtual router, you must configure it on the Device Management page (**Devices > Device Management > Virtual Routers**) after you finish setting up the routed interface. See [Adding Virtual Routers](#) on page 307.

9. Select the **Enabled** check box to allow the routed interface to handle traffic. If you clear the check box, the interface becomes disabled and administratively taken down. If you disable a physical interface, you also disable all of the logical interfaces associated with it.

10. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed. Note that the MTU is the Layer 2 MTU/MRU and not the Layer 3 MTU.

The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

11. Next to **ICMP**, select the **Enable Responses** check box to communicate updates or error information to other routers, intermediary devices, or hosts.
12. Next to **IPv6 NDP**, select the **Enable Router Advertisement** check box to enable the interface to broadcast router advertisements.
13. To add an IP address, click **Add**.
The Add IP Address pop-up window appears.
14. In the **Address** field, type the IP address in CIDR notation.
15. Optionally, if your organization uses IPv6 addresses, next to the **IPv6** field, select the **Address Autoconfiguration** check box to set the IP address of the interface automatically.
16. For **Type**, select either Normal or SFRP.
For SFRP options, see [Configuring SFRP](#) on page 304 for more information.
17. Click **OK**.

The IP address is added.

To edit an IP address, click the edit icon (✎). To delete an IP address, click the delete icon (🗑).

IMPORTANT! When you add an IP address to a routed interface of a clustered device, you must add a corresponding IP address to the routed interface on the cluster peer.

18. To add a static ARP entry, click **Add**.
The Add Static ARP Entry pop-up window appears.
19. In the **IP Address** field, type an IP address for the static ARP entry.
20. In the **MAC Address** field, type a MAC address to associate with the IP address. Enter the address using the standard format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).

21. Click **OK**.
The static ARP entry is added.

TIP! To edit a static ARP entry, click the edit icon (). To delete a static ARP entry, click the delete icon ().

22. Click **Save**.
The logical routed interface is added. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Deleting Logical Routed Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

When you delete a logical routed interface, you remove it from the physical interface where it resides, as well as its assigned virtual router and security zone.

To delete a routed interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to delete the routed interface, click the edit icon ().
The Interfaces tab for that device appears.
3. Next to the logical routed interface you want to delete, click the delete icon ().
4. When prompted, confirm that you want to delete the interface.
The interface is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Configuring SFRP

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can configure Sourcefire Redundancy Protocol (SFRP) to achieve network redundancy for high availability on either a device cluster or individual devices. SFRP provides gateway redundancy for both IPv4 and IPv6 addresses. You can configure SFRP on routed and hybrid interfaces.

If the interfaces are configured on individual devices, they must be in the same broadcast domain. You must designate at least one of the interfaces as master and an equal number as backup. The system supports only one master and one backup per IP address. If network connectivity is lost, the system automatically promotes the backup to master to maintain connectivity.

The options you set for SFRP must be the same on all interfaces in a group of SFRP interfaces. Multiple IP addresses in a group must be in the same master/backup state. Therefore, when you add or edit an IP address, the state you set for that address propagates to all the addresses in the group. For security purposes, you must enter values for **Group ID** and **Shared Secret** that are shared among the interfaces in the group.

To enable SFRP IP addresses on a virtual router, you must also configure at least one non-SFRP IP address.

For clustered devices, you designate the shared secret and the system copies it to the cluster peer along with the SFRP IP configuration. The shared secret authenticates peer data.

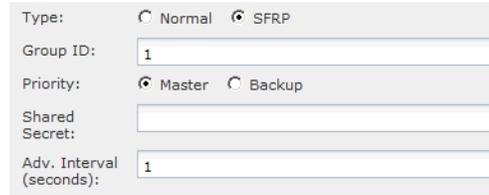
For more information about clustering devices, see [Clustering Devices](#) on page 220.

To configure SFRP:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to configure SFRP, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Next to the interface where you want to configure SFRP, click the edit icon (✎).
The Edit Interface pop-up window appears.
4. Select the type of interface where you want to configure SFRP:
 - Click **Routed** to display the routed interface options.
 - Click **Hybrid** to display the hybrid interface options.
5. You can configure SFRP while adding or editing an IP address:
 - To add an IP address, click **Add**.
 - To edit an IP address, click the edit icon (✎).The Add IP Address or Edit IP Address pop-up window appears.

6. For **Type**, select **SFRP** to display the SFRP options.



Type: Normal SFRP
Group ID:
Priority: Master Backup
Shared Secret:
Adv. Interval (seconds):

7. In the **Group ID** field, enter a value that designates a group of master or backup interfaces configured for SFRP.
8. For **Priority**, select either **Master** or **Backup** to designate the preferred interface:
- For individual devices, you must set one interface to master on one device and the other to backup on a second device.
 - For device clusters, when you set one interface as master, the other automatically becomes the backup.
9. In the **Shared Secret** field, type a shared secret.
The Shared Secret field populates automatically for a group in a device cluster.
10. In the **Adv. Interval (seconds)** field, enter an interval for route advertisements for Layer 3 traffic.
11. Click **OK**.
The IP address is added or edited.
12. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Configuring Virtual Routers

LICENSE: Control

SUPPORTED DEVICES: Series 3

Before you can use routed interfaces in a Layer 3 deployment, you must configure virtual routers and assign routed interfaces to them. A virtual router is a group of routed interfaces that route Layer 3 traffic.

See the following sections for more information about configuring virtual routers:

- [Viewing Virtual Routers](#) on page 307
- [Adding Virtual Routers](#) on page 307
- [Viewing Virtual Router Statistics](#) on page 340
- [Deleting Virtual Routers](#) on page 341

Viewing Virtual Routers

LICENSE: Control

SUPPORTED DEVICES: Series 3

The Virtual Routers tab of the Device Management page (**Devices > Device Management > Virtual Routers**) displays a list of all the virtual routers you have configured on a device. The table includes summary information about each router, as described in the following table.

Virtual Routers Table View Fields

FIELD	DESCRIPTION
Name	The name of the virtual router.
Interfaces	A list of all routed interfaces that are assigned to the virtual router. Disabling an interface from the Interfaces tab removes it.
Protocols	The protocols currently in use by the virtual router, which is one of the following: <ul style="list-style-type: none">• Static• Static, RIP• Static, OSPF

Adding Virtual Routers

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can add virtual routers from the Virtual Routers tab of the Device Management page. You can also add routers as you configure routed interfaces.

You can assign only routed and hybrid interfaces to a virtual router. If you want to create a virtual router before you configure the interfaces on your managed devices, you can create an empty virtual router and add interfaces to it later.

To maximize TCP security, you can enable strict enforcement, which blocks connections where the three-way handshake was not completed. Strict enforcement also blocks:

- non-SYN TCP packets for connections where the three-way handshake was not completed
- non-SYN/RST packets from the initiator on a TCP connection before the responder sends the SYN-ACK

- non-SYN-ACK/RST packets from the responder on a TCP connection after the SYN but before the session is established
- SYN packets on an established TCP connection from either the initiator or the responder

WARNING! Managed devices must be running at least Version 5.1 of the Sourcefire 3D System before you can enable strict TCP enforcement. Although you can apply a device configuration with strict TCP enforcement to pre-5.1 devices, those devices do not perform strict enforcement.

Note that if you change the configuration of a Layer 3 interface to a non-Layer 3 interface or remove a Layer 3 interface from the virtual router, the router may fall into an invalid state. For example, if it is used in DHCPv6, it may cause an upstream and downstream mismatch. Any changes you make to an existing virtual router may interrupt traffic on the device.

TIP! To edit an existing virtual router, click the edit icon () next to the router.

You can configure virtual routers in several different ways beyond the general options. See the following sections for more information about these configurations:

- [Setting up DHCP Relay](#) on page 310
- [Setting up Static Routes](#) on page 313
- [Setting up Dynamic Routing](#) on page 316
- [Setting up RIP Configuration](#) on page 316
- [Setting Up OSPF Configuration](#) on page 323
- [Setting up Virtual Router Filters](#) on page 335
- [Adding Virtual Router Authentication Profiles](#) on page 339

To add a virtual router:

ACCESS: Admin/Network Admin

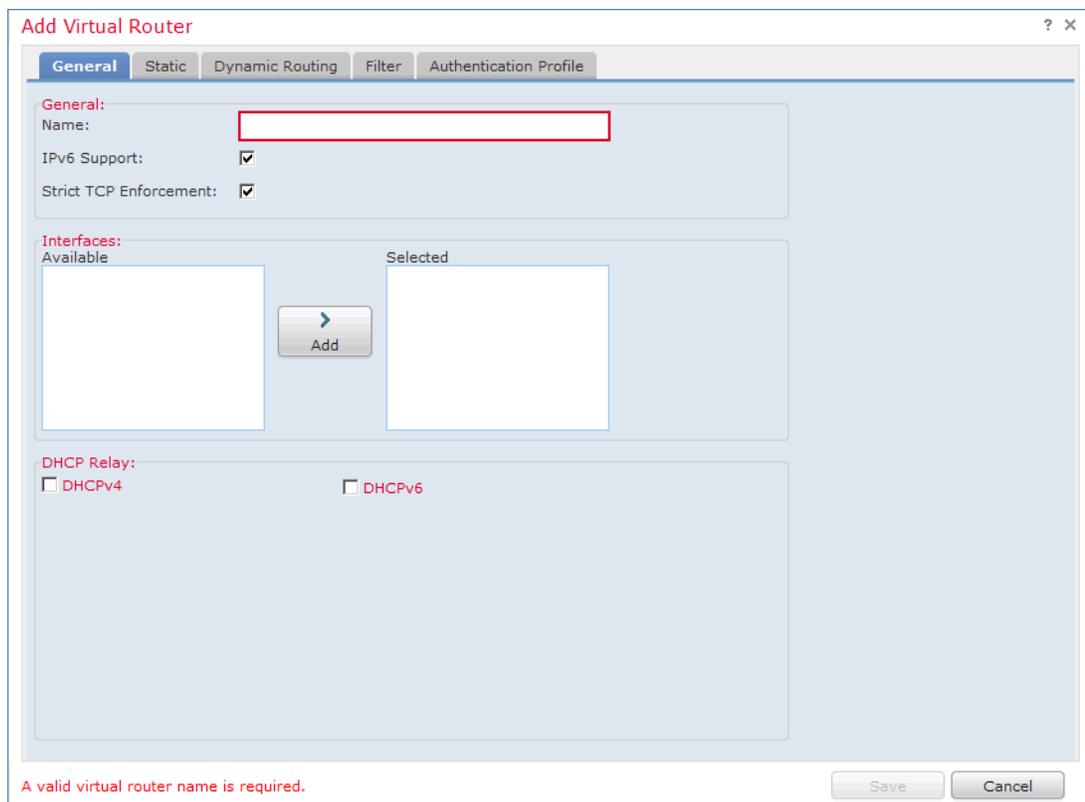
1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the virtual router, click the edit icon ().
The Interfaces tab for that device appears.

3. Click **Virtual Routers**.
The Virtual Routers tab appears.



TIP! If your devices are in a clustered stack deployment, select the stack you want to modify from the **Selected Device** drop-down list.

4. Click **Add Virtual Router**.
The Add Virtual Router pop-up window appears.



5. In the **Name** field, type a name for the virtual router. You can use alphanumeric characters and spaces.

6. To enable IPv6 static routing, OSPFv3, and RIPng on your virtual router, select the **IPv6 Support** check box. To disable these features, deselect the check box.
7. Optionally, clear **Strict TCP Enforcement** if you do not want to enable strict TCP enforcement.
This option is enabled by default.
8. Under **Interfaces**, the **Available** list contains all enabled Layer 3 interfaces, routed and hybrid, on the device that you can assign to the virtual router. Select one or more interfaces to assign to the virtual router and click **Add**.

TIP! To remove a routed or hybrid interface from the virtual router, click the delete icon (). Disabling a configured interface from the Interfaces tab also removes it.

9. Click **Save**.
The virtual router is added. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Setting up DHCP Relay

LICENSE: Control

SUPPORTED DEVICES: Series 3

DHCP provides configuration parameters to Internet hosts. A DHCP client that has not yet acquired an IP address cannot communicate directly with a DHCP server outside its broadcast domain. To allow DHCP clients to communicate with DHCP servers, you can configure DHCP relay instances to handle cases where the client is not on the same broadcast domain as the server.

You can set up DHCP relay for each virtual router you configure. By default, this feature is disabled. You can enable either DHCPv4 relay or DHCPv6 relay.

See the following sections for more information:

- [Setting up DHCPv4 Relay](#) on page 310
- [Setting up DHCPv6 Relay](#) on page 312

Setting up DHCPv4 Relay

LICENSE: Control

SUPPORTED DEVICES: Series 3

The following procedure explains how to set up DHCPv4 relay on a virtual router.

To set up DHCPv4 relay:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to set up DHCP relay, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to set up DHCP relay, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. To set up DHCP relay for DHCPv4, select the **DHCPv4** check box.



6. Under the **Servers** field, type a server IP address.
7. Click **Add**.
The IP address is added to the **Servers** field. You can add up to four DHCP servers.

TIP! To delete a DHCP server, click the delete icon (🗑) next to the server IP address.

8. In the **Max Hops** field, type the maximum number of hops from 1 to 255.
9. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Setting up DHCPv6 Relay

LICENSE: Control

SUPPORTED DEVICES: Series 3

The following procedure explains how to set up DHCPv6 relay on a virtual router.

IMPORTANT! You cannot run a DHCPv6 Relay chain through two or more virtual routers running on the same device.

To set up DHCPv6 relay:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

2. Next to the device where you want to set up DHCP relay, click the edit icon (✎).

The Interfaces tab for that device appears.

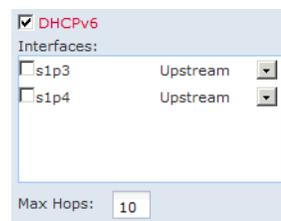
3. Click **Virtual Routers**.

The Virtual Routers tab appears.

4. Next to the virtual router where you want to set up DHCP relay, click the edit icon (✎).

The Edit Virtual Router pop-up window appears.

5. To set up DHCP relay for DHCPv6, select the **DHCPv6** check box.



✓ DHCPv6

Interfaces:

<input type="checkbox"/> s1p3	Upstream	▼
<input type="checkbox"/> s1p4	Upstream	▼

Max Hops: 10

6. In the **Interfaces** field, select the check boxes next to one or more interfaces that have been assigned to the virtual router.

TIP! You cannot disable an interface from the Interfaces tab while it is configured for DHCPv6 Relay. You must first clear the DHCPv6 Relay interfaces check box and save the configuration.

7. Next to a selected interface, click the drop-down icon and select whether the interface relays DHCP requests **Upstream**, **Downstream**, or **Both**.

Note that you must include at least one downstream interface and one upstream interface. Selecting both means that the interface is both downstream and upstream.

8. In the **Max Hops** field, type the maximum number of hops from 1 to 255
9. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Setting up Static Routes

LICENSE: Control

SUPPORTED DEVICES: Series 3

Static routing allows you to write rules about the IP addresses of traffic passing through a router. It is the simplest way of configuring path selection of a virtual router because there is no communication with other routers regarding the current topology of the network.

See the following sections for more information:

- [Understanding the Static Routes Table View](#) on page 313
- [Adding Static Routes](#) on page 314

Understanding the Static Routes Table View

LICENSE: Control

SUPPORTED DEVICES: Series 3

The Static Routes tab of the Virtual Router editor displays a list of all the static routes you have configured on a virtual router. The table includes summary information about each route, as described in the following table.

Static Routes Table View Fields

FIELD	DESCRIPTION
Enabled	Specifies whether this route is currently enabled or disabled.
Name	The name of the static route.
Destination	The destination network where traffic is routed.

Static Routes Table View Fields (Continued)

FIELD	DESCRIPTION
Type	Specifies the action that is taken for this route, which will be one of the following: <ul style="list-style-type: none"> IP — designates that the route forwards packets to the address of a neighboring router. Interface — designates that the route forwards packets to an interface through which traffic is routed to hosts on a directly connected network. Discard — designates that the static route drops packets.
Gateway	The target IP address if you selected IP as the static route type or the interface if you selected Interface as the static route type.
Preference	Determines the route selection. If you have multiple routes to the same destination, the system selects the route with the higher preference.

Adding Static Routes

LICENSE: Control

SUPPORTED DEVICES: Series 3

The following procedure explains how to add a static route.

To edit a static route, click the edit icon (✎). To delete a static route, click the delete icon (🗑).

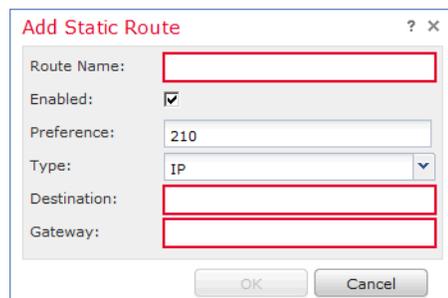
To add a static route:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the static route, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the static route, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Static** to display the static route options.

6. Click **Add Static Route**.

The Add Static Route pop-up window appears.



7. In the **Route Name** field, type a name for the static route. You can use alphanumeric characters and spaces.
8. For **Enabled**, select the check box to specify that the route is currently enabled.
9. In the **Preference** field, type a numerical value between 1 and 65535 to determine the route selection.
If you have multiple routes to the same destination, the system selects the route with the higher preference.
10. From the **Type** drop-down list, select the type of static route you are configuring.
11. In the **Destination** field, type the IP address for the destination network where traffic should be routed.
12. In the **Gateway** field, you have two options:
 - If you selected **IP** as the selected static route type, type an IP address.
 - If you selected **Interface** as the selected static route type, select an enabled interface from the drop-down list.

TIP! Interfaces you have disabled from the Interfaces tab are not available; disabling an interface you have added removes it from the configuration.

13. Click **OK**.

The static route is added.

14. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Setting up Dynamic Routing

LICENSE: Control

SUPPORTED DEVICES: Series 3

Dynamic, or adaptive, routing uses a routing protocol to alter the path that a route takes in response to a change in network conditions. The adaptation is intended to allow as many routes as possible to remain valid, that is, have destinations that can be reached in response to the change. This allows the network to “route around” damage, such as loss of a node or a connection between nodes, so long as other path choices are available. You can configure a router with no dynamic routing, or you can configure the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) routing protocol.

See the following sections for more information:

- [Setting up RIP Configuration](#) on page 316
- [Setting Up OSPF Configuration](#) on page 323

Setting up RIP Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

Routing Information Protocol (RIP) is a dynamic routing protocol, designed for small IP networks, that relies on hop count to determine routes. The best routes use the fewest number of hops. The maximum number of hops allowed for RIP is 15. This hop limit also limits the size of the network that RIP can support.

See the following sections for more information on configuring RIP:

- [Adding Interfaces for RIP Configuration](#) on page 316
- [Configuring Authentication Settings for RIP Configuration](#) on page 318
- [Configuring Advanced Settings for RIP Configuration](#) on page 319
- [Adding Import Filters for RIP Configuration](#) on page 320
- [Adding Export Filters for RIP Configuration](#) on page 322

Adding Interfaces for RIP Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

While configuring RIP, you must select interfaces from those already included in the virtual router, where you want to configure RIP. Disabled interfaces are not available.

To edit a RIP interface, click the edit icon (✎). To delete a RIP interface, click the delete icon (🗑).

To add an interface for RIP configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the RIP interface, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the RIP interface, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **RIP** to display the RIP options.

The screenshot shows the 'Edit Virtual Router, "My Virtual Router"' window with the 'Dynamic Routing' tab selected. The 'Routing Type' is set to 'RIP'. The 'Interfaces' table is empty. The 'Authentication' section is expanded, showing a 'Profile' dropdown set to 'Select a Profile...'. The 'Advanced' section includes fields for 'Preference' (120), 'Infinity' (16), 'Period' (30), 'Honor' (Never), 'Timeout Time' (120), and 'Garbage Time' (300). The 'Import Filters' and 'Export Filters' sections are also empty. 'Save' and 'Cancel' buttons are at the bottom right.

Name	Metric	Mode
------	--------	------

Routing Type: OFF RIP OSPF

Authentication Profile: Select a Profile...

Advanced Preference: 120 Infinity: 16
Period: 30 Honor: Never
Timeout Time: 120
Garbage Time: 300

Name	Action
------	--------

Name	Action
------	--------

Save Cancel

7. Under **Interfaces**, click the add icon ().
The Add an Interface pop-up window appears.
8. From the **Name** drop-down list, select the interface where you want to configure RIP.

TIP! Interfaces you have disabled from the Interfaces tab are not available; disabling an interface you have added removes it from the configuration.

9. In the **Metric** field, type a metric for the interface. When routes from different RIP instances are available and all of them have the same preference, the route with the lowest metric becomes the preferred route.
10. From the **Mode** drop-down list, select one of the following options:
 - **Multicast** — default mode where RIP multicasts the entire routing table to all adjacent routers at a specified address.
 - **Broadcast** — forces RIP to use broadcast (for example, RIPv1) even though multicast mode is possible.
 - **Quiet** — RIP will not transmit any periodic messages to this interface.
 - **No Listen** — RIP will send to this interface but not listen to it.
11. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Configuring Authentication Settings for RIP Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

RIP authentication uses one of the authentication profiles you configured on the virtual router. For more information about configuring authentication profiles, see [Adding Virtual Router Authentication Profiles](#) on page 339.

To configure authentication settings for RIP configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the RIP authentication profile, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.

4. Next to the virtual router where you want to add the RIP authentication profile, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **RIP** to display the RIP options.
7. Under **Authentication**, use the Profile drop-down list to select an existing virtual router authentication profile or select **None**.
8. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Configuring Advanced Settings for RIP Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can configure several advanced RIP settings pertaining to various timeout values and other features that affect the behavior of the protocol.

WARNING! Changing any of the advanced RIP settings to incorrect values can prevent the router from communicating successfully with other RIP routers.

To configure advanced settings for RIP configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to edit the RIP advanced settings, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to edit the RIP advanced settings, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **RIP** to display the RIP options.

7. In the **Preference** field, type a numerical value (higher is better) for the preference of the routing protocol. The system prefers routes learned through RIP over static routes.
8. In the **Period** field, type the interval, in seconds, between periodic updates. A lower number determines faster convergence, but larger network load.
9. In the **Timeout Time** field, type a numerical value that specifies how old routes must be, in seconds, before being considered unreachable.
10. In the **Garbage Time** field, type a numerical value that specifies how old routes must be, in seconds, before being discarded.
11. In the **Infinity** field, type a numerical value that specifies a value for infinity distance in convergence calculations. Larger values will make protocol convergence slower.
12. From the **Honor** drop-down list, select one of the following options to designate when requests for dumping routing tables should be honored:
 - **Always** — always honor requests
 - **Neighbor** — only honor requests sent from a host on a directly connected network
 - **Never** — never honor requests
13. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Import Filters for RIP Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can add an import filter to designate which routes are accepted or rejected from RIP into the route table. Import filters are applied in the order they appear in the table.

When adding an import filter, you use one of the filters you configured on the virtual router. For more information about configuring filters, see [Setting up Virtual Router Filters](#) on page 335.

TIP! To edit a RIP import filter, click the edit icon (✎). To delete a RIP import filter, click the delete icon (🗑).

To add an import filter for RIP configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the RIP virtual router filter, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the RIP virtual router filter, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **RIP** to display the RIP options.
7. Under **Import Filters**, click the add icon (⊕).
The Add an Import Filter pop-up window appears.
8. From the **Name** drop-down list, select the filter you want to add as an import filter.
9. Next to **Action**, select **Accept** or **Reject**.
10. Click **OK**.
The import filter is added.

TIP! To change the order of the import filters, click the move up (▲) and move down (▼) icons as needed. You can also drag the filters up or down in the list.

11. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Export Filters for RIP Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can add an export filter to define which routes will be accepted or rejected from the route table to RIP. Export filters are applied in the order they appear in the table.

When adding an export filter, you use one of the filters you configured on the virtual router. For more information about configuring filters, see [Setting up Virtual Router Filters](#) on page 335.

To add an export filter for RIP configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the RIP virtual router filter, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the RIP virtual router filter, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **RIP** to display the RIP options.
7. Under **Export Filters**, click the add icon (+).
8. From the **Name** drop-down list, select the filter you want to add as an export filter.
9. Next to **Action**, select **Accept** or **Reject**.
10. Click **OK**.
The export filter is added.

TIP! To change the order of the export filters, click the move up (▲) and move down (▼) icons as needed. You can also drag the filters up or down in the list.

11. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Setting Up OSPF Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

Open Shortest Path First (OSPF) is an adaptive routing protocol that defines routes dynamically by obtaining information from other routers and advertising routes to other routers using link state advertisements. The router keeps information about the links between it and the destination to make routing decisions. OSPF assigns a cost to each routed interface, and considers the best routes to have the lowest costs.

See the following sections for more information:

- [Setting Up OSPF Routing Areas](#) on page 323
- [Adding Import Filters for OSPF Configuration](#) on page 333
- [Adding Export Filters for OSPF Configuration](#) on page 334

Setting Up OSPF Routing Areas

LICENSE: Control

SUPPORTED DEVICES: Series 3

An OSPF network may be structured, or subdivided, into routing areas to simplify administration and optimize traffic and resource use. Areas are identified by 32-bit numbers, expressed either simply in decimal or often in octet-based dot-decimal notation.

By convention, area zero or 0.0.0.0 represents the core or backbone region of an OSPF network. You may choose to identify other areas. Often, administrators select the IP address of a main router in an area as the area's identification. Each additional area must have a direct or virtual connection to the backbone OSPF area. Such connections are maintained by an interconnecting router, known as the area border router (ABR). An ABR maintains separate link state databases for each area it serves and maintains summarized routes for all areas in the network.

See the following sections for more information on setting up OSPF areas:

- [Adding OSPF Areas](#) on page 324
- [Adding OSPF Area Interfaces](#) on page 327
- [Adding OSPF Area Vlinks](#) on page 331

Adding OSPF Areas

LICENSE: Control

SUPPORTED DEVICES: Series 3

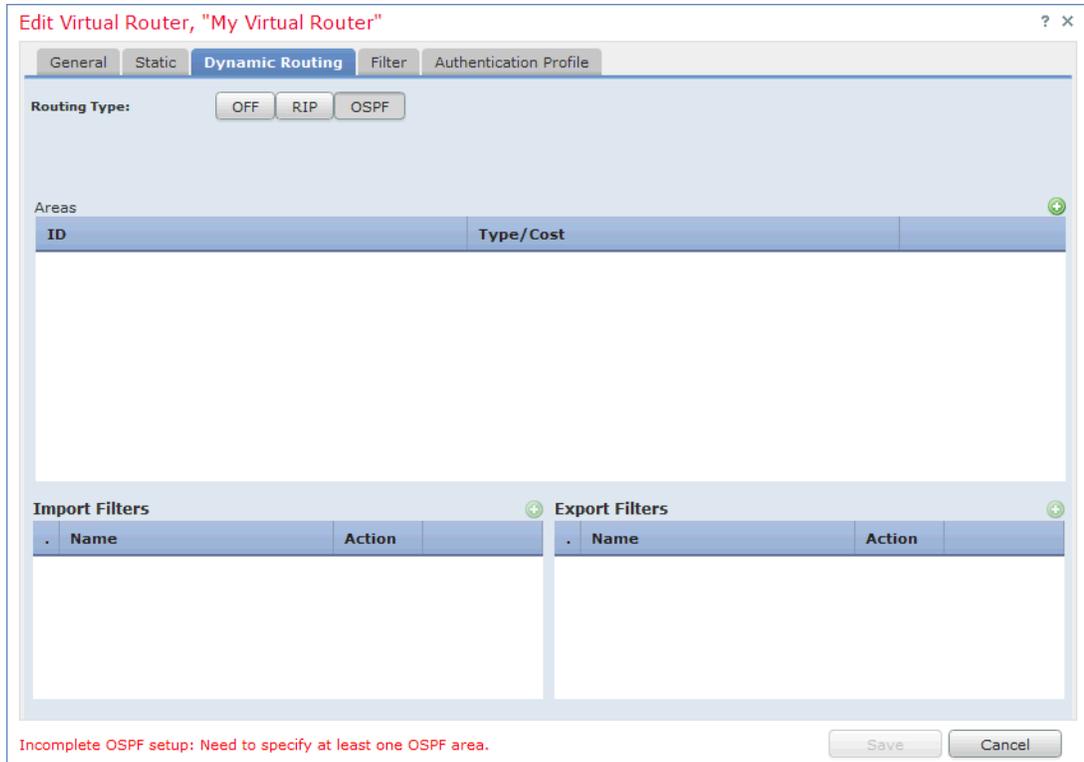
The following procedure explains how to add an OSPF area and configure general settings.

To add an OSPF area:

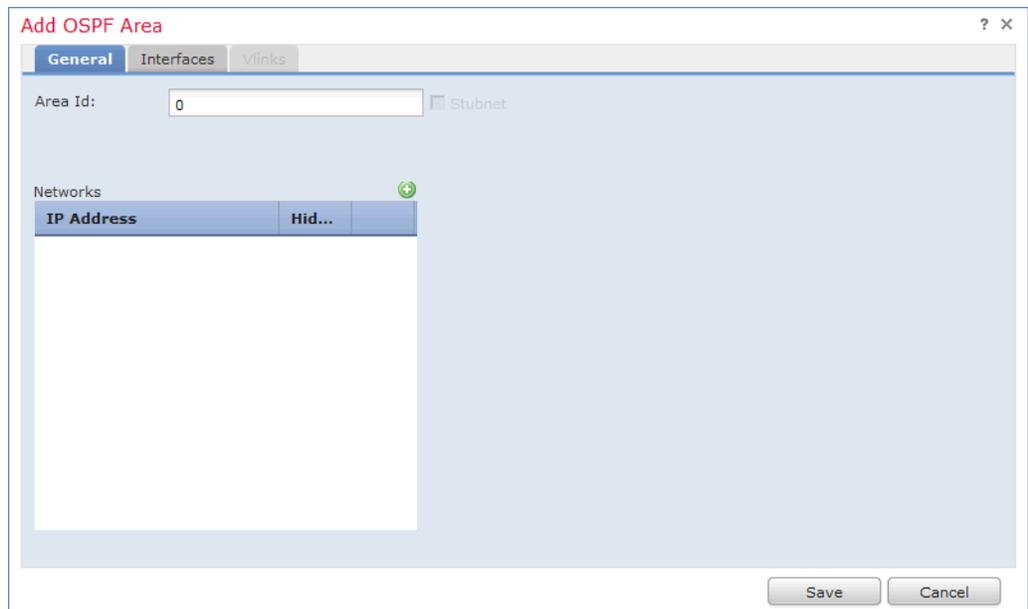
ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to edit the OSPF general options, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to edit the OSPF general options, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.

6. Click **OSPF** to display the OSPF options.



7. Under **Areas**, click the add icon (+).
The Add OSPF Area pop-up window appears.



8. In the **Area Id** field, type a numerical value for the area. This value can be either an integer or an IPv4 address.
9. Optionally, select the **Stubnet** check box to designate that the area does not receive router advertisements external to the autonomous system and routing from within the area is based entirely on a default route. If you clear the check box, the area becomes a backbone area or otherwise non-stub area.

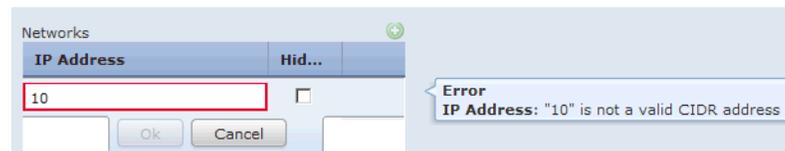
The Default cost field and Stubnet field appear.

10. In the **Default cost** field, type a cost associated with the default route for the area.
11. Under **Stubnets**, click the add icon ().
12. In the **IP Address** field, type an IP address in CIDR notation.
13. Select the **Hidden** check box to indicate that the stubnet is hidden. Hidden stubnets are not propagated into other areas.
14. Select the **Summary** check box to designate that default stubnets that are subnetworks of this stubnet are suppressed.
15. In the **Stub cost** field, type a value that defines the cost associated with routing to this stub network.
16. Click **OK**.

The stubnet is added.

TIP! To edit a stubnet, click the edit icon (). To delete a stubnet, click the delete icon ().

17. Optionally, under **Networks**, click the add icon ().



18. In the **IP Address** field, type an IP address in CIDR notation for the network.
19. Select the **Hidden** check box to indicate that the network is hidden. Hidden networks are not propagated into other areas.

20. Click **OK**.

The network is added.

TIP! To edit a network, click the edit icon (). To delete a network, click the delete icon ().

21. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding OSPF Area Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can configure a subset of the interfaces assigned to the virtual router for OSPF. The following list describes the options you can specify on each interface.

Interfaces

Select the interface where you want to configure OSPF. Interfaces you have disabled from the Interfaces tab are not available.

Type

Select the type of OSPF interface from the following choices:

- **Broadcast** — On broadcast networks, flooding and hello messages are sent using multicasts, a single packet for all the neighbors. The option designates a router to be responsible for synchronizing the link state databases and originating network link state advertisements. This network type cannot be used on physically non-broadcast multiple-access (NBMP) networks and on unnumbered networks without proper IP prefixes.
- **Point-to-Point (PtP)** — Point-to-point networks connect just two routers together. No election is performed and no network link state advertisement is originated, which makes it simpler and faster to establish. This network type is useful not only for physically PtP interfaces, but also for broadcast networks used as PtP links. This network type cannot be used on physically NBMP networks.

- Non-Broadcast — On NBMP networks, the packets are sent to each neighbor separately because of the lack of multicast capabilities. Similar to broadcast networks, the option designates a router, which plays a central role in the propagation of link state advertisements. This network type cannot be used on unnumbered networks.
- Autodetect — The system determines the correct type based on the specified interface.

Cost

Specify the output cost of the interface.

Stub

Specify whether the interface should listen for OSPF traffic and transmit its own traffic.

Priority

Enter a numerical value that specifies the priority value used in designated router election. On every multiple access network, the system designates a router and backup router. These routers have some special functions in the flooding process. Higher priority increases preferences in this election. You cannot configure a router with a priority of 0.

Nonbroadcast

Specify whether hello packets are sent to any undefined neighbors. This switch is ignored on any NBMA network.

Authentication

Select the OSPF authentication profile that this interface uses from one of the authentication profiles you configured on the virtual router or select **None**. For more information about configuring authentication profiles, see [Adding Virtual Router Authentication Profiles](#) on page 339.

Hello Interval

Type the interval, in seconds, between the sending of hello messages.

Poll

Type the interval, in seconds, between the sending of hello messages for some neighbors on NBMA networks.

Retrans Interval

Type the interval, in seconds, between retransmissions of unacknowledged updates.

Retrans Delay

Type the estimated number of seconds it takes to transmit a link state update packet over the interface.

Wait Time

Type the number of seconds that the router waits between starting election and building adjacency.

Dead Interval

Type the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it. If this value is defined, it overrides the value calculated from dead count.

Dead Count

Type a numerical value that when multiplied by the hello interval specifies the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it.

To edit an OSPF area interface, click the edit icon (✎). To delete an OSPF area interface, click the delete icon (🗑). Disabling a configured interface from the Interfaces tab also deletes it.

IMPORTANT! You can select only one interface for use in an OSPF area.

To add an OSPF area interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the OSPF interface, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the OSPF interface, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **OSPF** to display the OSPF options.

7. Under **Areas**, click the add icon (+).
The Add OSPF Area pop-up window appears.
8. Click **Interfaces**.
The Interfaces tab appears.
9. Click the add icon (+).
The Add OSPF Area Interface pop-up window appears.

10. Take any of the actions as described in [Adding OSPF Area Interfaces](#) on page 327.
11. Optionally under **Neighbors**, click the add icon (+).

12. In the **IP address** field, type an IP address for the neighbor receiving hello messages on non-broadcast networks from this interface.

13. Select the **Eligible** check box to indicate that the neighbor is eligible to receive messages.
14. Click **OK**.
The neighbor is added.

TIP! To edit a neighbor, click the edit icon (✎). To delete a neighbor, click the delete icon (🗑).

15. Click **OK**.
The OSPF area interface is added.
16. Click **Save**.
The OSPF area is saved.
17. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding OSPF Area Vlinks

LICENSE: Control

SUPPORTED DEVICES: Series 3

All areas in an OSPF autonomous system must be physically connected to the backbone area. In some cases where this physical connection is not possible, you can use a vlink to connect to the backbone through a non-backbone area. Vlinks can also be used to connect two parts of a partitioned backbone through a non-backbone area.

You must add a minimum of two OSPF areas before you can add a vlink.

To add an OSPF area vlink:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the OSPF vlink, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.

4. Next to the virtual router where you want to add the OSPF interface, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **OSPF** to display the OSPF options.
7. Under **Areas**, click the add icon (+).
8. Click **Vlinks**.
The Vlinks tab appears.
9. Click the add icon (+).
The Add OSPF Area Vlink pop-up window appears.

The screenshot shows a dialog box titled "Add OSPF Area Vlink". It contains the following fields and values:

Router ID:	<input type="text"/>
Authentication:	none
Timing	
Hello Interval:	10
Retrans Interval:	5
Wait Time:	40
Dead Interval:	40
Dead Count:	4

Buttons: OK, Cancel

10. In the **Router ID** field, type an IP address for the router.
11. From the **Authentication** drop-down list, select the authentication profile the vlink will use.
12. In the **Hello Interval** field, type the interval, in seconds, between sending of hello messages.
13. In the **Retrans Interval** field, type the interval, in seconds, between retransmissions of unacknowledged updates.
14. In the **Wait Time** field, type the number of seconds that the router waits between starting election and building adjacency.
15. In the **Dead Interval** field, type the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it. If this value is defined, it overrides the value calculated from dead count.
16. In the **Dead Count** field, type a numerical value that when multiplied by the hello interval, specifies the number of seconds that the router waits before declaring a neighbor down when not receiving messages from it.

17. Click **OK**.
The OSPF area vlink is added.
18. Click **Save**.
The OSPF area is saved.
19. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Import Filters for OSPF Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can add an import filter to define which routes are accepted or rejected from OSPF into the route table. Import filters are applied in the order they appear in the table.

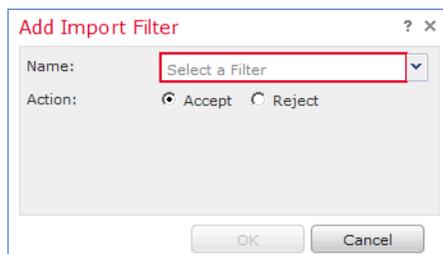
When adding an import filter, you use one of the filters you configured on the virtual router. For more information about configuring filters, see [Setting up Virtual Router Filters](#) on page 335.

To add an import filter for OSPF configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the OSPF virtual router filter, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the OSPF virtual router filter, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **OSPF** to display the OSPF options.

7. Under **Import Filters**, click the add icon ().
The Add Import Filter pop-up window appears.



8. From the **Name** drop-down list, select the filter you want to add as an import filter.
9. Next to **Action**, select **Accept** or **Reject**.
10. Click **OK**.
The import filter is added.

TIP! To change the order of the import filters, click the move up () and move down () icons as needed. You can also drag the filters up or down in the list.

11. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Export Filters for OSPF Configuration

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can add an export filter to define which routes will be accepted or rejected from the route table to OSPF. Export filters are applied in the order they appear in the table.

When adding an export filter, you use one of the filters you configured on the virtual router. For more information about configuring filters, see [Setting up Virtual Router Filters](#) on page 335.

To add an export filter for OSPF configuration:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.

2. Next to the device where you want to add the OSPF virtual router filter, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the OSPF virtual router filter, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Dynamic Routing** to display the dynamic routing options.
6. Click **OSPF** to display the OSPF options.
7. Under **Export Filters**, click the add icon (⊕).
The Add an Export Filter pop-up window appears.
8. From the **Name** drop-down list, select the filter you want to add as an export filter.
9. Next to **Action**, select **Accept** or **Reject**.
10. Click **OK**.
The export filter is added.

TIP! To change the order of the export filters, click the move up (▲) and move down (▼) icons as needed. You can also drag the filters up or down in the list.

11. Click **Save**.
Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Setting up Virtual Router Filters

LICENSE: Control

SUPPORTED DEVICES: Series 3

Filters provide a way to match routes for importing into the virtual router's route table and for exporting routes to dynamic protocols. You can create and manage a

list of filters. Each filter defines specific criteria to look for in routes that are defined statically or received from a dynamic protocol.

TIP! To edit a virtual router filter, click the edit icon (). To delete a virtual router filter, click the delete icon ().

The Filter tab of the Virtual Router editor displays a table listing of all the filters you have configured on a virtual router. The table includes summary information about each filter, as described in the following table.

Virtual Router Filters Table View Fields

FIELD	DESCRIPTION
Name	The name of the filter.
Protocol	The protocol that the route originates from: <ul style="list-style-type: none"> • Static — The route originates as a local static route. • RIP — The route originates from a dynamic RIP configuration. • OSPF — The route originates from a dynamic OSPF configuration.
From Router	The router IP addresses that this filter attempts to match in a router. You must enter this value for static and RIP filters.
Next Hop	The next hop where packets using this route are forwarded. You must enter this value for static and RIP filters.
Destination Type	The type of destination where packets are sent: <ul style="list-style-type: none"> • Router • Device • Discard
Destination Network	The networks that this filter attempts to match in a route.
OSPF Path Type	Applies only to OSPF protocol. The path type can be one of the following: <ul style="list-style-type: none"> • Ext-1 • Ext-2 • Inter Area • Intra Area
OSPF Router ID	Applies only to OSPF protocol. The router ID of the router advertising that route/network.

To add a virtual router filter:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the virtual filter router, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the virtual filter router, click the edit icon (✎).
The Edit Virtual Router pop-up window appears.
5. Click **Filter** to display the Filter options.
6. Click **Add Filter**.
The Create Filter pop-up window appears.

The screenshot shows the 'Create Filter' dialog box. It has a title bar with 'Create Filter' and a close button. The main area is divided into sections. At the top is a 'Name:' field. Below that is the 'General' section, which includes a 'Protocol:' section with radio buttons for 'All', 'Static', 'RIP', and 'OSPF'. There are also 'Destination Type:' checkboxes for 'Router', 'Device', and 'Discard'. To the right of these are 'From Router:' and 'Next Hop:' fields, each with an 'Add' button. Below these is a 'Destination Network:' field with an 'Add' button. The 'OSPF' section includes a 'Path Type:' section with checkboxes for 'Ext-1', 'Ext-2', 'Inter-Area', and 'Intra-Area', and a 'Router ID:' field with an 'Add' button. At the bottom of the dialog, there is a red message that says 'Incomplete filter setup.' and two buttons: 'OK' and 'Cancel'.

7. In the **Name** field, type a name for the filter. You can use alphanumeric characters only.

8. Under **Protocol**, select **All** or select the protocol that applies to the filter.
9. If you selected All, Static, or RIP as the Protocol, under **From Router**, type the router IP addresses that this filter will attempt to match in a route.

Note that you can also enter a /32 CIDR block for IPv4 addresses and a /128 prefix length for IPv6 addresses. All other address blocks are invalid for this field.
10. Click **Add**.

The **From Router** field is populated.
11. If you selected All, Static, or RIP as the Protocol, under **Next Hop**, type the IP addresses for the gateways that this filter will attempt to match in a route.

Note that you can also enter a /32 CIDR block for IPv4 addresses and a /128 prefix length for IPv6 addresses. All other address blocks are invalid for this field.
12. Click **Add**.

The **Next Hop** field is populated.
13. Under **Destination Type**, select the options that apply to the filter.
14. Under **Destination Network**, type the IP address of the network that this filter will attempt to match in a route.
15. Click **Add**.

The **Destination Network** field is populated.
16. If you selected All or OSPF as the Protocol, under **Path Type**, select the options that apply to the filter.

You must select at least one path type.
17. If you selected OSPF as the Protocol, under **Router ID**, type the IP address that serves as the router ID of the router advertising the route/network.
18. Click **Add**.

The **Router ID** field is populated.
19. Click **OK**.

The filter is added.
20. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Virtual Router Authentication Profiles

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can set up Authentication Profiles for use in RIP and OSPF configurations. You can configure a simple password or specify a shared cryptographic key. Simple passwords allow for every packet to carry eight bytes of the password. The system ignores received packets lacking this password. Cryptographic keys allow for validation, a 16-byte long digest generated from a password to be appended to every packet.

Note that for OSPF, each area can have a different authentication method. Therefore, you create authentication profiles that can be shared among many areas. You cannot add authentication for OSPFv3.

TIP! To edit an authentication profile, click the edit icon (). To delete an authentication profile, click the delete icon ().

To add a virtual router authentication profile:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the virtual router authentication profile, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router where you want to add the virtual router authentication profile, click the edit icon ().
The Edit Virtual Router pop-up window appears.
5. Click **Authentication Profile**.
The Authentication Profile tab appears.

6. Click **Add Authentication Profile**.

The Add Authentication Profile pop-up window appears.



7. In the **Authentication Profile Name** field, type a name for the authentication profile.
8. From the **Authentication Type** drop down list, select **simple** or **cryptographic**.
9. In the **Password** field, type a secure password.
10. In the **Confirm Password** field, type the password again to confirm it.
11. Click **OK**.

The authentication profile is added.

12. Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Viewing Virtual Router Statistics

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can view runtime statistics for each virtual router. The statistics display unicast packets, packets dropped, and separate routing tables for IPv4 and IPv6 addresses.

To view virtual router statistics:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.

2. Next to the device where you want to view the virtual router statistics, click the edit icon (✎).

The Interfaces tab for that device appears.

3. Click **Virtual Routers**.

The Virtual Routers tab appears.

4. Next to the virtual router where you want to view the router statistics, click the view icon ().
The Statistics pop-up window appears.
5. Click **OK** to close the window.

Deleting Virtual Routers

LICENSE: Control

SUPPORTED DEVICES: Series 3

When you delete a virtual router, any routed interfaces assigned to the router become available for inclusion in another router.

To delete a virtual router:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to delete the virtual router, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **Virtual Routers**.
The Virtual Routers tab appears.
4. Next to the virtual router that you want to delete, click the delete icon ().
5. When prompted, confirm that you want to delete the virtual router.
The virtual router is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Configuring NAT on a Virtual Router

LICENSE: Control

SUPPORTED DEVICES: Series 3

In a Layer 3 deployment, you can configure routing with network address translation (NAT). You can configure NAT in different ways to manage specific network needs.

IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices from the NAT tab on the Device Management page.

You can configure NAT to expose an internal server to an external network. In this configuration, you define a static translation from an external IP address to an internal IP address so that the system can access an internal server from outside the network. Traffic sent to the server targets the external IP address or IP address and port, and is translated into the internal IP address or IP address and port. Return traffic from the server is translated back to the external address.

You can configure NAT to allow an internal host or server to connect to an external application. In this configuration, you define a static translation from an internal address to an external address. This definition allows the internal host or server to initiate a connection to an external application that is expecting the internal host or server to have a specific IP address and port. Therefore, the system cannot dynamically allocate the address of the internal host or server.

You can configure NAT to hide private network addresses from an external network by using a block of IP addresses. This becomes useful if you want to obscure your internal network addresses and have sufficient external IP addresses to satisfy your internal network needs. In this configuration, you create a dynamic translation that automatically converts the source IP address of any outgoing traffic to an unused IP address from your externally facing IP addresses.

You can configure NAT to hide private network addresses from an external network using a limited block of IP addresses and port translation. This becomes useful if you want to obscure your internal network addresses, but have an insufficient number of external IP addresses to satisfy your internal network needs. In this configuration, you create a dynamic translation that automatically converts the source IP address and port of outgoing traffic to an unused IP address and port from your externally facing IP addresses.

WARNING! Sourcefire strongly recommends that if you configure dynamic NAT on an interface, you should also add an access control rule to block all inbound connections through that interface. In addition, if you configure static NAT on the same port, you must configure an access control rule with a lower rule number to allow inbound connections for the targeted internal host or server. For more information, see [Creating and Editing Access Control Rules](#) on page 474.

See the following sections for more information:

- [Understanding the NAT Entry List](#) on page 343
- [Adding NAT Entries](#) on page 344
- [Deleting NAT Entries](#) on page 350

Understanding the NAT Entry List

LICENSE: Control

SUPPORTED DEVICES: Series 3

The NAT entry list displays a table listing of all the NAT entries you have configured on a device. The table includes summary information about each NAT entry, as described in the NAT Entry Table View Fields table.

IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices while configuring virtual routers.

NAT Entry Table View Fields

FIELD	DESCRIPTION
Name	The name of the NAT entry.
Interface	The hybrid or routed interface where packets are translated. The type of NAT entry you configure determines how the packets are translated. If you specify an external interface, the system translates all traffic matching the NAT entry type. If you specify an internal interface, then the system only translates traffic passing through that interface.
Type	The type of NAT entry configured, which is one of the following: <ul style="list-style-type: none">• Dynamic IP Only — indicates a dynamic translation of all outbound packets using a range of IP addresses.• Dynamic IP + Port — indicates a dynamic translation of all outbound packets using a reduced range of IP addresses and port translation.• Static — indicates a one-to-one translation. The direction of the original packet informs the translation. An inbound packet will have the destination IP address and optionally the port translated, while an outbound packet will have the source IP address and optionally the port translated. Both inbound and outbound packets will also have the return packets translated in reverse so that the traffic routes properly.
Original IP Address	The IP address to be translated.
Original Service	The service protocol and port number that are being translated.
Translated IP Address	The translated address applied to the original IP address. This is a single address for static NAT entries, and a range or CIDR block for dynamic NAT entries.
Translated Service	The translated port number applied to the original port number, if applicable.

If the list of NAT entries becomes too long, you can find a specific entry by either searching or filtering the table. You can also combine filters with a search.

You can search for an existing NAT entry in the table by any of the columns in the NAT entry table. If a match is found, the matching entry becomes highlighted in the table. You can move up and down between searched values using the arrows next to the search box.

To filter the NAT entry table:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to find the NAT entry, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **NAT**.
The NAT tab appears.
4. Hover your pointer over the **Interface** or **Type** column header.
An arrow appears.
5. Click the arrow.
The **Filters** option appears.
6. Hover your pointer over the **Filters** option.
The **Filters** criteria appear.
7. Select one or more options from the list.
The NAT entry table displays only the entries that matched the selected filters.

TIP! To delete the filters, hover your pointer over the column header and clear the **Filters** check box or an individual filter's check box.

Adding NAT Entries

LICENSE: Control

SUPPORTED DEVICES: Series 3

Sourcefire strongly recommends that if you configure dynamic NAT on an interface, you should also add an access control rule to block all inbound connections through that interface. In addition, if you configure static NAT on the same port, you must configure an access control rule with a lower rule number to allow inbound connections for the targeted internal host or server. For more

information, see [Creating and Editing Access Control Rules](#) on page 474.

IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices while configuring virtual routers.

See the following sections for more information about adding NAT entries:

- [Adding Static NAT Entries](#) on page 345
- [Adding Dynamic IP Only NAT Entries](#) on page 347
- [Adding Dynamic IP and Port NAT Entries](#) on page 348

Adding Static NAT Entries

LICENSE: Control

SUPPORTED DEVICES: Series 3

The following procedure explains how to add a static NAT entry. To edit an existing NAT entry, click the edit icon () next to the entry. You can also right-click the entry and select **Edit**.

IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices while configuring virtual routers.

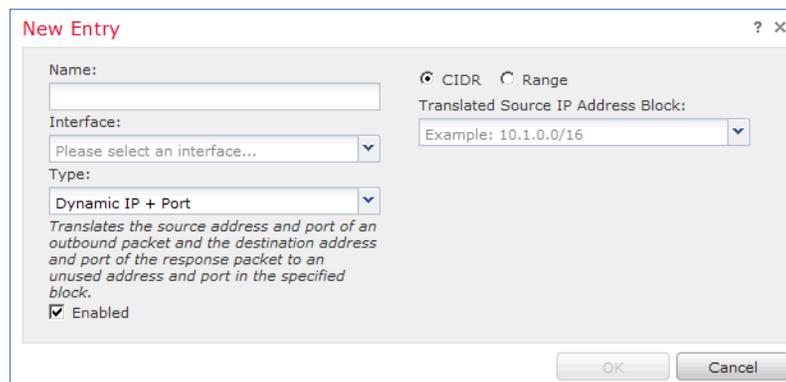
To add a NAT entry:

ACCESS: Admin/Network Admin

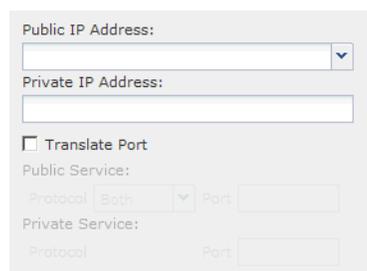
1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the NAT entry, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **NAT**.
The NAT tab appears.

4. Click **New Entry**.

The New Entry pop-up window appears.



5. In the **Name** field, type a name for the NAT entry.
6. From the **Interface** drop-down list, select a routed or hybrid interface.
7. From the **Type** drop-down list, select **Static** to display the static NAT options.



8. Optionally, you can clear the **Enabled** check box, which means the entry does not affect traffic while disabled.

TIP! You can also enable or disable a NAT entry from the entry list by right-clicking the entry, selecting **State**, and the appropriate option.

9. In the **Public IP Address** field, you can do one of the following:
 - Select the IP address to be translated from the drop-down list.
 - Type the IP address to be translated.
10. In the **Private IP Address** field, type the translated IP address that will be applied to the public IP address.
11. Optionally, select the **Translate Port** check box if you also want to translate port numbers statically.
 - Under **Public Service**, select a **Protocol** and type a **Port** to be translated.
 - Under **Private Service**, type the translated **Port** that will be applied to the original port number.

12. Click **OK**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Dynamic IP Only NAT Entries

LICENSE: Control

SUPPORTED DEVICES: Series 3

Sourcefire strongly recommends that if you configure dynamic NAT on an interface, you should also add an access control rule to block all inbound connections through that interface. In addition, if you configure static NAT on the same port, you must configure an access control rule with a lower rule number to allow inbound connections for the targeted internal host or server. For more information, see [Creating and Editing Access Control Rules](#) on page 474.

The following procedure explains how to add a dynamic IP only NAT entry. To edit an existing NAT entry, click the edit icon () next to the entry. You can also right-click the entry and select **Edit**.

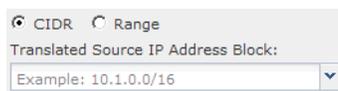
IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices while configuring virtual routers.

To add a dynamic IP only NAT entry:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the NAT entry, click the edit icon ()
The Interfaces tab for that device appears.
3. Click **NAT**.
The NAT tab appears.
4. Click **New Entry**.
The New Entry pop-up window appears.
5. In the **Name** field, type a name for the NAT entry.
6. From the **Interface** drop-down list, select a routed or hybrid interface.

7. From the **Type** drop-down list, select **Dynamic IP Only** to display the dynamic IP only NAT options.



The screenshot shows a configuration interface with two radio buttons: 'CIDR' (selected) and 'Range'. Below them is a text field labeled 'Translated Source IP Address Block:' containing the text 'Example: 10.1.0.0/16' and a downward-pointing arrow indicating a dropdown menu.

8. Optionally, you can clear the **Enabled** check box, which means the entry does not affect traffic while disabled.

TIP! You can also enable or disable a NAT entry from the entry list by right-clicking the entry, selecting **State**, and the appropriate option.

9. Select either **CIDR** or **Range**:
 - If you selected **CIDR**, in the **Translated Source IP Address Block** field, either select an IP address from the drop-down list or type an IP address in CIDR notation. See [IP Address Conventions](#) on page 57.
 - If you selected **Range**, type a **First Translated Source IP Address** for the beginning of the range and a **Last Translated Source IP Address** for the end of the range.

IMPORTANT! On a 3D8250, you must have at least two IP addresses.

10. Click **OK**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Adding Dynamic IP and Port NAT Entries

LICENSE: Control

SUPPORTED DEVICES: Series 3

Sourcefire strongly recommends that if you configure dynamic NAT on an interface, you should also add an access control rule to block all inbound connections through that interface. In addition, if you configure static NAT on the same port, you must configure an access control rule with a lower rule number to allow inbound connections for the targeted internal host or server. For more information, see [Creating and Editing Access Control Rules](#) on page 474.

The following procedure explains how to add a dynamic IP and port NAT entry. To edit an existing NAT entry, click the edit icon () next to the entry. You can also right-click the entry and select **Edit**.

IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices while configuring virtual routers.

To add a NAT entry:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to add the NAT entry, click the edit icon ().
The Interfaces tab for that device appears.
3. Click **NAT**.
The NAT tab appears.
4. Click **New Entry**.
The New Entry pop-up window appears.
5. In the **Name** field, type a name for the NAT entry.
6. From the **Interface** drop-down list, select a routed or hybrid interface.
7. From the **Type** drop-down list, select **Dynamic IP + Port** to display the dynamic IP and port NAT options.
8. Optionally, clear the **Enabled** check box, which means the entry does not affect traffic while disabled.

TIP! You can also enable or disable a NAT entry from the entry list by right-clicking the entry, selecting **State**, and the appropriate option.

9. Select either **CIDR** or **Range**:
 - If you selected **CIDR**, in the **Translated Source IP Address Block** field, either select an IP address from the drop-down list or type an IP address in CIDR notation. See [IP Address Conventions](#) on page 57.
 - If you selected **Range**, type a **First Translated Source IP Address** for the beginning of the range and a **Last Translated Source IP Address** for the end of the range.

IMPORTANT! On a 3D8250, you must have at least two IP addresses.

10. Click **OK**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Deleting NAT Entries

LICENSE: Control

SUPPORTED DEVICES: Series 3

The following procedure explains how to delete a NAT entry.

IMPORTANT! For managed devices running at least Version 5.2 of the Sourcefire 3D System, you configure NAT policies and apply them to the devices. See [Using NAT Policies](#) on page 383. This section refers to the NAT you must configure for pre-5.2 devices while configuring virtual routers.

To delete a NAT entry:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.

The Device Management page appears.
2. Next to the device where you want to delete the NAT entry, click the edit icon ().

The Interfaces tab for that device appears.
3. Click **NAT**.

The NAT tab appears.
4. Next to the NAT entry you want to delete, click the delete icon (). You can also right-click the entry and select **Delete**.

5. When prompted, confirm that you want to delete the entry.
The entry is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

CHAPTER 9

SETTING UP HYBRID INTERFACES

You can configure logical hybrid interfaces on managed devices that allow the Sourcefire 3D System to bridge traffic between virtual routers and virtual switches. If IP traffic received on interfaces in a virtual switch is addressed to the MAC address of an associated hybrid logical interface, the system handles it as Layer 3 traffic and either routes or responds to the traffic depending on the destination IP address. If the system receives any other traffic, it handles it as Layer 2 traffic and switches it appropriately. You cannot configure logical hybrid interfaces on a virtual managed device.

Note that hybrid interfaces that are not associated with both a virtual switch and a virtual router are not available for routing, and do not generate or respond to traffic.

For more information about setting up hybrid interfaces, see [Adding Logical Hybrid Interfaces](#) on page 352.

Adding Logical Hybrid Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

You must associate a logical hybrid interface with a virtual router and virtual switch to bridge traffic between Layer 2 and Layer 3. You can only associate a single hybrid interface with a virtual switch. However, you can associate multiple hybrid interfaces with a virtual router.

You can also configure Sourcefire Redundancy Protocol (SFRP) on a logical hybrid interface. See [Configuring SFRP](#) on page 304 for more information.

Note that disabling the **ICMP Enable Responses** option for hybrid interfaces does not prevent ICMP responses in all scenarios. You can add rules to an access control policy to drop packets where the destination IP is the hybrid interface's IP and the protocol is ICMP. For more information about creating access control rules, see [Understanding and Writing Access Control Rules](#) on page 472. If you have enabled the **Inspect Local Router Traffic** option on the managed device, it drops the packets before they reach the host, thereby preventing any response. For more information about inspecting local router traffic, see [Understanding Advanced Device Settings](#) on page 250.

WARNING! Changing the maximum transmission unit (MTU) interrupts routed or switched traffic on the device and packets are dropped. The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.

To edit an existing hybrid interface, click the edit icon (✎) next to the interface.

To add a logical hybrid interface:

ACCESS: Admin/Network Admin

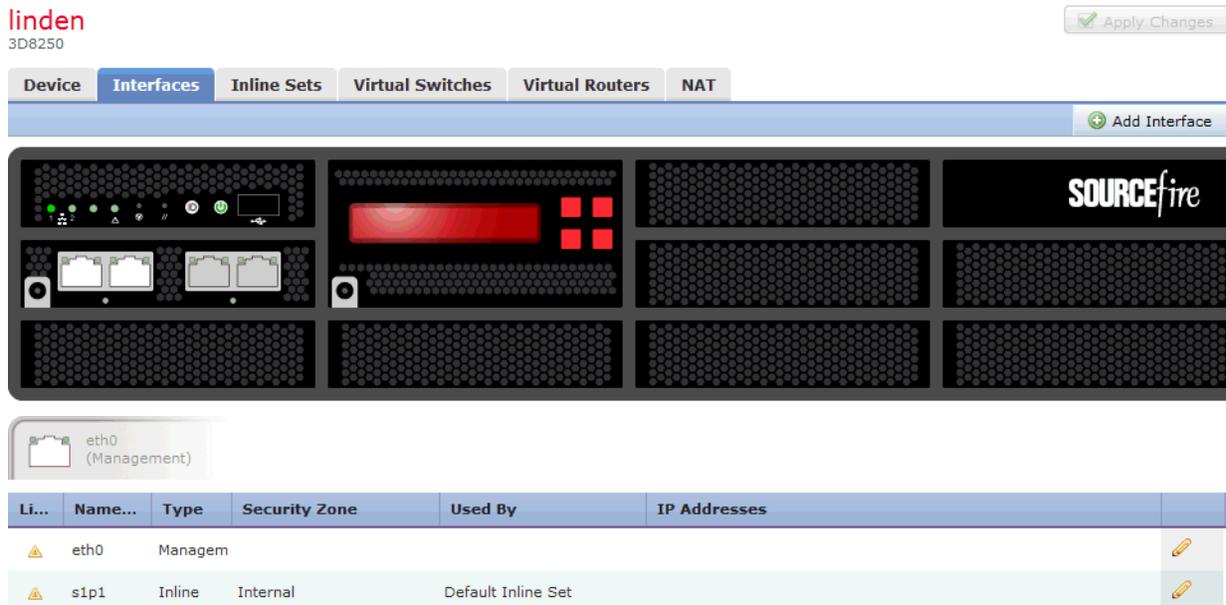
1. Select **Devices > Device Management**.

The Device Management page appears.

Name	License Type	Health Policy	System Policy	Access Control Pol...
4 📁 Ungrouped (1)				
linden 10.10.10.10 - 3D8250	Protect & Control w/URL Filtering	Blacklisted Power Supply	katsura system policy	My Access Control Policy

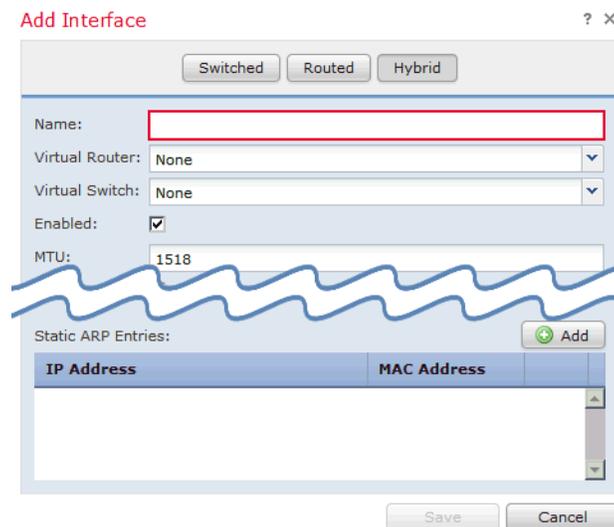
- Next to the device where you want to add the hybrid interface, click the edit icon (✎).

The Interfaces tab appears.



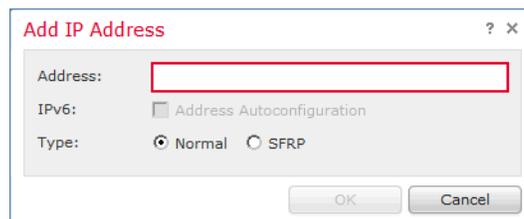
- Click **Add Interface**.

The Add Interface pop-up window appears.



- Click **Hybrid** to display the hybrid interface options.
- In the **Name** field, type a name for the interface. You can use alphanumeric characters and spaces.

6. From the **Virtual Router** drop-down list, select an existing virtual router, select **None**, or select **New** to add a new virtual router.
Note that if you add a new virtual router, you must configure it on the Device Management page (**Devices > Device Management > Virtual Routers**) after you finish setting up the hybrid interface. See [Adding Virtual Routers](#) on page 307.
7. From the **Virtual Switch** drop-down list, select an existing virtual switch, select **None**, or select **New** to add a new virtual switch.
Note that if you add a new virtual switch, you must configure it on the Device Management page (**Devices > Device Management > Virtual Switches**) after you finish setting up the hybrid interface. See [Adding Virtual Switches](#) on page 289.
8. Select the **Enabled** check box to allow the hybrid interface to handle traffic.
If you clear the check box, the interface becomes disabled and administratively taken down.
9. In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.
The range within which you can set the MTU can vary depending on the Sourcefire 3D System software version, device model, and interface type. See [Configuring the Interface MTU](#) on page 263 for more information.
10. Next to **ICMP**, select the **Enable Responses** check box to allow the interface to respond to ICMP traffic such as pings and traceroute.
11. Next to **IPv6 NDP**, select the **Enable Router Advertisement** check box to enable the interface to broadcast router advertisements.
You can only select this option if you added IPv6 addresses.
12. To add an IP address, click **Add**.
The Add IP Address pop-up window appears.

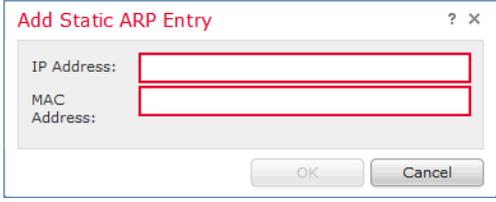


13. In the **Address** field, type the IP address and subnet mask.
14. Optionally if you have IPv6 addresses, next to the **IPv6** field, select the **Address Autoconfiguration** check box to set the IP address of the interface automatically.
15. For **Type**, select either Normal or SFRP.
For SFRP options, see [Configuring SFRP](#) on page 304 for more information.

16. Click **OK**.
The IP address is added.

TIP! To edit an IP address, click the edit icon (✎). To delete an IP address, click the delete icon (🗑).

17. To add a static ARP entry, click **Add**.
The Add Static ARP Entry pop-up window appears.



18. In the **IP Address** field, type an IP address for the static ARP entry.
19. In the **MAC Address** field, type a MAC address to associate with the IP address. Enter the address using the standard format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).
20. Click **OK**.
The static ARP entry is added.

TIP! To edit a static ARP entry, click the edit icon (✎). To delete a static ARP entry, click the delete icon (🗑).

21. Click **Save**.
The logical hybrid interface is added. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

Deleting Logical Hybrid Interfaces

LICENSE: Control

SUPPORTED DEVICES: Series 3

The following procedure explains how to delete a logical hybrid interface.

To delete a hybrid interface:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.

2. Next to the device where you want to delete the logical hybrid interface, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Next to the logical hybrid interface you want to delete, click the delete icon (🗑️).
4. When prompted, confirm that you want to delete the interface.
The interface is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

CHAPTER 10

USING GATEWAY VPN

A virtual private network (VPN) is a network connection that establishes a secure tunnel between endpoints via a public source, such as the Internet or other network. You can configure the Sourcefire 3D System to build secure VPN tunnels from the virtual routers of Sourcefire managed devices to remote devices or other third-party VPN endpoints. The system builds tunnels using the Internet Protocol Security (IPSec) protocol suite.

After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses and host names of the two gateways, the subnets behind them, and the shared secrets for the two gateways to authenticate to each other.

The VPN endpoints authenticate to each other with either the Internet Key Exchange (IKE) version 1 or version 2 protocol to create a security association for the tunnel. The system uses either the IPSec authentication header (AH) protocol or the IPSec encapsulating security payload (ESP) protocol to authenticate the data entering the tunnel. The ESP protocol encrypts the data as well as providing the same functionality as AH.

If you have access control policies in your deployment, the system does not send VPN traffic until it has passed through access control. In addition, the system does not send tunnel traffic to the public source when the tunnel is down.

To configure and apply VPN deployments, you must have a VPN license enabled on each of your target managed devices. Additionally, VPN features are only available on Series 3 devices.

See the following sections for more information on creating and managing VPN deployments:

- [Understanding IPSec](#) on page 359
- [Understanding VPN Deployments](#) on page 360
- [Managing VPN Deployments](#) on page 362

Understanding IPSec

The IPSec protocol suite defines how IP packets across a VPN tunnel are hashed, encrypted, and encapsulated in the ESP or AH security protocol. The Sourcefire 3D System uses the hash algorithm and encryption key of the Security Association (SA), which becomes established between the two gateways by the Internet Key Exchange (IKE) protocol.

Security associations (SA) establish shared security attributes between two devices and allow VPN endpoints to support secure communication. An SA allows two VPN endpoints to handle the parameters for how the VPN tunnel is secured between them.

The system uses the Internet Security Association and Key Management Protocol (ISAKMP) during the initial phase of negotiating the IPSec connection to establish the VPN between endpoints and the authenticated key exchange. The IKE protocol resides within ISAKMP. See [Understanding IKE](#) on page 359 for more information about the IKE protocol.

The AH security protocol provides protection for packet headers and data, but it cannot encrypt them. ESP provides encryption and protection for packets, but it cannot secure the outermost IP header. In many cases, this protection is not required, and most VPN deployments use ESP more frequently than AH because of its encryption capabilities. Since VPN only operates in tunnel mode, the system encrypts and authenticates the entire packet from Layer 3 and up in the ESP protocol. ESP in tunnel mode encrypts the data as well as providing the latter's encryption capabilities.

Understanding IKE

The Sourcefire 3D System uses the IKE protocol to mutually authenticate the two gateways against each other as well as to negotiate the SA for the tunnel. The process consists of two phases.

IKE phase 1 establishes a secure authenticated communication channel by using the Diffie-Hellman key exchange to generate a pre-shared key to encrypt further IKE communications. This negotiation results in a bidirectional ISAKMP security association. The system allows you to perform the authentication using a pre-shared key. Phase 1 operates in main mode, which seeks to protect all data during the negotiation, while also protecting the identity of the peers.

During IKE phase 2, the IKE peers use the secure channel established in phase 1 to negotiate security associations on behalf of IPSec. The negotiation results in a minimum of two unidirectional security associations, one inbound and one outbound.

Understanding VPN Deployments

A VPN deployment specifies the endpoints and networks that are included in a VPN and how they connect to each other. After you configure a VPN deployment, you can then apply it to your managed devices, as well as remote devices or other third-party VPN endpoints.

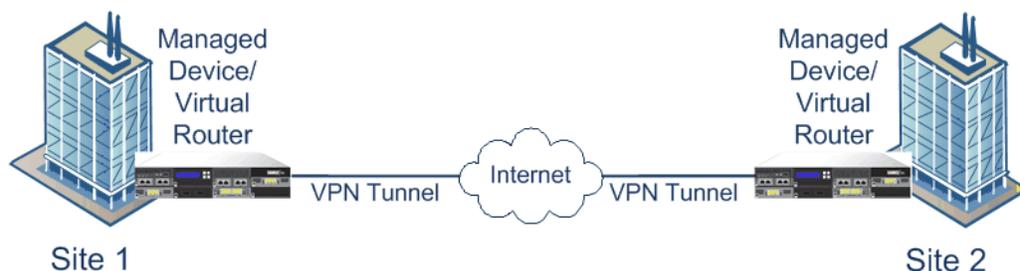
The system supports three types of VPN deployments: point-to-point, star, and mesh. See the following sections for more information about these VPN deployments:

- [Understanding Point-to-Point VPN Deployments](#) on page 360
- [Understanding Star VPN Deployments](#) on page 360
- [Understanding Mesh VPN Deployments](#) on page 361

Understanding Point-to-Point VPN Deployments

In a point-to-point VPN deployment, two endpoints communicate directly with each other. You configure the two endpoints as peer devices, and either device can initiate the secured connection. At least one of the devices in this configuration must be a VPN-enabled managed device.

The following diagram displays a typical point-to-point VPN deployment.



See [Configuring Point-to-Point VPN Deployments](#) on page 364 for more information.

Understanding Star VPN Deployments

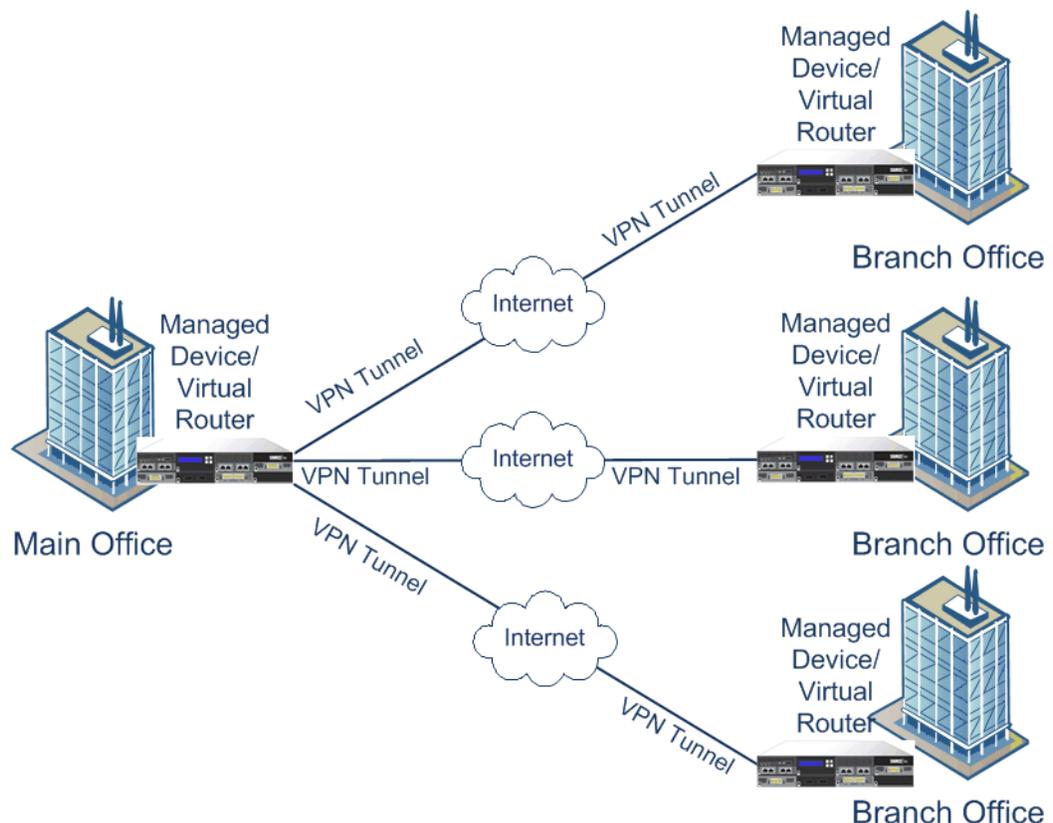
In a star VPN deployment, a central endpoint (hub node) establishes a secure connection with multiple remote endpoints (leaf nodes). Each connection between the hub node and an individual leaf node is a separate VPN tunnel. The

hosts behind any of the leaf nodes can communicate with each other through the hub node.

Star deployments commonly represent a VPN that connects an organization's main and branch office locations using secure connections over the Internet or other third-party network. Star VPN deployments provide all employees with controlled access to the organization's network.

In a typical star deployment, the hub node is a VPN-enabled managed device, located at the main office. Leaf nodes are located at branch offices and initiate most of the traffic. Star deployments only support IKE version 2.

The following diagram displays a typical star VPN deployment.



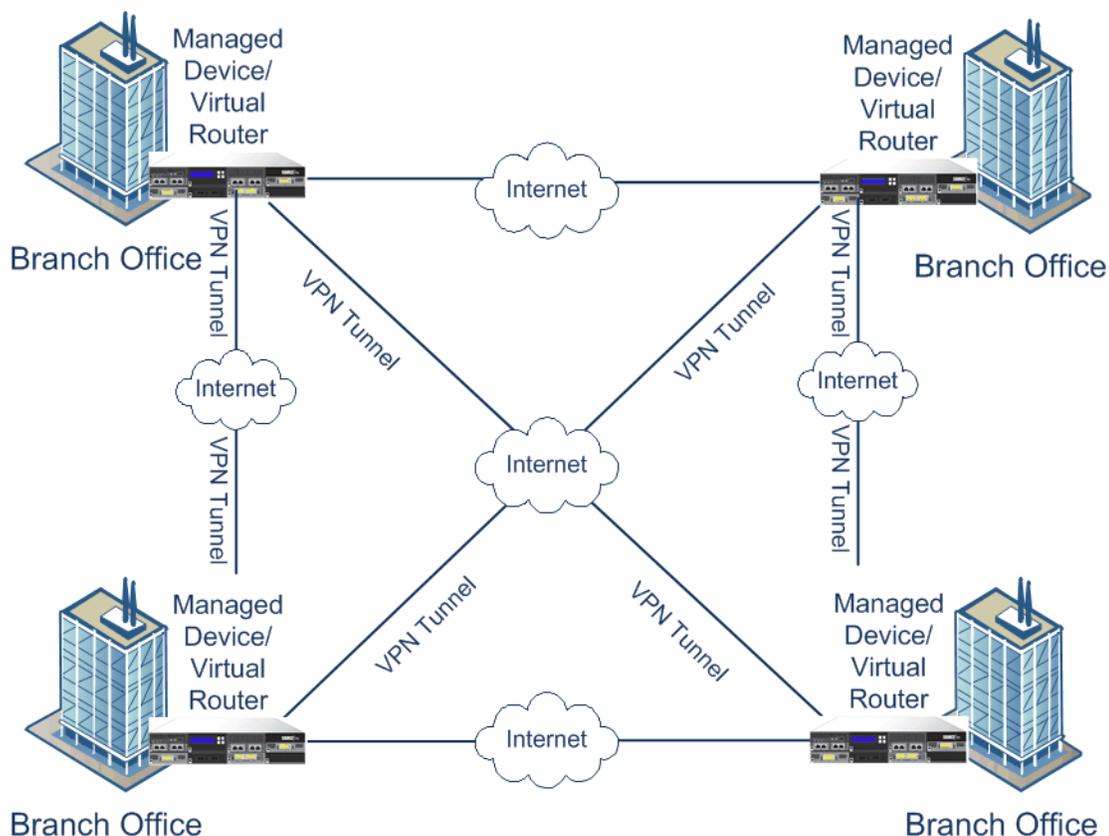
See [Configuring Star VPN Deployments](#) on page 368 for more information.

Understanding Mesh VPN Deployments

In a mesh VPN deployment, all endpoints can communicate with every other endpoint by means of an individual VPN tunnel. The mesh deployment offers redundancy so that when one endpoint fails, the remaining endpoints can still communicate with each other. This type of deployment commonly represents a VPN that connects a group of decentralized branch office locations. The number

of VPN-enabled managed devices you deploy in this configuration depends on the level of redundancy you require.

The following diagram displays a typical mesh VPN deployment.



See [Configuring Mesh VPN Deployments](#) on page 371 for more information.

Managing VPN Deployments

LICENSE: VPN

SUPPORTED DEVICES: Series 3

On the VPN page (**Devices > VPN**) you can view all of your current VPN deployments by name and the endpoints contained in the deployment. Options on this page allow you to view the status of a VPN deployment, create a new deployment, apply a deployment, and edit or delete a deployment.

WARNING! If you select the default access control policy when registering a device to your Defense Center, the default access control rule blocks all traffic. If you configure a VPN deployment on the device, the deployment fails.

The following table describes the actions you can take to manage your deployments on the VPN page.

VPN Deployment Management Actions

To...	You CAN...
create a new VPN deployment	click Add . See Configuring VPN Deployments on page 363 for more information.
modify the settings in an existing VPN deployment	click the edit icon (✎). See Configuring VPN Deployments on page 363 for more information.
view the status of an existing VPN deployment	click the status icon. See Viewing VPN Deployment Status on page 377 for more information.
apply a VPN deployment to all devices targeted in the deployment	click the apply icon (✔). See Applying a VPN Deployment on page 377 for more information.
delete a VPN deployment	click the delete icon (🗑), then click Yes , or click No if you decide not to delete the deployment.

Configuring VPN Deployments

LICENSE: VPN

SUPPORTED DEVICES: Series 3

When you create a new VPN deployment you must, at minimum, give it a unique name, specify a deployment type, and designate a pre-shared key. You can select from three types of deployment, each containing a group of VPN tunnels:

- Point-to-point (PTP) deployments establish a VPN tunnel between two endpoints.
- Star deployments establish a group of VPN tunnels connecting a hub endpoint to a group of leaf endpoints.
- Mesh deployments establish a group of VPN tunnels among a set of endpoints.

You must define a pre-shared key for VPN authentication. You can specify a default key to use in all of the VPN connections you generate in a deployment. For point-to-point deployments, you can specify a pre-shared key for each endpoint pair.

See the following sections for more information on creating each type of VPN deployment:

- [Configuring Point-to-Point VPN Deployments](#) on page 364
- [Configuring Star VPN Deployments](#) on page 368
- [Configuring Mesh VPN Deployments](#) on page 371

Configuring Point-to-Point VPN Deployments

LICENSE: VPN

SUPPORTED DEVICES: Series 3

When configuring a point-to-point VPN deployment, you define a group of endpoint pairs and then create a VPN between the two nodes in each pair. For more information, see [Understanding Point-to-Point VPN Deployments](#) on page 360.

The following list describes the options you can specify in your deployment.

Name

Give the deployment a unique name.

Type

Click **PTP** to specify that you are configuring a point-to-point deployment.

Pre-shared Key

Define a unique pre-shared key for authentication. The system uses this key for all the VPNs in your deployment, unless you specify a pre-shared key for each endpoint pair.

Device

You can select a managed device, including a device stack or cluster, as an endpoint for your deployment. For endpoints not managed by your Defense Center, select **Other** and then specify an IP address for the endpoint.

Virtual Router

If you selected a managed device as your endpoint, select a virtual router that is currently applied to the selected device. You cannot select the same virtual router for more than one endpoint.

Interface

If you selected a managed device as your endpoint, select a routed interface that is assigned to the selected virtual router.

IP Address

- If you selected a managed device as an endpoint, select an IP address that is assigned to the selected routed interface.
- If the managed device is a device cluster, you can only select from a list SFRP IP addresses.
- If you selected an endpoint **not** managed by the Defense Center, specify an IP address for the endpoint.

Protected Networks

Specify the networks in your deployment that are encrypted. Enter a subnet with CIDR block for each network. IKE version 1 only supports a single protected network.

Internal IP

Select the check box if the endpoint resides behind a firewall with network address translation.

Public IP

If you selected **Internal IP**, specify a public IP address for the firewall. If the endpoint is a responder, you must specify this value.

Public IKE Port

If you selected **Internal IP**, specify a numerical value for the UDP port on the firewall that is being port-forwarded to the internal endpoint. If the endpoint is a responder and the port on the firewall being forwarded is not 500 or 4500, you must specify this value.

Use Deployment Key

Select the check box to use the pre-shared key defined for the deployment. Clear the check box to specify a pre-shared key for VPN authentication for this endpoint pair.

Pre-shared Key

If you cleared the **Use Deployment Key** check box, specify a pre-shared key in this field.

TIP! To edit an existing point-to-point deployment, click the edit icon (✎) next to the deployment. You cannot edit the deployment type after you initially save the deployment.

To configure a point-to-point VPN deployment:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.
The VPN page appears.

2. Click **Add**.

The Create New VPN Deployment pop-up window appears.

The screenshot shows a dialog box titled "Create New VPN Deployment" with a close button (X) and a help button (?). It has two tabs: "General" (active) and "Advanced".

Fields and controls include:

- Name:** A text input field.
- Type:** Three buttons: "PTP" (selected), "Star", and "Mesh".
- Pre-shared Key:** A text input field.
- Node Pairs:** A section with a table for defining node pairs. The table has two columns: "Node A" and "Node B". A green plus icon is in the top right corner of this section.

At the bottom, there is a red error message: "Please enter a deployment name." and two buttons: "Save" and "Cancel".

3. Give the deployment a unique **Name**.
You can use all printable characters, including spaces and special characters.
4. Ensure that **PTP** is selected as the **Type**.
5. Give the deployment a unique **Pre-shared Key**.

- Next to **Node Pairs**, click the add icon (+).
- The Add New Endpoint Pair pop-up window appears.

The screenshot shows the 'Add New Endpoint Pair' dialog box. It is divided into two main sections for 'Node A' and 'Node B'. Each section contains the following fields: 'Device', 'Virtual Router', 'Interface', and 'IP Address', all with dropdown menus. Below these is a 'Protected Networks' section with a table header 'Network' and an empty table body. At the bottom of each section are 'Advanced' settings: 'Internal IP' (checkbox), 'Public IP' (text field), and 'public IKE Port' (text field with '500'). A 'Use Deployment Key' checkbox is checked at the bottom left. A red error message 'Incomplete endpoint.' is visible at the bottom left, and 'Save' and 'Cancel' buttons are at the bottom right.

- Configure the VPN deployment, as described earlier in this section.
- Under **Node A**, next to **Protected Networks**, click the add icon (+).
The Add Network pop-up window appears.
- Type a CIDR block for the protected network.
- Click **OK**.
The protected network is added.
- Repeat step 8 through step 10 for **Node B**.
- Click **Save**.
The endpoint pair is added to your deployment and the Create New VPN Deployment pop-up window appears again.
- Click **Save** to finish configuring your deployment and the VPN page appears again.
Note that you must apply the deployment for it to take effect; see [Applying a VPN Deployment](#) on page 377.

Configuring Star VPN Deployments

LICENSE: VPN

SUPPORTED DEVICES: Series 3

When configuring a star VPN deployment, you define a single hub node endpoint and a group of leaf node endpoints. You must define the hub node endpoint and at least one leaf node endpoint to configure the deployment. For more information, see [Understanding Star VPN Deployments](#) on page 360.

The following list describes the options you can specify in your deployment.

Name

Give the deployment a unique name.

Type

Click **Star** to specify that you are configuring a star deployment.

Pre-shared Key

Define a unique pre-shared key for authentication.

Device

You can select a managed device, including a device stack or cluster, as an endpoint for your deployment. For endpoints not managed by your Defense Center, select **Other** and then specify an IP address for the endpoint.

Virtual Router

If you selected a managed device as your endpoint, select a virtual router that is currently applied to the selected device. You cannot select the same virtual router for more than one endpoint.

Interface

If you selected a managed device as your endpoint, select a routed interface that is assigned to the selected virtual router.

IP Address

- If you selected a managed device as an endpoint, select an IP address that is assigned to the selected routed interface.
- If the managed device is a device cluster, you can only select from a list SFRP IP addresses.
- If you selected an endpoint **not** managed by the Defense Center, specify an IP address for the endpoint.

Protected Networks

Specify the networks in your deployment that are encrypted. Enter a subnet with CIDR block for each network.

Internal IP

Select the check box if the endpoint resides behind a firewall with network address translation.

Public IP

If you selected **Internal IP**, specify a public IP address for the firewall. If the endpoint is a responder, you must specify this value.

Public IKE Port

If you selected **Internal IP**, specify a numerical value for the UDP port on the firewall that is being port-forwarded to the internal endpoint. If the endpoint is a responder and the port on the firewall being forwarded is not 500 or 4500, you must specify this value.

TIP! To edit an existing star deployment, click the edit icon (✎) next to the deployment. You cannot edit the deployment type after you initially save the deployment. To change the deployment type, you must delete the deployment and create a new one.

To configure a star deployment:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.
The VPN page appears
2. Click **Add**.
The Create New VPN Deployment pop-up window appears.
3. Give the deployment a unique **Name**.
You can use all printable characters, including spaces and special characters.
4. Click **Star** to specify the **Type**.
5. Give the deployment a unique **Pre-shared Key**.

- Next to **Hub Node**, click the add icon (+).
- The Add Hub Node pop-up window appears.

The screenshot shows the 'Add Hub Node' dialog box. It includes fields for 'Device', 'Virtual Router', 'Interface', and 'IP Address', each with a dropdown arrow. Below these is a 'Protected Networks' section with a table header 'Network' and a plus icon. The 'Advanced' section is collapsed. The 'Internal IP' checkbox is unchecked. The 'Public IP' field is empty. The 'public IKE Port' field contains the value '500'. At the bottom, there is a red error message 'Incomplete endpoint.' and 'Save' and 'Cancel' buttons.

- Configure the VPN deployment, as described earlier in this section.
 - Next to **Protected Networks**, click the add icon (+).
- The Add Network pop-up window appears.
- Type an IP address for the protected network.
 - Click **OK**.
- The protected network is added.
- Click **Save**.
- The hub node is added to your deployment and the Create New VPN Deployment pop-up window appears again.

- Next to **Leaf Nodes**, click the add icon (+).
- The Add Leaf Node pop-up window appears.

The screenshot shows the 'Add Leaf Node' dialog box. It features a title bar with a question mark and a close button. The main area is divided into several sections: 'Device:', 'Virtual Router:', 'Interface:', and 'IP Address:', each with a dropdown menu. Below these is a 'Protected Networks' section with a table header 'Network' and a green plus icon. The 'Advanced:' section includes 'Internal IP:' with a checkbox, 'Public IP:' with a text field, and 'public IKE Port:' with a text field containing '500'. At the bottom, there is a red error message 'Incomplete endpoint.', a 'Save' button, and a 'Cancel' button.

- Repeat step 7 through step 10 to complete the leaf node, which has the same options as the hub node.
- Click **Save**.
The leaf node is added to your deployment and the Create New VPN Deployment pop-up window appears again.
- Click **Save** to finish configuring your deployment and the VPN page appears again.
Note that you must apply the deployment for it to take effect; see [Applying a VPN Deployment](#) on page 377.

Configuring Mesh VPN Deployments

LICENSE: VPN

SUPPORTED DEVICES: Series 3

When configuring a mesh VPN deployment, you define a group of VPNs to link any two points for a given set of endpoints. For more information, see [Configuring Mesh VPN Deployments](#) on page 371.

The following list describes the options you can specify in your deployment.

Name

Give the deployment a unique name.

Type

Click **Mesh** to specify that you are configuring a mesh deployment.

Pre-shared Key

Define a unique pre-shared key for authentication.

Device

You can select a managed device, including a device stack or cluster, as an endpoint for your deployment. For endpoints not managed by your Defense Center, select **Other** and then specify an IP address for the endpoint.

Virtual Router

If you selected a managed device as your endpoint, select a virtual router that is currently applied to the selected device. You cannot select the same virtual router for more than one endpoint.

Interface

If you selected a managed device as your endpoint, select a routed interface that is assigned to the selected virtual router.

IP Address

- If you selected a managed device as an endpoint, select an IP address that is assigned to the selected routed interface.
- If the managed device is a device cluster, you can only select from a list SFRP IP addresses.
- If you selected an endpoint **not** managed by the Defense Center, specify an IP address for the endpoint.

Protected Networks

Specify the networks in your deployment that are encrypted. Enter a subnet with CIDR block for each network. IKE version 1 only supports a single protected network.

Internal IP

Select the check box if the endpoint resides behind a firewall with network address translation.

Public IP

If you selected **Internal IP**, specify a public IP address for the firewall. If the endpoint is a responder, you must specify this value.

Public IKE Port

If you selected **Internal IP**, specify a numerical value for the UDP port on the firewall that is being port-forwarded to the internal endpoint. If the endpoint is a responder and the port on the firewall being forwarded is not 500 or 4500, you must specify this value.

TIP! To edit an existing mesh deployment, click the edit icon (✎) next to the deployment. You cannot edit the deployment type after you initially save the deployment. To change the deployment type, you must delete the deployment and create a new one.

To configure a mesh VPN deployment:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.

The VPN page appears

2. Click **Add**.

The Create New VPN Deployment pop-up window appears.

3. Give the deployment a unique **Name**.

You can use all printable characters, including spaces and special characters.

4. Click **Mesh** to specify the **Type**.

5. Give the deployment a unique **Pre-shared Key**.

- Next to **Nodes**, click the add icon (+).
The Add Endpoint pop-up window appears.

The screenshot shows the 'Add Endpoint' dialog box. It includes fields for Device, Virtual Router, Interface, IP Address, and IKE port (set to 500). There is a section for Protected Networks with a table header 'Network' and a plus icon. Below that is an Advanced section with Internal IP (checkbox), Public IP, and Public IKE Port (set to 500). At the bottom, there is a red error message 'Incomplete endpoint.' and 'Save' and 'Cancel' buttons.

- Configure the VPN deployment, as described earlier in this section.
- Next to **Protected Networks**, click the add icon (+).
The Add Network pop-up window appears.
- Type a CIDR block for the protected network.
- Click **OK**.
The protected network is added.
- Click **Save**.
The endpoint is added to your deployment and the Create New VPN Deployment pop-up window appears again.
- Repeat step 6 through step 11 to add more endpoints.
- Click **Save** to complete your deployment and the VPN page appears again.
Note that you must apply the deployment for it to take effect; see [Applying a VPN Deployment](#) on page 377.

Configuring Advanced VPN Deployment Settings

LICENSE: VPN

SUPPORTED DEVICES: Series 3

VPN deployments contain some common settings that can be shared among the VPNs in a deployment. Each VPN can use the default settings or you can override

the default settings. Advanced settings typically require little or no modification and are not common to every deployment.

The following list describes the advanced options you can specify in your deployment.

Other Algorithm Allowed

Select the check box to enable auto negotiation to an algorithm not listed in the Algorithm list, but proposed by the remote peer when the remote peer is a third party device.

Algorithm

Specify the phase one and phase two algorithm proposals to secure data in your deployment. Select **Cipher, Hash**, and Diffie-Hellman (**DH**) group authentication messages for both phases.

IKE Life Time

Specify a numerical value and select a time unit for the maximum IKE SA renegotiation interval. You can specify a minimum of 15 minutes and a maximum of 30 days.

IKE v2

Select the check box to specify that the system uses IKE version 2. This version supports the star deployment and multiple protected networks.

Life Time

Specify a numerical value and select a time unit for the maximum SA renegotiation interval. You can specify a minimum of 5 minutes and a maximum of 24 hours.

Life Packets

Specify the number of packets that can be transmitted over an IPsec SA before it expires. You can use any integer between 0 and 18446744073709551615.

Life Bytes

Specify the number of bytes that can be transmitted over an IPsec SA before it expires. You can use any integer between 0 and 18446744073709551615.

AH

Select the check box to specify that the system uses the authentication header security protocol for the data to be protected. Clear the check box to use encryption service payload (ESP) protocol. See [Understanding IPSec](#) on page 359 for guidance on when to use each protocol.

To configure advanced VPN deployment settings:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.
The VPN page appears.
2. Click **Add**.
The Create New VPN Deployment pop-up window appears.
3. Click the **Advanced** tab.

The screenshot shows the 'Create New VPN Deployment' dialog box with the 'Advanced' tab selected. It is divided into two main sections: IKE Algorithm Proposal and IPSec Algorithm Proposal. Each section has a table of algorithms and a set of settings below it.

IKE Algorithm Proposal:

Cipher	Hash	DH	
aes128	sha1	modp2048	
3des	sha1	modp1536	

IPSec Algorithm Proposal:

Cipher	Hash	DH	
aes128	sha1		
3des	sha1		

IKE Settings:

IKE Life Time: 3 hours

IKE v2:

IPSec Settings:

Life Time: 1 hours

Life Packets: 0

Life Bytes: 0

AH:

Please enter a deployment name.

Save Cancel

4. Configure the advanced settings, as described earlier in this section.
5. Next to **Algorithms**, click the add icon ().
The Add IKE Algorithm Proposal pop-up window appears.
6. Select **Cipher**, **Hash**, and Diffie-Hellman (**DH**) group authentication messages for both phases.
7. Click **OK**.
The IKE algorithm proposal is added.

8. Click **Save**.

Your changes are saved and the VPN page appears.

Note that you must apply the deployment for it to take effect; see [Applying a VPN Deployment](#) on page 377.

Applying a VPN Deployment

LICENSE: VPN

SUPPORTED DEVICES: Series 3

After configuring or making any changes to a VPN deployment, you must apply the deployment to one or more devices to implement the settings you designated for the deployment.

To apply a VPN deployment:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.

The VPN page appears.

2. Click the apply icon (✓) next to the VPN deployment that you want to apply.

3. When prompted, click **Yes**.

The VPN deployment is applied.

TIP! Optionally, from the Apply VPN deployment dialog box, click **View Changes**. The VPN Comparison View page appears in a new browser window. For more information, see [Using the VPN Deployment Comparison View](#) on page 381.

4. Click **OK**.

You are returned to the VPN page.

Viewing VPN Deployment Status

LICENSE: VPN

SUPPORTED DEVICES: Series 3

After you configure a VPN deployment, you can view the status of your configured VPN tunnels. The VPN page displays a status icon for each applied VPN deployment:

- The (✓) icon designates that all VPN endpoints are up.
- The (ⓘ) icon designates that all VPN endpoints are down.
- The (⚠) icon designates that some endpoints are up, while others are down.

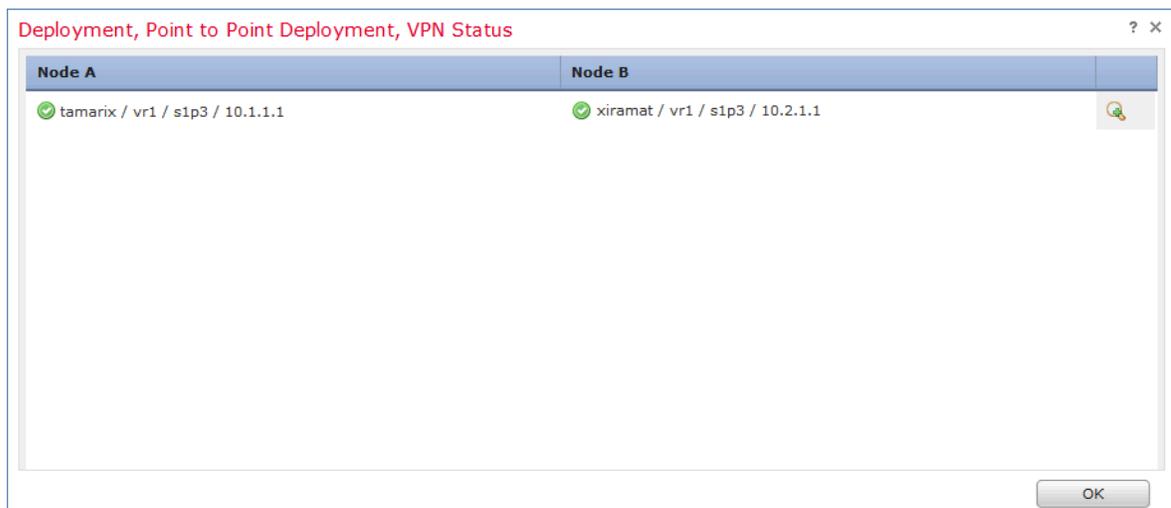
You can click a status icon to view the deployment status along with basic information about the endpoints in the deployment, such as endpoint name and IP address. The VPN status updates every minute or when a status change occurs, such as an endpoint going down or coming up.

To view VPN status:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.
The VPN page appears.
2. Click the VPN status icon next to the deployment where you want to view the status.

The VPN Status pop-up window appears.



3. Click **OK** to return to the VPN page.

Viewing VPN Statistics and Logs

LICENSE: VPN

SUPPORTED DEVICES: Series 3

After you configure a VPN deployment, you can view statistics about the data traversing your configured VPN tunnels. In addition, you can view the latest VPN system and IKE logs for each endpoint.

The system displays the following statistics.

Endpoint

The device path to the routed interface and IP address designated as the VPN endpoint.

Status

Whether the VPN connection is up or down.

Protocol

The protocol used for encryption, either ESP or AH.

Packets Received

The number of packets per interface the VPN tunnel receives during an IPsec SA negotiation.

Packets Forwarded

The number of packets per interface the VPN tunnel transmits during an IPsec SA negotiation.

Bytes Received

The number of bytes per interface the VPN tunnel receives during an IPsec SA negotiation.

Bytes Forwarded

The number of bytes per interface the VPN tunnel transmits during an IPsec SA negotiation.

Time Created

The date and time the VPN connection was created.

Time Last Used

The last time a user initiated a VPN connection.

NAT Traversal

If Yes is displayed, at least one of the VPN endpoints resides behind a device with network address translation.

IKE State

The state of the IKE SA: connecting, established, deleting, or destroying.

IKE Event

The IKE SA event: reauthentication or rekeying.

IKE Event Time

The time in seconds the next event should occur.

IKE Algorithm

The IKE algorithm being used by the VPN deployment.

IPSec State

The state of the IPSec SA: installing, installed, updating, rekeying, deleting, and destroying.

IPSec Event

Notification of when the IPSec SA event is rekeying.

IPSec Event Time

The time in seconds until the next event should occur.

IPSec Algorithm

IPSec algorithm being used by the VPN deployment.

To view VPN statistics:

ACCESS: Admin/Network Admin

1. Select **Devices > VPN**.

The VPN page appears.

2. Click the VPN status icon next to the deployment where you want to view the VPN statistics.

The VPN Status pop-up window appears.

- Click the view statistics icon (📊).
The VPN Statistics pop-up window appears.

	Node A	Node B
Endpoint:	tamarix/vr2/s3p1/10.1.1.3	xiramat/vr1/s3p1/10.1.1.1
Status:	UP	UP
Protocol:	ESP	ESP
Packets Received:	0	0
Packets Forwarded:	54567610	54431977
Bytes Received:	0	0
Bytes Forwarded:	2147483647	2147483647
Time Created:	Tue May 21 03:05:12 2013	Tue May 21 03:04:59 2013
Time Last Used:	Tue May 21 12:50:26 2013	Tue May 21 12:50:26 2013
NAT Traversal:	No	No
IKE Data		
State:	ESTABLISHED	ESTABLISHED
Event:	reauthentication	reauthentication
Event Time:	720	1020
Algorithm:	AES_CBC_128/HMAC_SHA1_96 /PRF_HMAC_SHA1/MODP_1024	AES_CBC_128/HMAC_SHA1_96 /PRF_HMAC_SHA1/MODP_1024
IPSec Data		
State:	INSTALLED	INSTALLED
Event:	Rekeying	Rekeying
Event Time:	2040	2160
Algorithm:	AES_CBC_128/HMAC_SHA1_96	AES_CBC_128/HMAC_SHA1_96

[View Recent Log](#)
[View Recent Log](#)

- Optionally, click **Refresh** to update the VPN statistics.
- Optionally, click **View Recent Log** to view the latest data log for each endpoint.
To view the log for clustered devices and stacked devices, you can select the link for either the active/primary or backup/secondary device.

Using the VPN Deployment Comparison View

LICENSE: VPN

SUPPORTED DEVICES: Series 3

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

The comparison view displays both deployments in a side-by-side format, with each deployment 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 deployment name.

Differences between the two deployments are highlighted:

- Blue indicates that the highlighted setting is different in the two deployments, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one deployment but not the other.

You can perform any of the actions in the following table.

VPN Deployment 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 deployment comparison report	click Comparison Report . The deployment comparison report creates a PDF document that lists only the differences between the two policies.

CHAPTER 11

USING NAT POLICIES

A network address translation (NAT) policy determines how the system achieves routing with network address translation. You can configure one or more NAT policies, which you can then apply to one or more managed devices. Each device can have one currently applied policy.

You add NAT rules to a policy to control how the system handles network address translations. Each rule contains a set of conditions that identify the specific traffic you want to translate. You can create the following types of rules:

- static, which provide one-to-one translations on destination networks and optionally port and protocol
- dynamic IP, which translate many-to-many source networks, but maintain port and protocol
- dynamic IP and port, which translate many-to-one or many-to-many source networks and port and protocol

The system matches traffic to static translations before dynamic translations are inspected. The system then matches traffic to dynamic NAT rules in order; the first-matched rules handle the traffic. See [Organizing Rules in a NAT Policy](#) on page 389 for more information.

If you have access control policies in your deployment, the system does not translate traffic until it has passed through access control.

To configure and apply NAT policies on your appliances, you must have a Control license enabled on each of your target managed devices. Additionally, you can

only apply NAT policies to Series 3 devices with configured virtual routers or hybrid interfaces.

IMPORTANT! Managed devices must be running at least Version 5.2 of the Sourcefire 3D System before you can apply NAT policies to them. You must configure NAT for pre-5.2 devices from the NAT tab on the Device Management page. See [Configuring NAT on a Virtual Router](#) on page 341.

After you have configured and deployed NAT policies, you can use the command line interface (CLI) for managed device targets to troubleshoot the deployment. The CLI displays three types of NAT information: configuration, rule definitions, and active translations. See [Command Line Reference](#) on page 2242 for more information.

See the following sections for more information on creating and managing NAT policies:

- [Planning and Implementing a NAT Policy](#) on page 384
- [Configuring NAT Policies](#) on page 385
- [Organizing Rules in a NAT Policy](#) on page 389
- [Managing NAT Policies](#) on page 391
- [Creating and Editing NAT Rules](#) on page 404
- [Understanding NAT Rule Types](#) on page 406
- [Understanding NAT Rule Conditions and Condition Mechanics](#) on page 409
- [Working with Different Types of Conditions in NAT Rules](#) on page 415

Planning and Implementing a NAT Policy

LICENSE: Any

You can configure NAT policies in different ways to manage specific network needs. This section provides information for some of the ways you can deploy NAT policies.

WARNING! In clustered configurations, only select an individual peer interface for a static NAT rule on a clustered device if all networks affected by the NAT translations are private. Do **not** use this configuration for static NAT rules affecting traffic between public and private networks.

You can configure NAT to expose an internal server to an external network. In this configuration, you define a static translation from an external IP address to an internal IP address so the system can access an internal server from outside the network. Traffic sent to the server targets the external IP address or IP address and port, and is translated into the internal IP address or IP address and port. Return traffic from the server is translated back to the external address.

You can configure NAT to allow an internal host or server to connect to an external application. In this configuration, you define a static translation from an internal address to an external address. This definition allows the internal host or server to initiate a connection to an external application that is expecting the internal host or server to have a specific IP address and port. Therefore, the system cannot dynamically allocate the address of the internal host or server.

You can configure NAT to hide private network addresses from an external network by using a block of IP addresses. This becomes useful if you want to obscure your internal network addresses and have sufficient external IP addresses to satisfy your internal network needs. In this configuration, you create a dynamic translation that automatically converts the source IP address of any outgoing traffic to an unused IP address from your externally facing IP addresses.

You can configure NAT to hide private network addresses from an external network using a limited block of IP addresses and port translation. This becomes useful if you want to obscure your internal network addresses, but have an insufficient number of external IP addresses to satisfy your internal network needs. In this configuration, you create a dynamic translation that automatically converts the source IP address and port of outgoing traffic to an unused IP address and port from your externally facing IP addresses.

Configuring NAT Policies

LICENSE: Control

SUPPORTED DEVICES: Series 3

To configure a NAT policy, you must give the policy a unique name and identify the devices, or *targets*, where you want to apply the policy. You can also add, edit, delete, enable, and disable NAT rules. After you create or modify a NAT policy, you can apply the policy to all or some targeted devices.

You can apply NAT policies to a device cluster, including clustered stacks, as you would a standalone device. However, you can define static NAT rules for interfaces on individual clustered devices or the entire cluster and use the interfaces in source zones. For dynamic rules, you can use only the interfaces on the entire cluster in source or destination zones.

WARNING! In clustered configurations, only select an individual peer interface for a static NAT rule on a clustered device if all networks affected by the NAT translations are private. Do **not** use this configuration for static NAT rules affecting traffic between public and private networks.

If you configure dynamic NAT on a device cluster without HA link interfaces established, both clustered devices independently allocate dynamic NAT entries, and the system cannot synchronize the entries between devices. See [Configuring HA Link Interfaces](#) on page 262 for more information.

You can apply NAT policies to a device stack as you would a standalone device. If you establish a device stack from devices that were included in a NAT policy and had rules associated with interfaces from the secondary device that was a member of the stack, the interfaces from the secondary device remain in the NAT policy. You can save and apply policies with the interfaces, but the rules do not provide any translation. See [Managing Stacked Devices](#) on page 237 for more information.

The following table summarizes the configuration actions you can take on the NAT policy Edit page.

NAT 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.
manage policy targets	find more information at Managing NAT Policy Targets on page 387.
save your policy changes	click Save .
save and apply your policy	click Save and Apply . See Applying a NAT Policy on page 401 for more information.
cancel your policy changes	click Cancel , then, if you have made changes, click OK .
add a rule to a policy	click Add Rule . See Creating and Editing NAT Rules on page 404 for more information. TIP! You can also right-click an existing rule and select Insert new rule .
edit an existing rule	click the edit icon () next to the rule. See Creating and Editing NAT Rules on page 404 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! To delete one or more selected rules, you can right-click a blank area in the row for a selected rule, select Delete , then click OK .
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 Source Network page for the selected rule. See Working with Different Types of Conditions in NAT Rules on page 415 for more information.

Managing NAT Policy Targets

LICENSE: Control

SUPPORTED DEVICES: Series 3

Before you can apply a NAT policy, you must identify the managed devices, including device stacks, clusters, or groups, where you want to apply the policy. You can identify the managed devices you want to target with your policy while creating or editing a policy. You can search a list of available devices, stacks, and clusters, and add them to a list of selected devices. You can also drag and drop selected devices, or add devices using the button between the two lists.

Note that you cannot target stacked devices running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). See [Managing Stacked Devices](#) on page 237 for more information.

The following table summarizes the actions you can take when managing targeted devices.

Targeted Device Management Actions

To...	You CAN...
search a list of available devices, stacks, and clusters	click inside the search field, then type a search string. The list of devices updates as you type to display matching device names.
clear a search for available devices	click the clear icon (✕) in the search field.
select available devices, stacks, or clusters to add to the list of selected targets	click the device name; use the Ctrl and Shift keys to select multiple devices. TIP! You can also right-click an available device, then click Select All .
add selected devices, stacks, or clusters	click Add to Policy . TIP! You can also drag and drop into the list of selected devices.
delete a single device, stack, or cluster from the Selected Devices list	click the delete icon (🗑) next to the device. TIP! You can also right-click the device and select Delete .
delete multiple devices from the Selected Devices list	use the Ctrl and Shift keys to select multiple devices, right-click to highlight the row for a selected device, then click Delete Selected .

Targeted Device Management Actions (Continued)

To...	You CAN...
save your configuration	click Save .
discard your configuration without saving your changes	click Cancel .

The following procedure explains how to configure a NAT policy to manage targeted devices. See [Editing a NAT Policy](#) on page 394 for the complete procedure for editing a NAT policy.

To manage targeted devices in a NAT policy:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.
The NAT page appears.
2. Click the edit icon () next to the NAT policy you want to configure.
The NAT Policy Editor page appears.
3. Click the **Targets** tab.
The Targets page appears.
4. Optionally, click the **Search** prompt above the Available Devices list, then type a name.
The list updates as you type to display matching devices. You can click the clear icon () to clear the list.
5. Click the device, stack, cluster, or device group you want to add. Use Ctrl and Shift to select multiple devices.

TIP! You can also right-click an available device, then click **Select All**.

6. Click **Add to Policy**.
Selected devices are added.

TIP! You can also drag and drop to add devices.

7. Optionally, click the delete icon () to delete a device from the list of selected devices; or, use the Ctrl and Shift keys to select multiple devices, right-click, then select **Delete Selected**.
8. Click **Save** to save your configuration, or click **Cancel** to discard it.

Organizing Rules in a NAT Policy

LICENSE: Any

The Edit page for the NAT policy lists static NAT rules and dynamic NAT rules separately. The system sorts static rules alphabetically by name, and you cannot change the display order. You cannot create static rules with identical matching values. The system inspects static translations for a match before it inspects any dynamic translations.

Dynamic rules are processed in numerical order. The numeric position of each dynamic rule appears on the left side of the page next to the rule. You can move or insert dynamic rules and otherwise change the rule order. For example, if you move dynamic rule 10 under dynamic rule 3, rule 10 becomes rule 4 and all subsequent numbers increment accordingly.

A dynamic rule's position is important because the system compares packets to dynamic rules in the rules' numeric order on the policy Edit page. When a packet meets all the conditions of a dynamic rule, the system applies the conditions of that rule to the packet and ignores all subsequent rules for that packet.

Optionally, you can specify a dynamic rule's numeric position when you add or edit a dynamic rule. You can also highlight a dynamic rule before adding a new dynamic rule to insert the new rule below the rule you highlighted. See [Creating and Editing NAT Rules](#) on page 404.

You can select one or more dynamic rules by clicking a blank space in the row for the rule. You can drag and drop selected dynamic rules into a new location, thereby changing the position of the rules you moved and all subsequent rules.

You can cut or copy selected rules and paste them above or below an existing rule. You can only paste static rules in the Static Translations list and only dynamic rules in the Dynamic Translations list. You can also delete selected rules and insert new rules into any location in the list of existing rules.

IMPORTANT! You can copy, but not cut static rules.

You can display explanatory warnings to identify rules that will never match because they are preempted by preceding rules.

If you have access control policies in your deployment, the system does not translate traffic until it has passed through access control.

The following table summarizes the actions you can take to organize your rules.

NAT Rule Organization Actions

To...	You CAN...
select a rule	click a blank area in the row for a rule. Use the Ctrl or Shift key to select multiple rules. Rules you select are highlighted.
clear rule selections	click the reload icon () on the lower right side of the page. To clear individual rules, click a blank area in a rule's row while holding the Ctrl key.
cut or copy selected rules	right-click a blank area in the row for a selected rule, then select Cut or Copy . TIP! You can copy, but not cut static rules.
paste rules you have cut or copied into the rule list	right-click a blank area in the row for a rule where you want to paste selected rules, then select Paste above or Paste below . TIP! You can only paste static rules in the Static Translations list and only dynamic rules in the Dynamic Translations list.
move selected rules	drag and drop selected rules beneath a new location, indicated by a horizontal blue line that appears above your pointer as you drag.
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.
show warnings	click Show Warnings ; see Working with NAT Rule Warnings and Errors on page 390.

Working with NAT Rule Warnings and Errors

LICENSE: Any

The conditions of a NAT rule may preempt a subsequent rule from matching traffic. 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 were different.

The following table summarizes the actions you can take to show and clear warnings.

Preempted Rule Warning Actions

To...	YOU CAN...
show warnings	click Show Warnings . The page updates with an warning icon (⚠) next to each preempted rule.
display the warning for a rule	hover your pointer over the warning icon (⚠) next to a rule. A message indicates which rule preempts the rule.
clear warnings	click Hide Warnings . The page refreshes and the warnings disappear. TIP! Any action that refreshes the page, such as adding or editing a rule, or clicking the reload icon (🔄), also clears warnings.

If you create a rule that causes the NAT policy to fail upon apply, an error icon (❗) appears next to the rule. An error occurs if there is a conflict in the static rules, or if you edit a network object used in the policy that now makes the policy invalid. For example, an error occurs if you change a network object to use only IPv6 addresses and the rule that uses that object no longer has any valid networks where at least one network is required. Error icons appear automatically; you do not have to click **Show Warnings**.

Managing NAT Policies

LICENSE: Control

SUPPORTED DEVICES: Series 3

On the NAT policy page (**Devices > NAT**), you can view all your current NAT policies by name with optional description and the following status information:

- when a policy is up to date on targeted devices, in green text
- when a policy is out of date on targeted devices, in red text

Options on this page allow you to compare policies, create a new policy, apply a policy to targeted devices, copy a policy, view a report that lists all of the most recently saved settings in each policy, and edit a policy.

IMPORTANT! After you have applied a NAT policy to a managed device, you cannot delete the policy, even if it is out of date. Instead, you must apply a NAT policy with no rules to remove the applied NAT rules from the managed device.

The following table describes the actions you can take to manage your policies on the NAT policy page.

NAT Policy Management Actions

To...	You CAN...
create a new NAT policy	click New Policy . See Creating a NAT Policy on page 392 for more information.
modify the settings in an existing NAT policy	click the edit icon (). See Editing a NAT Policy on page 394 for more information.
apply a NAT policy to all devices targeted for the policy	click the policy apply icon (). See Applying a NAT Policy on page 401 for more information.
copy a NAT policy	click the copy icon (). See Copying a NAT Policy on page 395 for more information.
view a PDF report that lists the current configuration settings in a NAT policy	click the report icon (). See Viewing a NAT Policy Report on page 396 for more information.
compare NAT policies	click Compare Policies . See Comparing Two NAT Policies on page 397 for more information.
delete a NAT policy	click the delete icon () , then click OK , or click Cancel if you decide not to delete the policy. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy. IMPORTANT! After you have applied a NAT policy to a managed device, you cannot delete the policy from the device. Instead, you must apply a NAT policy with no rules to remove the applied NAT rules from the managed device. You also cannot delete a policy that is the last applied policy on any of its target devices, even if it is out of date. Before you can delete the policy completely, you must apply a different policy to those targets.

Creating a NAT Policy

LICENSE: Control

SUPPORTED DEVICES: Series 3

When you create a new NAT policy you must, at minimum, give it a unique name. Although you are not required to identify policy targets at policy creation time, you must perform this step before you can apply the policy; see [Managing NAT Policy Targets](#) on page 387. If you apply a NAT policy with no rules to a device, the system removes all NAT rules from that device.

To create a NAT policy:

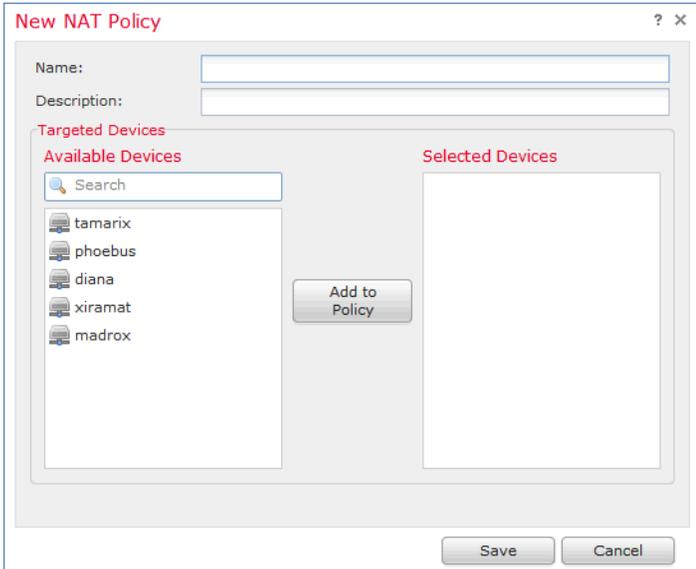
ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.

The NAT page appears.

2. Click **New Policy**.

The New NAT Policy pop-up window appears.



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

You can use all printable characters, including spaces and special characters.

4. Select the **Available Devices** where you want to apply the policy.

Use Ctrl and Shift to select multiple devices, or right-click to **Select All**. To narrow the devices that appear, type a search string in the **Search** field. To clear the search, click the clear icon (✕).

5. Add the **Selected Devices**. You can click and drag, or you can click **Add to Policy**.

6. Click **Save**.

The NAT policy Edit page appears. For information on configuring your new policy, including adding rules, see [Editing a NAT Policy](#) on page 394. Note that you must apply the policy for it to take effect; see [Applying a NAT Policy](#) on page 401.

Editing a NAT Policy

LICENSE: Control

SUPPORTED DEVICES: Series 3

On the NAT policy Edit page, you can configure your policy. See [Configuring NAT Policies](#) on page 385 and 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 NAT policy Edit page. If you attempt to exit the policy Edit page 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 Edit page.

To protect the privacy of your session, after 60 minutes of inactivity on the policy Edit page, changes to your policy are discarded and you are returned to the NAT page. After the first 30 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 resets the timer.

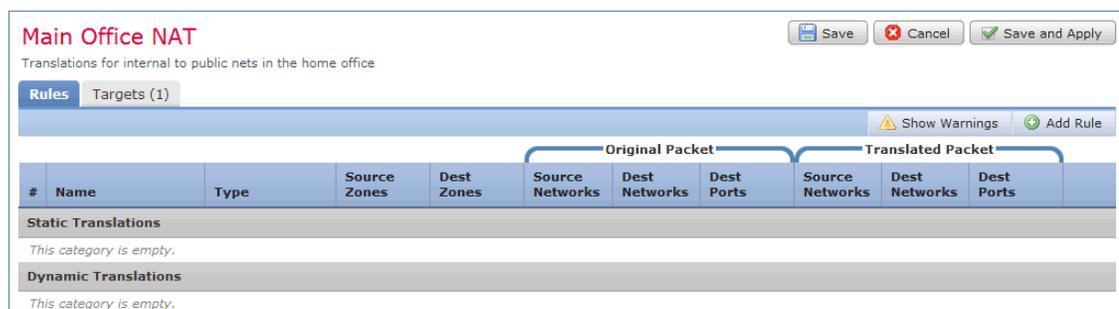
When you attempt to edit the same policy in two browser windows, you are prompted whether to resume your edit in the new window, discard your changes in the original window and continue editing in the new window, or cancel the second window and return to the policy Edit page.

When multiple users edit the same policy concurrently, a message for each user on the policy Edit page identifies other users who have unsaved changes. Any user who attempts to save changes is cautioned that saving changes will overwrite changes by other users. When multiple users save the same policy, the last saved changes are retained.

To edit a NAT policy:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.
The NAT page appears.
2. Click the edit icon (✎) next to the NAT policy you want to configure.
The NAT policy Edit page appears.



3. To configure your policy, take any of the actions described in [Configuring NAT Policies](#) on page 385.
4. Save or discard your configuration. You have the following choices:
 - To save your changes and continue editing, click **Save**.
 - To save your changes and apply your policy, click **Save and Apply**. See [Applying a NAT Policy](#) on page 401.
You must apply your policy to put your changes into effect.
 - To discard your changes, click **Cancel** and, if prompted, click **OK**.
Your changes are discarded and the NAT page appears.

Copying a NAT Policy

LICENSE: Control

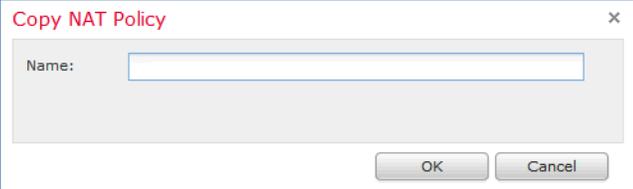
SUPPORTED DEVICES: Series 3

You can copy and rename a NAT policy. A policy you copy includes all policy rules and configurations.

To copy a NAT policy:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.
The NAT page appears.
2. Click the copy icon (📄) next to the NAT policy you want to configure.
The Copy NAT Policy pop-up window appears.



3. Enter a unique policy **Name**.
You can use any printable characters, including spaces and special characters.
4. Click **OK**.
Your copy appears on the NAT page in alphabetical order by name.

Viewing a NAT Policy Report

LICENSE: Control

SUPPORTED DEVICES: Series 3

A NAT 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 a NAT comparison report that compares a policy with the currently applied policy or with another policy. For more information, see [Comparing Two NAT Policies](#) on page 397.

A NAT policy report contains the sections described in the following table.

NAT 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 last modified it.
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. See Editing a NAT Policy on page 394.
Device Targets	Lists the managed devices targeted by the policy. See Managing NAT Policy Targets on page 387.
Rules	Provides the rule type and conditions for each rule in the policy. See Creating and Editing NAT Rules on page 404.
Referenced Objects	Provides the name and configuration of all individual objects and group objects used in the policy, by type of condition (Zones, Networks, and Ports) where the object is configured.

To view a NAT policy report:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.
The NAT page appears.

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 a NAT 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 Two NAT Policies

LICENSE: Control

SUPPORTED DEVICES: Series 3

To review policy changes, you can examine the differences between two NAT 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 policy comparison tools, see the following sections:

- [Using the NAT Policy Comparison View](#) on page 397
- [Using the NAT Policy Comparison Report](#) on page 399

Using the NAT Policy Comparison View

LICENSE: Control

SUPPORTED DEVICES: Series 3

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.

Main Office NAT (2013-05-23 09:22:27 by admin)	TX Sales Office (2013-05-23 09:24:18 by admin)
↔	
Policy Information	
Name	Main Office NAT
Description	Translations for internal to p
Modified	2013-05-23 09:22:27 by adi
Applied To	tamarix
Target Sensors	
tamarix	
Rules	
Static Translations	
Rule 1	
Name	Static Rule 1
Enabled	true
Source Zones and Interfaces	
s1p3	
Original Destination Networks	
"192.168.1.1/32"	
Dynamic Translations	
Rule 1	
Name	Dynamic 1
Enabled	true
Source Zones and Interfaces	
s1p3	

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.

NAT 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.

NAT Policy Comparison View Actions (Continued)

To...	YOU CAN...
generate a new policy comparison view	click New Comparison . The Select Comparison window appears. See Using the NAT Policy Comparison Report on page 399 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 NAT Policy Comparison Report

LICENSE: Control

SUPPORTED DEVICES: Series 3

A NAT policy comparison report is a record of all differences between two NAT 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 a NAT 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, while the policy comparison report lists only those configurations that differ between the policies. A NAT policy comparison report contains the sections described in the [NAT Policy Report Sections table](#) on page 396.

To compare two NAT policies:

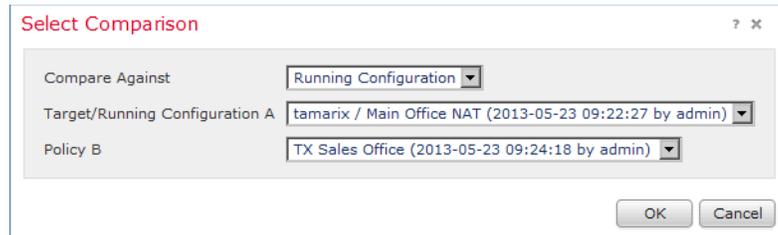
ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.

The NAT page appears.

2. Click **Compare Policies**.

The Select Comparison window appears.



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 two different revisions, select **Other Revision**.
The page refreshes and the Policy, Revision A and Revision 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.

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 different revisions, select the policy, then select the revisions you want to compare from the **Revision A** and **Revision 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.

5. Click **OK** to display the policy comparison view.

The comparison view appears.

6. Optionally, click **Comparison Report** to generate the NAT policy comparison report.

The NAT 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.

Applying a NAT Policy

LICENSE: Control

SUPPORTED DEVICES: Series 3

After making any changes to a NAT policy, you must apply the policy to one or more devices to implement the configuration changes on the networks monitored by the devices. You must target devices where you want to apply the policy before you can apply the policy. See [Managing NAT Policy Targets](#) on page 387.

Keep the following points in mind when applying NAT policies:

- You can configure and maintain multiple NAT policies on a Defense Center, but only one policy can be applied to a device at a time.
- You can apply two different NAT policies to different devices, even though they are both targets for multiple policies.
- You cannot apply a NAT policy to stacked devices running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). See [Managing Stacked Devices](#) on page 237 for more information.
- If you apply a device configuration that affects the interfaces in a NAT policy, the system reapplies the NAT policy on the device, including the interface changes. However, the policy remains unchanged on the DC and the interface displays an error icon ().

IMPORTANT! Applying an empty NAT policy removes all NAT rules from a device.

See the following sections for more information:

- [Applying a Complete NAT Policy](#) on page 401 explains how to use the quick-apply option to apply the NAT policy.
- [Applying Selected Policy Configurations](#) on page 402 explains how to select and apply configurations within the NAT policy.

Applying a Complete NAT Policy

LICENSE: Control

SUPPORTED DEVICES: Series 3

You can apply a NAT policy at any time. Applying a NAT policy also applies any associated rule configurations, objects, and policy changes to the devices targeted by the policy. A pop-up window allows you to apply all changes together as a single quick-apply action.

To quick-apply a complete NAT policy:

ACCESS: Admin/Network Admin

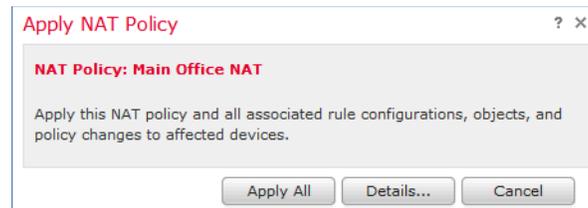
1. Select **Devices > NAT**.

The NAT page appears.

NAT Policy	Status	
APAC Regional Office	Targeting 1 devices Out-of-date on 1 targeted devices	✓ 📄 ✎ 🗑️
Main Office NAT Translations for internal to public nets in the home office	Targeting 1 devices Up-to-date on all targeted devices	✓ 📄 ✎ 🗑️
TX Sales Office	Targeting 1 devices Up-to-date on all targeted devices	✓ 📄 ✎ 🗑️

2. Click the apply icon (✓) next to the policy you want to apply.

The Apply NAT Rules pop-up window appears.



Alternatively, you can click **Save and Apply** on the policy Edit page; see [Editing a NAT Policy](#) on page 394.

3. Click **Apply All**.

Your policy apply task is queued. Click **OK** to return to the NAT page.

TIP! You can monitor the progress of the policy apply task on the Task Status page (**System > Monitoring > Task Status**).

Applying Selected Policy Configurations

LICENSE: Control

SUPPORTED DEVICES: Series 3

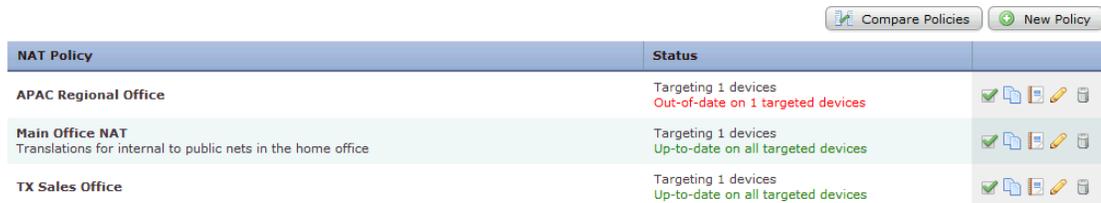
You can use the detailed policy apply page to apply changes to your NAT policy and to any designated targeted devices. The detailed page lists each device targeted by the policy and provides a column for the NAT policy by device. You can specify whether to apply changes to a NAT policy for each targeted device that is out of date.

To apply selected NAT policy configurations:

ACCESS: Admin/Network Admin

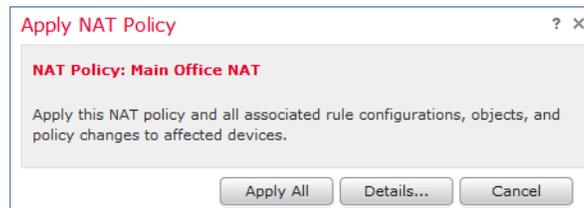
1. Select **Devices > NAT**.

The NAT page appears.



NAT Policy	Status	
APAC Regional Office	Targeting 1 devices Out-of-date on 1 targeted devices	<input checked="" type="checkbox"/>
Main Office NAT Translations for internal to public nets in the home office	Targeting 1 devices Up-to-date on all targeted devices	<input checked="" type="checkbox"/>
TX Sales Office	Targeting 1 devices Up-to-date on all targeted devices	<input checked="" type="checkbox"/>

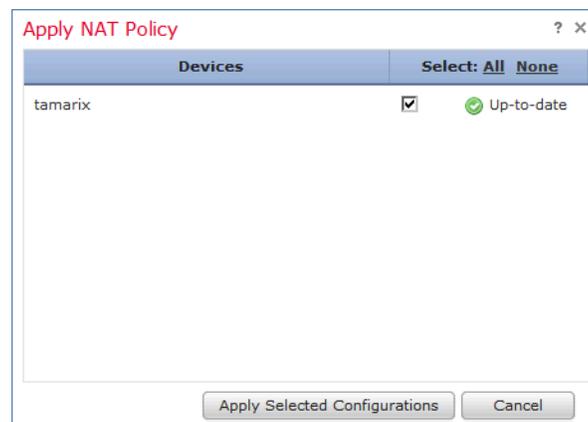
2. Click the apply icon () next to the policy you want to apply.
The Apply NAT Rules pop-up window appears.



Alternatively, you can click **Save and Apply** on the policy Edit page; see [Editing a NAT Policy](#) on page 394.

3. Click **Details**.

The detailed Apply NAT Rules pop-up window appears.



TIP! You can also open the pop-up window from the NAT page (**Devices > NAT**) by clicking on an out-of-date message in the **Status** column for the policy.

4. Select or clear the **NAT policy** check box next to the device name to specify whether to apply the NAT policy to a targeted device.

5. Click **Apply Selected Configurations**.

Your policy apply task is queued. Click **OK** to return to the NAT page.

TIP! You can monitor the progress of the policy apply task on the Task Status page (**System > Monitoring > Task Status**).

Creating and Editing NAT Rules

LICENSE: Control

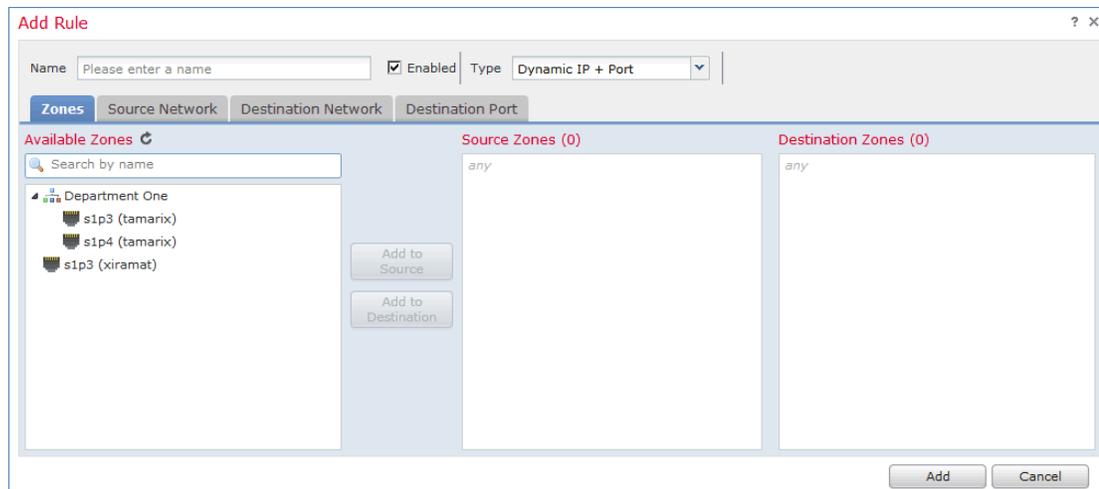
SUPPORTED DEVICES: Series 3

A NAT rule is simply a set of configurations and conditions that:

- qualifies network traffic
- specifies how the traffic that matches those qualifications is translated

You create and edit NAT rules from within an existing NAT policy. Each rule belongs to only one policy.

The web interface for adding or editing a rule is similar. You specify the rule name, state, type, and position (if dynamic) at the top of the page. You build conditions using the tabs on the left side of the page; each condition type has its own tab.



The following list summarizes the configurable components of a NAT rule.

Name

Give each rule a unique name. You can use up to thirty printable characters, including spaces and special characters, with the exception of the colon (:).

Rule State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic for translation. When viewing the list of rules in a NAT policy, disabled rules are grayed out, although you can still modify them.

Type

A rule's type determines how the system handles traffic that matches the rule's conditions. When you create and edit NAT rules, the configurable components vary according to rule type.

For detailed information on rule types and how they affect translation and traffic flow, see [Understanding NAT Rule Types](#) on page 406.

Position (Dynamic Rules Only)

Dynamic rules in a NAT policy are numbered, starting at 1. The system matches traffic to NAT rules in top-down order by ascending rule number.

When you add a rule to a policy, you specify its position by placing it **above** or **below** a specific rule, using rule numbers as a reference point. When editing an existing rule, you can **Move** the rule in a similar fashion. For more information, see [Organizing Rules in a NAT Policy](#) on page 389.

Conditions

Rule conditions identify the specific traffic you want to translate. Conditions can match traffic by any combination of multiple attributes, including security zone, network, and transport protocol port.

For detailed information on adding conditions, see [Understanding NAT Rule Conditions and Condition Mechanics](#) on page 409 and [Working with Different Types of Conditions in NAT Rules](#) on page 415.

To create or edit a NAT rule:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.
The NAT page appears.
2. Click the edit icon () next to the NAT policy where you want to add a rule.
The NAT policy Edit page appears.

3. Add a new rule or edit an existing rule:
 - 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.

Either the Add Rule or the Editing Rule page appears.

TIP! You can use the right-click context menu to perform many rule creation and management actions; see [Using the Context Menu](#) on page 63. You can also drag and drop rules to change their order.

4. Configure the rule components, as described earlier in this section. You can configure the following, or accept the defaults:
 - You must provide a unique rule **Name**.
 - Specify whether the rule is **Enabled**.
 - Select a rule **Type**.
 - Specify the rule position (dynamic rules only).
 - Configure the rule's conditions.
 - Static rules must include an original destination network.
 - Dynamic rules must include a translated source network.
5. Click **Add** or **Save**.

Your changes are saved. You must apply the NAT policy for your changes to take effect; see [Applying a NAT Policy](#) on page 401.

Understanding NAT Rule Types

LICENSE: Any

Every NAT rule has an associated type that:

- qualifies network traffic
- specifies how the traffic that matches those qualifications is translated

The following list summarizes the NAT rule types.

Static

Static rules provide one-to-one translations on destination networks and optionally port and protocol. When configuring static translations, you can configure source zones, destination networks, and destination ports. You cannot configure destination zones or source networks.

You **must** specify an original destination network. For destination networks, you can only select network objects and groups containing a single IP address or

enter literal IP addresses that represent a single IP address. You can only specify a single original destination network and a single translated destination network.

Optionally, you can specify a single original destination port and a single translated destination port. You must specify an original destination network before you can specify an original destination port. In addition, you cannot specify a translated destination port unless you also specify an original destination port, and the translated value must match the protocol of the original value.

WARNING! For static NAT rules on a clustered device, only select an individual peer interface if all networks affected by the NAT translations are private. Do **not** use this configuration for static NAT rules affecting traffic between public and private networks.

Dynamic IP Only

Dynamic IP Only rules translate many-to-many source networks, but maintain port and protocol. When configuring dynamic IP only translations, you can configure zones, source networks, original destination networks, and original destination ports. You cannot configure translated destination networks or translated destination ports.

You **must** specify at least one translated source network. If the number of translated source network values is less than the number of original source networks, the system displays a warning on the rule that it is possible to run out of translated addresses before all original addresses are matched.

If there are multiple rules with conditions that match the same packet, the low priority rules become dead, meaning they can never be triggered. The system also displays warnings for dead rules. You can view tooltips to determine which rule supersedes the dead rule.

IMPORTANT! You can save and apply policies with dead rules, but the rules cannot provide any translation.

In some instances, you may want to create rules with limited scope preceding rules with a broader scope. For example:

Rule 1: Match on address A and port A/Translate to address B
Rule 2: Match on address A/Translate to Address C

In this example, rule 1 matches some packets that also match rule 2. Therefore, rule 2 is not completely dead.

Optionally, you can specify only original destination ports. You cannot specify translated destination ports.

Dynamic IP + Port

Dynamic IP and port rules translate many-to-one or many-to-many source networks and port and protocol. When configuring dynamic IP and port translations, you can configure zones, source networks, original destination networks, and original destination ports. You cannot configure translated destination networks or translated destination ports.

You **must** specify at least one translated source network. If there are multiple rules with conditions that match the same packet, the low priority rules become dead, meaning they can never be triggered. The system also displays warnings for dead rules. You can view tool tips to determine which rule supersedes the dead rule.

IMPORTANT! You can save and apply policies with dead rules, but the rules cannot provide any translation.

Optionally, you can specify only original destination ports. You cannot specify translated destination ports.

IMPORTANT! If you create a dynamic IP and port rule, and the system passes traffic that does not use a port, no translation occurs for the traffic. For example, a ping (ICMP) from an IP address that matches the source network does not map, because ICMP does not use a port.

The following table summarizes the NAT rule condition types that can be configured based on the specified NAT rule type:

Available NAT Rule Condition Types per NAT Rule Type

CONDITION	STATIC	DYNAMIC (IP ONLY OR IP + PORT)
Source Zones	Optional	Optional
Destination Zones	Not allowed	Optional
Original Source Networks	Not allowed	Optional
Translated Source Networks	Not allowed	Required
Original Destination Networks	Required	Optional

Available NAT Rule Condition Types per NAT Rule Type (Continued)

CONDITION	STATIC	DYNAMIC (IP ONLY OR IP + PORT)
Translated Destination Networks	Optional; single address only	Not allowed
Original Destination Ports	Optional; single port only, and only allowed if you define the original destination network	Optional
Translated Destination Ports	Optional; single port only, and only allowed if you define the original destination port	Not allowed

Understanding NAT Rule Conditions and Condition Mechanics

LICENSE: Any

You can add conditions to NAT rules to identify the type of traffic that matches the rule. For each condition type, you select conditions you want to add to a rule from a list of available conditions. When applicable, condition filters allow you to constrain available conditions. Lists of available and selected conditions may be as short as a single condition or many pages long. You can search available conditions and display only those matching a typed name or value in a list that updates as you type.

Depending on the type of condition, lists of available conditions may be comprised of a combination of conditions provided directly by Sourcefire or configured using other Sourcefire 3D System features, including objects created using the object manager (**Objects > Object Management**), objects created directly from individual conditions pages, and literal conditions.

See the following sections for information on specifying rule conditions:

- [Understanding NAT Rule Conditions](#) on page 410 defines the different types of rule conditions.
- [Adding Conditions to NAT Rules](#) on page 411 describes the controls used to select and add rule conditions.
- [Searching NAT Rule Condition Lists](#) on page 413 explains how to search available conditions and display only those matching a typed name or value in a list that updates as you type.
- [Adding Literal Conditions to NAT Rules](#) on page 414 explains how to add literal conditions to a rule.
- [Using Objects in NAT Rule Conditions](#) on page 415 explains how to add individual objects to the system from the configuration pages for relevant condition types.

Understanding NAT Rule Conditions

LICENSE: Any

You can set a NAT rule to match traffic meeting any of the conditions described in the following table:

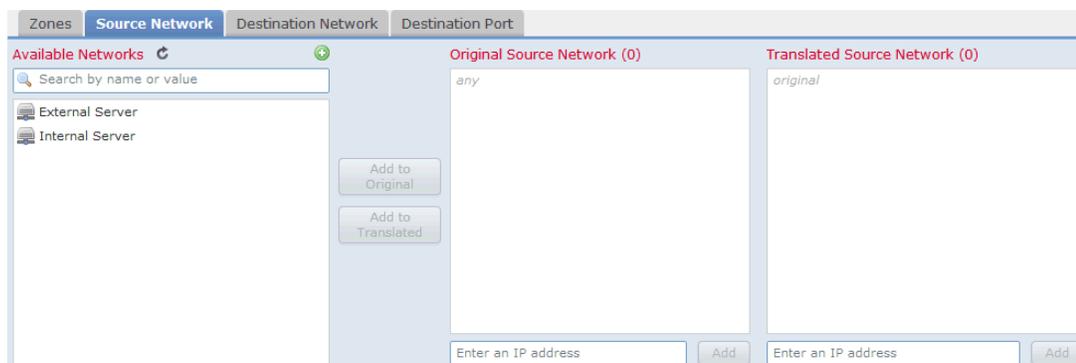
NAT Rule Condition Types

CONDITION	DESCRIPTION	SUPPORTED DEFENSE CENTERS	SUPPORTED DEVICES
Zones	A configuration of one or more routed interfaces where you can apply NAT policies. Zones provide a mechanism for classifying traffic on source and destination interfaces, and you can add source and destination zone conditions to rules. See Working with Security Zones on page 187 for information on creating zones using the object manager.	Any	Series 3
Networks	Any combination of individual IP addresses, CIDR blocks, and prefix lengths, either specified explicitly or using network objects and groups (see Working with Network Objects on page 161). You can add source and destination network conditions to NAT rules.	Any	Series 3
Destination Ports	Transport protocol ports, including individual and group port objects you create based on transport protocols. See Working with Port Objects on page 173 for information on creating individual and group transport protocol objects using the object manager.	Any	Series 3

Adding Conditions to NAT Rules

LICENSE: Any

Adding conditions to NAT rules is essentially the same for each type of condition. You select from a list of available conditions on the left, and add the selected conditions to one or two lists of selected conditions on the right.



For all condition types, you select one or more individual available conditions by clicking on them to highlight them. You can either click a button between the two types of lists to add available conditions that you select to your lists of selected conditions, or drag and drop available conditions that you select into the list of selected conditions.

You can add up to 50 conditions of each type to a list of selected conditions. For example, you can add up to 50 source zone conditions, up to 50 destination zone conditions, up to 50 source network conditions, and so on, until you reach the upper limit for the appliance.

The following table describes the actions you can take to select and add conditions to a rule.

Adding Conditions to a NAT Rule

To...	You CAN...
select available conditions to add to a list of selected conditions	click the available condition; use the Ctrl and Shift keys to select multiple conditions.
select all listed available conditions	right-click the row for any available condition, then click Select All .
search a list of available conditions or filters	click inside the Search field and type a search string. See Searching NAT Rule Condition Lists on page 413 for more information.

Adding Conditions to a NAT Rule (Continued)

To...	YOU CAN...
clear a search when searching available conditions or filters	click the reload icon () above the Search field or the clear icon () in the Search field.
add selected zone conditions from a list of available conditions to a list of selected source or destination conditions	click Add to Source or Add to Destination . See Adding Zone Conditions to NAT Rules on page 415 for more information.
add selected network and port conditions from a list of available conditions to a list of selected original or translated conditions	click Add to Original or Add to Translated . See Adding Source Network Conditions to Dynamic NAT Rules on page 418, Adding Destination Network Conditions to NAT Rules on page 419, or Adding Port Conditions to NAT Rules on page 421 for more information.
drag and drop selected available conditions into a list of selected conditions	click a selected condition, then drag and drop into the list of selected conditions.
add a literal condition to a list of selected conditions using a literal field	click to remove the prompt from the literal field, type the literal condition, then click Add . Network conditions provide a field for adding literal conditions.
add a literal condition to a list of selected conditions using a drop-down list	select a condition from the drop-down list, then click Add . Port conditions provide a drop-down list for adding literal conditions. See Adding Port Conditions to NAT Rules on page 421 for more information.
add an individual object or condition filter so you can then select it from the list of available conditions	click the add icon (). See Using Objects and Security Zones on page 158 for information on adding objects using the object manager.
delete a single condition from a list of selected conditions	click the delete icon () next to the condition
delete a condition from a list of selected conditions	right-click to highlight the row for a selected condition, then click Delete .
delete multiple conditions from a list of selected conditions	use the Shift and Ctrl keys to select multiple conditions, or right-click and Select All ; next, right-click to highlight the row for a selected condition, then click Delete Selected .

On the relevant condition page, and also on the policy Edit page, you can hover your pointer over an individual object to display the contents of the object, and over a group object to display the number of individual objects in the group.

The following basic procedure explains how to add conditions to a new rule. See [Creating and Editing NAT Rules](#) on page 404 for complete instructions on adding and modifying rules.

To add available conditions to a list of selected conditions:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.
The NAT page appears.
2. Click the edit icon () next to the NAT policy you want to modify.
The policy Edit page appears.
3. Click **Add Rule**.
The Add Rule page appears.



4. Click the tab for the type of condition you want to add to the rule.
The conditions page appears for the type of condition you selected.
5. Take any of the available actions in the [Adding Conditions to a NAT Rule table](#) on page 411.
6. Click **Add** to save your configuration.
Your rule is added and the policy Edit page appears.

Searching NAT Rule Condition Lists

LICENSE: Any

You can filter a list of available NAT rule conditions to limit the number of items displayed in the list. The list updates as you type to display matching items.

Optionally, you can search on object names and on the values configured for 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**.

The following basic procedure explains how to filter a list in a new rule. See [Creating and Editing NAT Rules](#) on page 404 for complete instructions on adding and modifying rules.

To search a list of available conditions:

ACCESS: Admin/Network Admin

1. Select **Devices > NAT**.

The NAT page appears.

2. Click the edit icon (✎) next to the NAT policy you want to modify.

The policy Edit page appears.

3. Click **Add Rule**.

The Add Rule page appears. The following graphic shows the search field on the Source Network and Destination Network conditions page.



4. To search a list, click inside the search field to clear the prompt, then type a search string.

The list updates as you type to display matching items and a clear list icon (✕) appears in the search field. The list updates and no items are listed when none match the search string.

5. Optionally, click the reload icon (↻) above the **Search** field or click the clear icon (✕) in the **Search** field to clear the search string.

The complete list appears.

6. Click **Add** to save your configuration.

Your rule is added and the policy Edit page appears.

Adding Literal Conditions to NAT Rules

LICENSE: Any

You can add a literal value to the list of original and translated conditions for the following condition types:

- Networks
- Ports

For network conditions, you type the literal value in a configuration field below the list of original or translated conditions.

In the case of port conditions, you select a protocol from a drop-down list. When the protocol is **ATM** and, optionally, when the protocol is **TCP** or **UDP**, you type a port number in a configuration field.

Each relevant conditions page provides the controls needed to add literal values. Values you type in a configuration field appear as red text if the value is invalid, or until it is recognized as valid. Typed values change to blue text as you type when they are recognized as valid. A grayed **Add** button activates when a valid value is

recognized. Literal values you add appear immediately in the list of selected conditions.

See the following sections for specific details on adding each type of literal value:

- [Adding Source Network Conditions to Dynamic NAT Rules](#) on page 418
- [Adding Destination Network Conditions to NAT Rules](#) on page 419
- [Adding Port Conditions to NAT Rules](#) on page 421

Using Objects in NAT Rule Conditions

LICENSE: Any

Objects that you create in the object manager (**Objects > Object Management**) are immediately available for you to select from relevant lists of available NAT rule conditions. See [Using Objects and Security Zones](#) on page 158 for information.

You can also create objects on-the-fly from the NAT policy. A control on relevant conditions pages provides access to the same configuration controls that you use in the object manager.

Individual objects created on-the-fly appear immediately in the list of available objects. You can add them to the current rule, and to other existing and future rules. On the relevant conditions page, and also on the policy Edit page, you can hover your pointer over an individual object to display the contents of the object, and over a group object to display the number of individual objects in the group.

Working with Different Types of Conditions in NAT Rules

LICENSE: Any

You can match traffic with one or more rule conditions. See the following sections for more information:

- [Adding Zone Conditions to NAT Rules](#) on page 415 explains how to match traffic by security zones that you create using the object manager.
- [Adding Source Network Conditions to Dynamic NAT Rules](#) on page 418 and [Adding Destination Network Conditions to NAT Rules](#) on page 419 explain how to match traffic by IP address or address block.
- [Adding Port Conditions to NAT Rules](#) on page 421 explains how to match traffic by specified transport protocol ports.

Adding Zone Conditions to NAT Rules

LICENSE: Any

The security zones on your system are comprised of interfaces on your managed devices. Zones that you add to a NAT rule target the rule to devices on your network that have routed or hybrid interfaces in those zones. You can only add security zones with routed or hybrid interfaces as conditions for NAT rules. See

[Working with Security Zones](#) on page 187 for information on creating security zones using the object manager.

You can add either zones or standalone interfaces that are currently assigned to a virtual router to NAT rules. If there are devices with unapplied device configurations, the Zones page displays a warning icon (⚠) at the top of the available zones list, indicating that only applied zones and interfaces are displayed. You can click the arrow icon (▸) next to a zone to collapse or expand the zone to hide or view its interfaces.

If an interface is on a clustered device, the available zones list displays an additional branch from that interface with the other interfaces in the cluster as children of the primary interface on the active device in the cluster. You can also click the arrow icon (▸) to collapse or expand the clustered device interfaces to hide or view its interfaces.

IMPORTANT! You can save and apply policies with disabled interfaces, but the rules cannot provide any translation until the interfaces are enabled.

The two lists on the right are the source and destination zones used for matching purposes by the NAT rules. If the rule already has values configured, these lists display the existing values when you edit the rule. If the source zones list is empty, the rule matches traffic *from* any zone or interface. If the destination zones list is empty, the rule matches traffic *to* any zone or interface.

The system displays warnings for rules with zone combinations that never trigger on a targeted device.

IMPORTANT! You can save and apply policies with these zone combinations, but the rules will not provide any translation.

You can add individual interfaces by selecting an item in a zone or by selecting a standalone interface. You can only add interfaces in a zone if the zone it is assigned to has not already been added to a source zones or destination zones list. These individually selected interfaces are not affected by changes to zones, even if you remove them and add them to a different zone. If an interface is the primary member of a cluster and you are configuring a dynamic rule, you can add only the primary interface to the source zones or destination zones list. For static rules, you can add individual cluster member interfaces to the source zones list. You can only add a primary cluster interface to a list if none of its children have been added, and you can only add individual cluster interfaces if the primary has not been added.

If you add a zone, the rule uses all interfaces associated with the zone. If you add or remove an interface from the zone, the rule will not use the updated version of

the zone until the device configuration has been reapplied to the devices where the interfaces reside.

IMPORTANT! In a static NAT rule, you can add only source zones. In a dynamic NAT rule, you can add both source and destination zones.

The following procedure explains how to add source and destination zone conditions while adding or editing a NAT rule. See [Understanding NAT Rule Conditions and Condition Mechanics](#) on page 409 for more detailed information.

To add zone conditions to a NAT rule:

ACCESS: Admin/Network Admin

1. Select the **Zones** tab on the rule Edit page.

The Zones page appears.



2. Optionally, click the **Search by name** prompt above the **Available Zones** list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching NAT Rule Condition Lists](#) on page 413 for more information.

3. Click a zone or interface in the **Available Zones** list. Use the Shift and Ctrl keys to select multiple conditions, or right-click and then click **Select All**.

Conditions you select are highlighted.

4. You have the following choices:

- To match traffic by source zone, click **Add to Source**.
- To match traffic by destination zone, click **Add to Destination**.

Optionally, you can drag and drop selected conditions into the **Source Zones** or **Destination Zones** lists.

Selected conditions are added. Note that while you can add disabled interfaces to a NAT rule, the rule does not provide any translation.

IMPORTANT! You can add only source zones to static NAT rules.

5. Save or continue editing the rule.

You must apply the NAT policy for your changes to take effect; see [Applying a NAT Policy](#) on page 401.

Adding Source Network Conditions to Dynamic NAT Rules

LICENSE: Any

You configure the matching values and translation values of the source IP address for packets. If the original source network is not configured, then any source IP address matches the dynamic NAT rule. Note that you cannot configure source networks for static NAT rules. If a packet matches the NAT rule, the system uses the values in the translated source network to assign the new value for the source IP address. For dynamic rules, you must configure a translated source network with at least one value.

WARNING! If a network object or object group is being used by a NAT rule, and you change or delete the object or group, it can cause the rule to become invalid.

You can add any of the following kinds of source network conditions to a dynamic NAT rule:

- individual and group network objects that you have created using the object manager
See [Working with Network Objects](#) on page 161 for information on creating individual and group network objects using the object manager.
- individual network objects that you add from the Source Network conditions page, and can then add to your rule and to other existing and future rules
See [Using Objects in NAT Rule Conditions](#) on page 415 for more information.
- literal, single IP addresses, ranges, or address blocks
See [Adding Literal Conditions to NAT Rules](#) on page 414 for more information.

The following procedure explains how to add source network conditions while adding or editing a dynamic NAT rule. See [Understanding NAT Rule Conditions and Condition Mechanics](#) on page 409 for more detailed information.

To add network conditions to a dynamic NAT rule:

ACCESS: Admin/Network Admin

1. Select the **Source Networks** tab on the rule Edit page.
The Source Network page appears.



2. Optionally, click the **Search by name or value** prompt above the **Available Networks** list, then type a name or value.
The list updates as you type to display matching conditions. See [Searching NAT Rule Condition Lists](#) on page 413 for more information.
3. Click a condition in the **Available Networks** list. Use the Shift and Ctrl keys to select multiple conditions, or right-click and then click **Select All**.
Conditions you select are highlighted.
4. You have the following choices:
 - To match traffic by original source network, click **Add to Original**.
 - To specify the translation value for traffic that matches the translated source network, click **Add to Translated**.Alternatively, you can drag and drop selected conditions into the **Original Source Network** or **Translated Source Network** lists.
Conditions you selected are added.
5. Optionally, click the add icon () above the **Available Networks** list to add an individual network object.
You can add multiple IP addresses, CIDR blocks, and prefix lengths to each network object.
Optionally, you can then select the object you added. See [Working with Network Objects](#) on page 161 and [Using Objects in NAT Rule Conditions](#) on page 415 for more information.
6. Optionally, click the **Enter an IP address** prompt below the **Original Source Network** or **Translated Source Network** list; then type an IP address, range, or address block and click **Add**.
You add ranges in the following format: lower IP address-upper IP address.
For example: 179.13.1.1-179.13.1.10.
The list updates to display your entry. See [Adding Literal Conditions to NAT Rules](#) on page 414 for more information.
7. Save or continue editing the rule.
You must apply the NAT policy for your changes to take effect; see [Applying a NAT Policy](#) on page 401.

Adding Destination Network Conditions to NAT Rules

LICENSE: Any

You configure the matching values and translation values of the destination IP address for packets. Note that you cannot configure translated destination networks for dynamic NAT rules.

Because static NAT rules are one-to-one translations, the **Available Networks** list contains only network objects and groups that contain only a single IP address.

For static translations, you can add only a single object or literal value to both the **Original Destination Network** or **Translated Destination Network** lists.

WARNING! If a network object or object group is being used by a NAT rule, and you change or delete the object or group, it can cause the rule to become invalid.

You can add any of the following kinds of destination network conditions to a NAT rule:

- individual and group network objects that you have created using the object manager
See [Working with Network Objects](#) on page 161 for information on creating individual and group network objects using the object manager.
- individual network objects that you add from the Destination Network conditions page, and can then add to your rule and to other existing and future rules
See [Using Objects in NAT Rule Conditions](#) on page 415 for more information.
- literal, single IP addresses, range, or address blocks
For static NAT rules, you can add only a CIDR with subnet mask /32, and only if there is not already a value in the list.
See [Adding Literal Conditions to NAT Rules](#) on page 414 for more information.

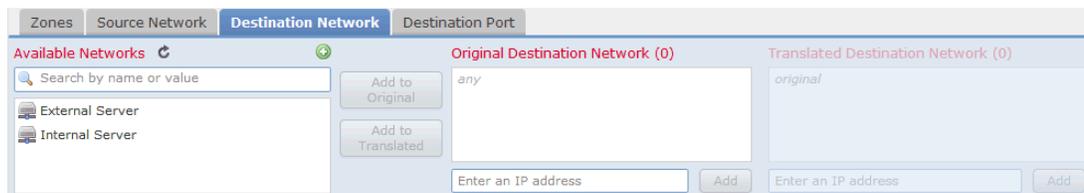
The following procedure explains how to add destination network conditions while adding or editing a NAT rule. See [Understanding NAT Rule Conditions and Condition Mechanics](#) on page 409 for more detailed information.

To add destination network conditions to a NAT rule:

ACCESS: Admin/Network Admin

1. Select the **Destination Network** tab on the rule Edit page.

The Destination Network page appears.



2. Optionally, click the **Search by name or value** prompt above the **Available Networks** list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching NAT Rule Condition Lists](#) on page 413 for more information.

3. Click a condition in the **Available Networks** list. Use the Shift and Ctrl keys to select multiple conditions, or right-click and then click **Select All**.

Conditions you select are highlighted.

4. You have the following choices:
 - To match traffic by original destination network, click **Add to Original**.
 - To specify the translation value for traffic that matches the translated destination network, click **Add to Translated**.

Alternatively, you can drag and drop selected conditions into the **Original Destination Network** or **Translated Destination Network** lists.

Conditions you selected are added.

5. Optionally, click the add icon (+) above the **Available Networks** list to add an individual network object.

For dynamic rules, you can add multiple IP addresses, CIDR blocks, and prefix lengths to each network object. For static rules, you can add only a single IP address. Optionally, you can then select the object you added. See [Working with Network Objects](#) on page 161 and [Using Objects in NAT Rule Conditions](#) on page 415 for more information.

6. Optionally, click the **Enter an IP address** prompt below the **Original Destination Network** or **Translated Destination Network** list, then type an IP address or address block and click **Add**.

The list updates to display your entry. See [Adding Literal Conditions to NAT Rules](#) on page 414 for more information.

7. Save or continue editing the rule.

You must apply the NAT policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Adding Port Conditions to NAT Rules

LICENSE: Any

You can add a port condition to a rule to match network traffic based on the original and translated destination port and transport protocol for translation. If the original port is not configured, any destination port matches the rule. If a packet matches the NAT rule and a translated destination port is configured, the system translates the port into that value. Note that for dynamic rules, you can specify only the original destination port. For static rules, you can define a translated destination port, but only with an object with the same protocol as the original destination port object or literal value.

The system matches the destination port against the value of the port object or literal port in the original destination port list for static rules, or multiple values for dynamic rules.

Because static NAT rules are one-to-one translations, the **Available Ports** list contains only port objects and groups that contain only a single port. For static translations, you can add only a single object or literal value to both the **Original Port** or **Translated Port** lists.

For dynamic rules, you can add a range of ports. For example, when specifying the original destination port, you can add **1000-1100** as a literal value.

WARNING! If a port object or object group is being used by a NAT rule, and you change or delete the object or group, it can cause the rule to become invalid.

You can add any of the following kinds of port conditions to a NAT rule:

- individual and group port objects that you have created using the object manager
See [Working with Port Objects](#) on page 173 for information on creating individual and group port objects using the object manager.
- individual port objects that you add from the Destination Ports conditions page, and can then add to your rule and to other existing and future rules
See [Using Objects in NAT Rule Conditions](#) on page 415 for more information.
- literal port values, consisting of a TCP, UDP, or All (TCP and UDP) transport protocol and a port
See [Adding Literal Conditions to NAT Rules](#) on page 414 for more information.

The following procedure explains how to add port conditions while adding or editing a NAT rule. See [Understanding NAT Rule Conditions and Condition Mechanics](#) on page 409 for more detailed information.

To add destination port conditions to a NAT rule:

ACCESS: Admin/Network Admin

1. Select the **Destination Port** tab on the rule Edit page.

The Destination Port page appears.



2. Optionally, click the **Search by name or value** prompt above the **Available Ports** list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching NAT Rule Condition Lists](#) on page 413 for more information.

3. Click a condition in the **Available Ports** list. Use the Shift and Ctrl keys to select multiple conditions, or right-click to select all conditions. Note that you can add a maximum of 50 conditions.

Conditions you select are highlighted.

4. You have the following choices:
 - Click **Add to Original** to add the selected port to the Original Ports list.
 - Click **Add to Translated** to add the selected port to the Translated Ports list.
 - Drag and drop available ports into a list.

5. Optionally, to create and add an individual port object click the add icon () above the **Available Ports** list.

You can identify a single port or a port range in each port object that you add. You can then select objects you added as conditions for your rule. See [Using Objects in NAT Rule Conditions](#) on page 415 for more information.

For static rules, you can use only port objects with single ports.

6. Optionally, to add a literal port select an entry from the **Protocol** drop-down list beneath the **Original Port** or **Translated Port** lists.

Enter a port, then click **Add**. You can specify a port number from 0 through 65535. For dynamic rules, you can specify a single port or a range.

The list updates to display your selection. See [Adding Literal Conditions to NAT Rules](#) on page 414 for more information.

Conditions you selected are added

7. Save or continue editing the rule.

You must apply the NAT policy for your changes to take effect; see [Applying a NAT Policy](#) on page 401.

CHAPTER 12

USING ACCESS CONTROL POLICIES

An *access control policy* determines how the system handles non-fast-pathed traffic on your network. You can configure one or more access control policies, which you can then apply to one or more managed devices. Each device can have one currently applied policy.

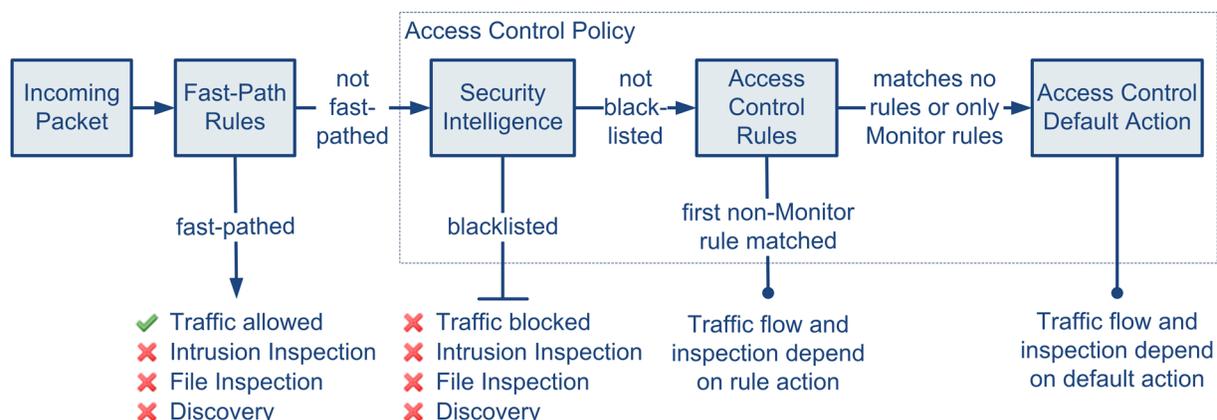
A simple access control policy can filter (blacklist or monitor) traffic based on Security Intelligence data, then use the policy's *default action* to handle non-blacklisted traffic in one of the following ways:

- block all traffic from entering your network
- trust all traffic to enter your network without further inspection
- allow all traffic to enter your network, and inspect the traffic with a network discovery policy only
- allow all traffic to enter your network, and inspect the traffic with intrusion and network discovery policies

Optionally, you can add *access control rules* to a policy, which provide granular control over how you handle and log network traffic. For each rule, you specify a rule *action*, that is, whether to trust, monitor, block, or inspect matching traffic with an intrusion or file policy. Each rule contains a set of conditions that identify the specific traffic you want to control. Rules can be simple or complex, matching traffic by any combination of security zone, network, VLAN, Active Directory LDAP user or group, application, transport protocol port, or URL.

The system matches traffic to access control rules in order; the first matched rule handles the traffic. (An exception occurs with Monitor rules, which allow traffic to continue to be evaluated.)

The diagram below illustrates traffic flow through the Sourcefire 3D System, and provides some details on the types of inspection performed on that traffic. Notice that the system does not inspect fast-pathed or blacklisted traffic. For traffic handled by an access control rule or default action, flow and inspection depend on the rule action. Although rule actions are not shown in the diagram for simplicity, the system does not perform any kind of inspection on trusted or blocked traffic. Additionally, file inspection is not supported with the default action.



This chapter contains information on creating a basic access control policy (including Security Intelligence filtering) and adding rules to that policy. For detailed information on associated components of the Sourcefire 3D System, see the following documentation:

- [Configuring Fast-Path Rules](#) on page 253
- [Understanding and Writing Access Control Rules](#) on page 472
- [Understanding and Creating File Policies](#) on page 1206
- [Configuring Intrusion Policies](#) on page 667
- [Introduction to Network Discovery](#) on page 1250

Although you can create access control policies regardless of the licenses on your Defense Center, certain aspects of access control require that you enable specific licensed capabilities on target devices before you can apply the policy. Additionally, some features are only available on certain appliance models. The Defense Center uses warning icons (⚠) and confirmation dialog boxes to designate unsupported features for your deployment. For details, hover your pointer over a warning icon.

The following table explains the license and appliance model requirements to apply access control policies. Note that Series 2 devices automatically have most

Protection capabilities; you do not have to explicitly enable Protection on those devices.

License and Appliance Requirements for Access Control

TO APPLY A POLICY THAT...	ADD THIS LICENSE...	TO ONE OF THESE DEFENSE CENTERS...	AND ENABLE IT ON ONE OF THESE DEVICES...
performs access control based on zone, network, VLAN, or port, or that performs URL filtering using literal URLs and URL objects	Any	Any	Any, except Series 2 devices cannot perform URL filtering using literal URLs and URL objects
performs intrusion detection and prevention, file control, or Security Intelligence filtering	Protection	Any	Any, except Series 2 devices cannot perform Security Intelligence filtering
performs advanced malware protection, that is, network-based malware detection and blocking	Malware	Any except DC500	Series 3, virtual
performs user or application control	Control	Any, except the DC500 cannot perform user control	Series 3, virtual
performs URL filtering using category and reputation data	URL Filtering	Any except DC500	Series 3, virtual

See the following sections for more information on creating and managing access control policies:

- [Configuring Policies](#) on page 426
- [Organizing Rules in a Policy](#) on page 449
- [Managing Access Control Policies](#) on page 456

Configuring Policies

LICENSE: Any

To configure an access control policy, you must give the policy a unique name, specify a default action, and identify the devices, or *targets*, where you want to apply the policy.

You can also:

- blacklist (deny without further inspection) traffic based on Security Intelligence data before that traffic can be inspected by any access control rules; optionally you can monitor traffic based on that same data
- add, edit, delete, enable, and disable access control rules
- configure an HTML page (called the *HTTP response page*) that users see when an access control rule blocks their HTTP request
- configure advanced settings, such as the number of URL characters to store in connection events, the depth or duration of file and malware inspection, and the duration of bypasses for interactively blocked sessions
- log traffic that is handled by the default action

After you create or modify an access control policy, you can apply the policy to all or some targeted devices. 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 policy Edit page.

Access Control 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 the Default Action on page 428.
log connections for the default action	find more information at Logging Connections for the Default Action on page 431.
assign different rights to different users	find more information at Using Custom User Roles with Access Control Policies on page 433.
manage policy targets	find more information at Managing Policy Targets on page 434.
save your policy changes	click Save .
save and apply your policy	click Save and Apply . See Applying an Access Control Policy on page 466 for more information. TIP! You can also click the edit icon () next to your policy on the Access Control page.
cancel your policy changes	click Cancel , then, if you have made changes, click OK .

Access Control Policy Configuration Actions (Continued)

To...	You CAN...
add a rule to a policy	click Add Rule . See Understanding and Writing Access Control Rules on page 472 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 Creating and Editing Access Control Rules on page 474 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 Working with Different Types of Conditions on page 491 for more information.
configure a response page to blocked HTTP requests	find more information at Adding an HTTP Response Page on page 436.
filter traffic based on Security Intelligence data	find more information in Filtering Traffic Based on Security Intelligence Data on page 438.
configure advanced settings	find more information in Configuring Advanced Access Control Policy Settings on page 447.

Setting the Default Action

LICENSE: Any

The default action for an access control policy determines how the system handles traffic that:

- is not blacklisted by Security Intelligence
- does not match any non-Monitor rule in the policy

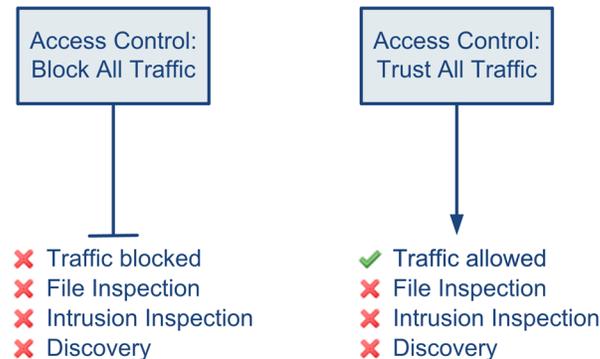
When you apply an access control policy that does not contain any access control rules or Security Intelligence configurations, the default action determines how all traffic on your network is handled.

The following table lists the default actions you can choose, as well as their effect on traffic and the types of inspection performed on traffic handled by each option.

Access Control Policy Default Actions

DEFAULT ACTION	EFFECT ON TRAFFIC	INSPECTION
Access Control: Block All Traffic	block without further inspection	none
Access Control: Trust All Traffic	trust (allow without further inspection)	none
Network Discovery Only	allow	network discovery
Intrusion Prevention	allow, as long as it is passed by the intrusion policy you specify	intrusion and network discovery

The diagram below illustrates the block and trust Access Control default actions. Notice that the system does not perform any kind of inspection on traffic blocked or trusted by the default action.



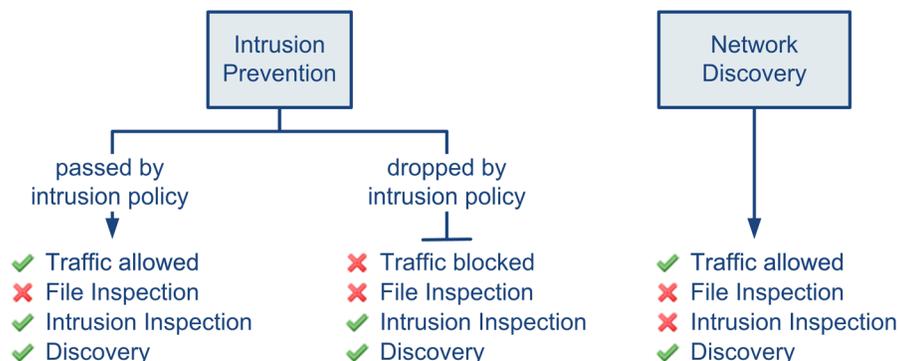
You can also set the default action so that it inspects default-action traffic with network discovery, an intrusion policy, or both. If you are performing neither intrusion detection nor access control, selecting a default action of **Network Discovery Only** can improve Defense Center performance. Note that to take advantage of this performance improvement, you must make sure your access

control rules do not contain: application, user, or URL conditions; or file and intrusion inspection options.

IMPORTANT! Selecting a default action of **Network Discovery Only** does not automatically guarantee discovery inspection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy. For more information, see [Introduction to Network Discovery](#) on page 1250.

If you inspect default-action traffic with an intrusion policy, the system can also inspect it using network discovery, depending on the settings in your network discovery policy. See [Intrusion Policies and Access Control Rules](#) on page 513 for a discussion on associating intrusion policies with access control rules.

The diagram below illustrates the **Intrusion Prevention** and **Network Discovery Only** default actions. Notice that although file inspection is supported in access control rules, you cannot perform file inspection on traffic handled by the default action.



The following procedure explains how to set the default action for an access control policy while editing the policy. See [Editing an Access Control Policy](#) on page 459 for the complete procedure for editing an access control policy.

To set the default action of an access control policy:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon (🔧) next to the access control policy you want to configure.
The policy Edit page appears.
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 allow all traffic and inspect it with network discovery, select **Network Discovery Only**.
- To inspect all traffic with both network discovery and intrusion policies, select an intrusion policy, all of which begin with the label **Intrusion Prevention**. Keep in mind that an intrusion policy can block traffic.

WARNING! Do **not** use **Experimental Policy 1** unless instructed to do so by a Sourcefire representative. Sourcefire uses this policy for testing.

4. Configure logging options for the default action as described in the next section, [Logging Connections for the Default Action](#).

Logging Connections for the Default Action

LICENSE: Any

You must decide whether you want to log connection data for the traffic that is handled by the default action. The options for logging connections handled by the policy default action largely parallel the options for logging connection handled by individual access control rules. However, there are some differences:

- The default action has no file logging options because you cannot perform file control or malware protection 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-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.

For a comprehensive discussion of connection logging, see [Logging Connection, File, and Malware Information](#) on page 515.

In general, if you want to perform any kind of detailed analysis on connection data, you should log the end of connections. If you want to view connection summaries in custom workflows, view connection data in graphical format, or create and use traffic profiles, you **must** log connection events at the end of connections. Note that for the **Block All Traffic** default action, you can log only beginning-of-connection events, because traffic is denied without further inspection.

Logging connection events to the Defense Center database allows you to take advantage of the analysis, reporting, and correlation features in the Sourcefire 3D System. Optionally, you can send most connection events to the syslog or an SNMP trap server.

The following procedure explains how to configure an access control policy to log connections. See [Editing an Access Control Policy](#) on page 459 for the complete procedure for editing an access control policy.

To log connections in traffic handled by the default action:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

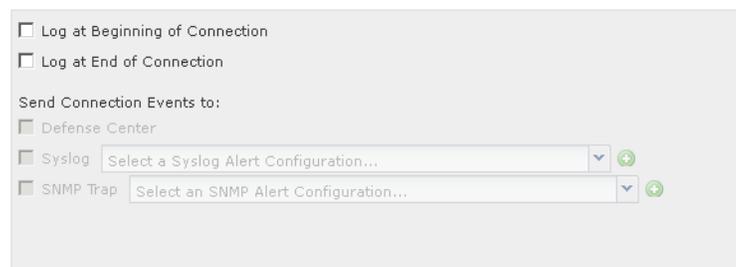
The Access Control page appears.

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

The policy Edit page appears.

3. Click the logging icon (📄) next to the **Default Action** drop-down list.

The Logging pop-up window appears.



4. Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

You cannot log end-of-connection events for blocked traffic.

5. Specify where to send connection events. You have the following choices:

- To send connection events to the Defense Center, select **Defense Center**.
- To send connection events to syslog, 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](#) on page 530.
- 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 configure an SNMP alert response by clicking the add icon (+); see [Creating an SNMP Alert Response](#) on page 528.

6. Save your changes.

You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Using Custom User Roles with Access Control Policies

LICENSE: Any

As described in [Managing Custom User Roles](#) on page 1908, you can create custom user roles with specialized access privileges. Custom user roles can have any set of menu-based and system permissions, and may be completely original or based on a predefined user role. Custom roles for access control-related features determine whether users can view, modify, and apply access control, intrusion, and file policies, as well as insert or modify rules in the Administrator Rules or Root Rules categories.

The following table shows five example custom roles that determine how Sourcefire 3D System users interact with access control features. The table lists, in the order they appear when creating custom user roles, the privileges required for each custom role.

Example Access Control Custom Roles

CUSTOM ROLE PERMISSION	ACCESS CONTROL EDITOR	INTRUSION EDITOR	FILE POLICY EDITOR	POLICY APPLIER (ALL)	INTRUSION POLICY APPLIER
Access Control	X			X	X
Access Control List	X			X	X
Modify Access Control Policy	X				
Apply Intrusion Policies				X	X
Apply Access Control Policies				X	
Intrusion		X			
Intrusion Policy		X			
Modify Intrusion Policy		X			
File Policy			X		
Modify File Policy			X		

Note that the system can render the web interface differently depending on whether a user can apply both access control policies and intrusion policies, only intrusion policies, or neither. For example, the Intrusion Policies Applier in the

table above can view access control policies and apply intrusion policies, but cannot edit access control policies or intrusion policies, cannot apply access control policies, and cannot view file policies. In the web interface:

- the edit icon () does not appear on the Access Control page
- the delete icon () does not appear on the Access Control page
- the quick-apply pop-up window applies only the intrusion policy
- access control policy check boxes in the detailed apply pop-up window are disabled

Managing Policy Targets

LICENSE: Any

Before you can apply an access control policy, you must identify the managed devices, including device groups, where you want to apply the policy. You can identify the managed devices you want to target with your policy while creating or editing a policy. You can search a list of available devices and add devices to a list of selected devices. You can drag and drop selected devices, or add devices using the button between the two lists.

Note that you cannot target stacked devices running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). You can target a device stack, but not individual devices within the stack. See [Managing Stacked Devices](#) on page 237 for more information.

The following table summarizes the actions you can take when managing targeted devices.

Targeted Device Management Actions

To...	You CAN...
search a list of available devices	click inside the search field, then type a search string. The list of devices updates as you type to display matching device names.
clear a search for available devices	click the clear icon () in the search field.
select available devices to add to the list of selected targets	click the device name; use the Ctrl and Shift keys to select multiple devices. TIP! You can also right-click an available device, then click Select All .
add selected devices	click Add to Policy . TIP! You can also drag and drop into the list of selected devices.

Targeted Device Management Actions (Continued)

To...	You CAN...
delete a single device from the Selected Devices list	click the delete icon () next to the device. TIP! You can also right-click the device and select Delete .
delete multiple devices from the Selected Devices list	use the Ctrl and Shift keys to select multiple devices, right-click to highlight the row for a selected device, then click Delete Selected .
save your configuration	click OK .
discard your configuration without saving your changes	click Cancel .

The following procedure explains how to configure an access control policy to manage targeted devices. See [Editing an Access Control Policy](#) on page 459 for the complete procedure for editing an access control policy.

To manage targeted devices in an access control policy:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon () next to the access control policy you want to configure.
The policy Edit page appears.
3. Click the device targets link, then click **Manage Targets**.
The Manage Device Targets pop-up window appears.
4. Optionally, click the **Search** prompt above the **Available Devices** list, then type a name.
The list updates as you type to display matching devices. You can click the clear icon () to clear the list.
5. Click the device or device group you want to add. Use Ctrl and Shift to select multiple devices.

TIP! You can also right-click an available device, then click **Select All**.

6. Click **Add to Policy**.
Selected devices are added.

TIP! You can also drag and drop.

7. Optionally, click the delete icon () to delete a device from the list of selected devices; or, use the Ctrl and Shift keys to select multiple devices, right-click, then select **Delete Selected**.
8. Click **OK** to save your configuration, or click **Cancel** to discard it.
If you click **OK**, your configuration is added to the policy and the policy Edit page appears.

Adding an HTTP Response Page

LICENSE: FireSIGHT

When an access control rule blocks a user's HTTP request, what the user sees in a web browser depends on how you block the session. When choosing a rule action, select:

- **Block** or **Block with reset** if you want 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 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. Note that users **must** enable JavaScript and third-party cookies in their browsers before they can dismiss these warnings and click to continue. Users may have to refresh after bypassing the response page to load page elements that did not load.

You can either display a generic Sourcefire-provided response page, or you can enter custom HTML. When you enter custom text, a counter shows how many characters you have used. For a blocked session custom page you can use up to 1353 characters; click-to-continue custom pages are limited to 1273 characters.

Note that HTTP response pages do not appear for traffic blocked because of a Security Intelligence blacklist or an application detected based on a Secure Sockets Layer (SSL) certificate.

TIP! To quickly disable interactive blocking for all rules in an access control policy, display neither the Sourcefire-provided page nor a custom page.

In each access control policy, you configure the response page for blocking and interactively blocking rules separately. The rule action determines the page displayed to your users. For example, you could display the Sourcefire-provided page to users whose sessions are blocked, but a custom page to users who can click to continue. For detailed information on rule actions, see [Understanding Rule Actions](#) on page 478.

Note that in rare cases, when a blocking rule is preceded by another rule that contains an application condition, an HTTP response page may not appear even though the system successfully blocks traffic.

To configure HTTP response pages:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

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

The policy Edit page appears.

3. Select the **HTTP Responses** tab.

HTTP response page settings for the access control policy appear.



Block Response Page
This page will be displayed when HTTP traffic is blocked.
None

Interactive Block Response Page
This page will be displayed when HTTP traffic is blocked, but the user may choose to continue.
Sourcefire-provided 🔍

4. For the **Block Response Page** and the **Interactive Block Response Page**, select responses from the drop-down lists. For each page, you have the following choices:

- To use a generic response, select **Sourcefire-provided**. You can click the view icon (🔍) to view the HTML code for this page.
- To create a custom response, select **Custom**.
A pop-up window appears, pre-populated with Sourcefire-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 (✎).
- To prevent the system from displaying an HTTP response page, select **None**. Note that selecting this option for interactively blocked sessions prevents users from clicking to continue.

5. Click **Save** to save your configuration.

You must apply the access control policy for your changes to take effect. For more information, see [Applying an Access Control Policy](#) on page 466.

Filtering Traffic Based on Security Intelligence Data

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

The Security Intelligence feature allows you to specify the traffic that can traverse your network, per access control policy, 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.

Note that you could create access control rules that perform a similar function to Security Intelligence filtering. However, access control rules are wider in scope, more complex to configure, and cannot automatically update using dynamic feeds. In contrast, Security Intelligence filtering can immediately blacklist connections based on the latest intelligence, removing the need for a more resource-intensive, in-depth analysis.

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.

To help you build blacklists, Sourcefire provides the Sourcefire Intelligence Feed, which is comprised of several regularly updated collections of IP addresses determined by the VRT to have a poor reputation. To augment the intelligence feed, you can use third-party feeds and custom lists of IP addresses, including a global blacklist. You can also blacklist IP addresses using network objects and groups. These configurations are collectively called *Security Intelligence objects*.

IMPORTANT! 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 an access control policy reapply. For more information, see the [Security Intelligence Object Capabilities table](#) on page 165.

Choosing IP Addresses to Blacklist

The easiest way to construct a blacklist is to use the Sourcefire 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:

- a *list* is a static list of IP addresses that you upload to the Defense Center
- a *feed* is a dynamic list of IP addresses that the Defense Center downloads from the internet on a regular basis; the Sourcefire Intelligence Feed is a special kind of feed

For detailed information on configuring Security Intelligence lists and feeds, including high availability and internet access requirements, see [Working with Security Intelligence Lists and Feeds](#) on page 162.

Also, in the course of your analysis, you can build a *global blacklist* by selecting any IP address in an event view, the Context Explorer, or a dashboard. For example, if you notice a set of routable IP addresses in intrusion events associated with exploit attempts, you can immediately blacklist those IP addresses. The Defense Center 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](#) on page 166.

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](#) on page 161.

IMPORTANT! Although they have all other Protection capabilities by default, Series 2 devices cannot perform Security Intelligence filtering. You cannot apply an access control policy that uses a populated global whitelist or blacklist to Series 2 devices (or to unlicensed Series 3 devices). 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.

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 might want to 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. 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, Sourcefire recommends that you always use monitor-only settings. This is because managed devices that are deployed passively cannot affect traffic flow; there is no advantage to configuring the system to block traffic.

Logging Blacklisted Connections

Logging blacklisted connections allows you to generate a connection event when the system detects network traffic to or from a blacklisted IP address. Events generated by Security Intelligence filtering represent the decision made by the system to either deny (blacklist) or inspect (blacklist set to monitor-only) the connection. This logging configuration is independent of the logging configurations for access control rules or the default action.

You must enable logging for Security Intelligence if you want to set blacklisted objects to monitor-only. Note that for those matching connections that go on to be inspected by access control rules, the system may generate additional connection events, depending on the logging settings in the access control rule or default action that later handles the connection.

Health Monitoring

The default health policy includes the Security Intelligence module (see [Configuring Security Intelligence Monitoring](#) on page 2139), which warns you if:

- the Defense Center cannot update a feed, or if feed data is corrupt or contains no recognizable IP addresses
- a managed device had a problem receiving updated Security Intelligence data from the Defense Center
- a managed device cannot load all of the Security Intelligence data provided to it by the Defense Center, due to memory issues

For detailed information on configuring your access control policy to perform Security Intelligence filtering, see the following sections:

- [Building the Security Intelligence Whitelist and Blacklist](#) on page 441
- [Searching for Objects to Whitelist or Blacklist](#) on page 444
- [Creating Objects to Whitelist or Blacklist](#) on page 444
- [Logging Blacklisted Connections](#) on page 445

Building the Security Intelligence Whitelist and Blacklist

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

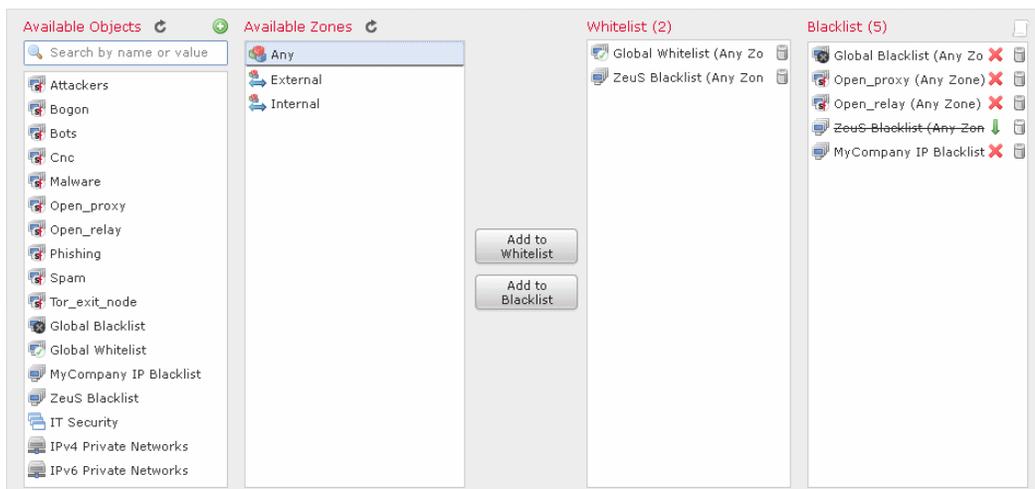
SUPPORTED DEFENSE CENTERS: Any except DC500

To build a whitelist and blacklist, 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 Defense Center's global whitelist and blacklist, which apply to any zone. These lists are populated by your analysts, who can quickly add individual IP addresses using the context menu. You can opt not to use these global lists on a per-policy basis. For more information, see [Working with the Global Whitelist and Blacklist](#) on page 166.

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 Sourcefire icon () represent the different categories in the Sourcefire 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 ranges using a /0 netmask in those objects will be ignored and whitelist and blacklist filtering will not occur based on those addresses. Address ranges with a /0 netmask from security intelligence feeds will also be 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.

TIP! The general mechanics of constructing Security Intelligence whitelists and blacklists are the same as those for constructing access control rules. For detailed information, see [Understanding Rule Conditions and Condition Mechanics](#) on page 482.

To build the Security Intelligence whitelist and blacklist for an access control policy:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon () next to the access control policy you want to configure.
The policy Edit page appears.
3. Select the **Security Intelligence** tab.
Security Intelligence settings for the access control policy appear.
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 Blacklisted Connections](#) on page 445.
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](#) on page 444 and [Creating Objects to Whitelist or Blacklist](#) on page 444.

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.
7. Click **Add to Whitelist** or **Add to Blacklist**.
You can also click and drag the selected objects to either list.

TIP! To remove an object, 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 the object from the Defense Center.

8. Repeat steps 5 through 7 until you are finished adding objects to your whitelist and blacklist.

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, Sourcefire recommends you set all blacklisted objects to monitor-only. Note, however, that you cannot set the global blacklist to monitor-only.

10. Click **Save**.

You must apply the access control policy for your changes to take effect. For more information, see [Applying an Access Control Policy](#) on page 466.

Searching for Objects to Whitelist or Blacklist

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

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:

ACCESS: Admin/Access Admin/Network Admin

- ▶ Type in the **Search by name or value** field.

The Available Objects list 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.

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**.

Creating Objects to Whitelist or Blacklist

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

While editing an access control policy, you can create an object on-the-fly to use in its whitelist and blacklist: either a network object or a Security Intelligence list or feed. Note that to group network objects or create network object groups, you must use the object manager.

To create objects to whitelist or blacklist:

ACCESS: Admin/Access Admin/Network Admin

- ▶ Click the add icon (), then select the type of object you want to create:
 - Select **Add IP List** to create a Security Intelligence list or feed; see [Working with Security Intelligence Lists and Feeds](#) on page 162.
 - Select **Add Network Object** to add a network object; see [Working with Network Objects](#) on page 161.

Logging Blacklisted Connections

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

Logging blacklisted connections allows you to generate a connection event when the system detects network traffic to or from a blacklisted IP address. You can save these connection events to the Defense Center database, and can also log the events to the syslog or to an SNMP trap server using alert responses. For information on setting up alert responses, see [Working with Alert Responses](#) on page 526.

IMPORTANT! You **must** send events to the Defense Center if you want to set blacklisted objects to monitor-only, or perform any other Defense Center-based analysis on connection events generated by Security Intelligence filtering.

Unlike the logging options for access control rules or the default action, you cannot choose whether to generate beginning- or end-of-connection events. Events generated by Security Intelligence filtering always represent the beginning of a connection and the decision made by the system to either:

- deny the traffic without further inspection (blacklist)
- perform further analysis on the connection (blacklist set to monitor-only)

This decision is logged as a connection event's reason: either **IP Block** or **IP Monitor**. The decision is also reflected in the connection event's action, which for a blacklisted connection is **Block**. Contrast with a monitored connection, where the action is that of the first non-Monitor access control rule triggered by the connection, or the default action.

The system also logs a Security Intelligence category, which qualifies the reason with the name of the blacklisted object — whether network object, feed, or list — that contains the blacklisted IP address.

In the event viewer, so that you can identify the blacklisted IP address in the connection, host icons next to blacklisted and monitored IP addresses look slightly different.

Because the decision to blacklist a connection occurs before the network traffic is evaluated by access control rules, connection events generated by Security

Intelligence filtering do not contain information that must be determined by examining traffic over the duration of the session, nor do they contain application data. For details on the information in connection events, see [Information Available in Connection Events](#) on page 550.

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.

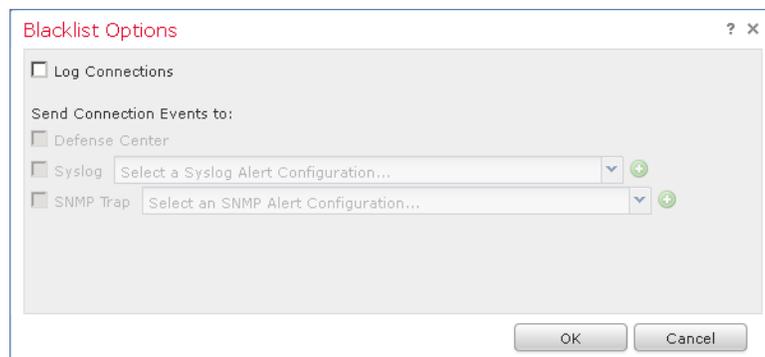
Note that the system may generate additional events for monitored connections, depending on the logging settings in the access control rule or default action that later handles the connection. For similar reasons, the system does not generate a special connection event when it detects a connection to or from a whitelisted IP address. That is, whitelisted connections generate events depending on how the system later handles the connection.

To log blacklisted connections:

ACCESS: Admin/Access Admin/Network Admin

1. On the Security Intelligence tab in an access control policy, click the logging icon ().

The Blacklist Options dialog box appears.



2. Select the **Log Connections** check box to log beginning-of-connection events when traffic meets Security Intelligence conditions.

3. Specify where to send connection events. You have the following choices:
 - To send connection events to the Defense Center, select **Defense Center**.
 - To send connection events to 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](#) on page 530.
 - 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 responses by clicking the add icon (+); see [Creating an SNMP Alert Response](#) on page 528.
4. Click **OK** to set your logging options.
The Security Intelligence tab appears again.
5. Click **Save**.
You must apply the access control policy for your changes to take effect. For more information, see [Applying an Access Control Policy](#) on page 466.

Configuring Advanced Access Control Policy Settings

LICENSE: Any

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments.

URL Logging Advanced Options

When you log an end-of-connection event to the Defense Center database (see [Logging Connection, File, and Malware Information](#) on page 515) 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. Using **Maximum URL characters to store in connection events**, 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.

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 [Adding URL Conditions](#) on page 506.

File and Malware Detection Options

If you use file policies to perform file control and malware detection or blocking, you can:

- using **Limit the number of bytes inspected when doing file type detection**, specify the number of bytes inspected when performing file type detection. By default, this option is set to 1460 bytes, or one packet. In most cases, the system can identify common file types using the first packet.
- using **Do not perform malware detection on files larger than (in bytes)**, prevent the system from performing malware detection on files larger than a certain size. By default, the option is set to 10485760 (10MB)
- using **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 set duration elapses without the system obtaining a disposition, the system allows the file to pass through into traffic. By default, the option is set to 2 seconds.

Note that dispositions of Unavailable are not cached.

WARNING! Although this option will accept values of up to 30 seconds, Sourcefire recommends that you use the default value to avoid blocking traffic because of connection failures. Do **not** set a value of 0 for this option without first contacting Sourcefire Support.

You can set either of the first two options as high as 4294967295 (4GB). Setting the first two options to zero removes the restriction altogether. However, keep in mind that increasing the file sizes can affect the performance of the system.

Interactive Block Options

When traffic matches access control rules with **Interactive Block** or **Interactive Block with Reset** as the action, the user can click through a response page to bypass the block. Using **Allow an Interactive Block to bypass blocking for (seconds)**, you can set how long the system allows a user to resume a session using the same URL from the same IP address without clicking through the response page again. The default setting is 600 seconds (equivalent to 10 minutes). You can set the duration to as long as 31536000 seconds (equivalent to 365 days). Set this option to zero to make any user bypass of the response page for an interactive block a permanent bypass for sessions with the same URL and user's IP address.

To configure advanced access control policy options:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

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

The policy Edit page appears.

3. Select the **Advanced** tab.

Advanced settings for the access control policy appear.

Maximum URL characters to store in connection events	1024
Limit the number of bytes inspected when doing file type detection	1460
Do not perform malware detection on files larger than (in bytes)	10485760
Allow file if cloud lookup for Block Malware takes longer than (seconds)	2
Allow an Interactive Block to bypass blocking for (seconds)	600

4. Configure the advanced options, as described above.

5. Click **Save**.

You must apply the access control policy for your changes to take effect. For more information, see [Applying an Access Control Policy](#) on page 466.

Organizing Rules in a Policy

LICENSE: Any

The Edit page for the access control policy lists access control rules in numerical order. The numeric position of each rule appears on the left side of the page next to the rule. You can move or insert rules and otherwise change the rule order. For example, if you move rule 10 under rule 3, rule 10 becomes rule 4 and all subsequent numbers increment accordingly.

A rule's position is important because the system compares packets to rules in the numeric order in which the rules are arranged on the policy Edit page. When a packet meets all the conditions of a rule, the system applies the conditions of that rule to the packet and ignores all subsequent rules for that packet.

Optionally, you can specify a rule's numeric position when you add or edit a rule. You can also highlight a rule before adding a new rule to predetermine the default position of the new rule to be below the rule you highlighted. See [Creating and Editing Access Control Rules](#) on page 474.

To locate specific rules, you can use partial or complete strings to search for rules by rule name or by a name or value in configured rule conditions. You can also filter rules to display only rules for selected devices targeted by your policy.

You can select one or more rules by clicking a blank space in the row for the rule. You can drag and drop selected rules into a new location, thereby changing the position of the rules you moved and all subsequent rules. You can cut or copy

selected rules and paste them above or below an existing rule. You can delete selected rules and insert new rules into any location in the list of existing rules.

You can further organize rules by adding custom categories between the administrative and root categories. You can delete or rename custom categories that you add.

You can display explanatory warnings to identify rules that will never match because they are preempted by preceding rules.

The following table summarizes the actions you can take to organize your rules.

Access Control Rule Organization Actions

To...	You CAN...
add a category to a policy	click Add Category . See Working with Rule Categories on page 451 for more information. TIP! You can also right-click a blank area in the row for a rule and select Insert new category .
search rule names and conditions for a string	click the Search Rules prompt, type a name or value, then press the Enter key. See Searching for Rules on page 453 for more information.
clear rule search	click the clear icon (✕) in the search field.
display rules for selected devices	find more information at Filtering Rules by Device on page 454.
select a rule	click a blank area in the row for a rule. Use the Ctrl or Shift key to select multiple rules. Rules you select are highlighted. Note that you can select rules in multiple categories.
clear rule selections	click the reload icon (↻) on the lower right side of the page.
cut or copy selected rules	right-click a blank area in the row for a selected rule, then select Cut or Copy .
paste rules you have cut or copied into the rule list	right-click a blank area in the row for a rule where you want to paste selected rules, then select Paste above or Paste below .
enable an inactive rule	right-click the rule and select State > Enable .
disable an active rule	right-click the rule and select State > Disable .
move selected rules	drag and drop selected rules beneath a new location indicated by a horizontal blue line that appears above your pointer as you drag.

Access Control Rule Organization Actions (Continued)

To...	You CAN...
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.
read warnings or errors	hover over the warning icon () or error icon () to read the warning or error text; see Working with Warnings and Errors on page 455 for more information.
determine if an intrusion policy or file policy is selected for a rule	view the intrusion policy icon () or the file policy icon (). If the icon for a policy is active (yellow) a policy is selected; if it is inactive (white), no policy of that type is selected for the rule.
view the intrusion policy or file policy selected for a rule	click the intrusion policy icon () or the file policy icon ().

Working with Rule Categories

LICENSE: Any

The following three predefined access control rule categories on the policy Edit page can help you organize your rules:

- Administrator Rules
- Standard Rules
- Root Rules

You cannot move, delete, or rename predefined categories. By default, any predefined user role that allows you to modify access control policies also allows you to move rules into and from, and modify rules in, any of these categories. You can create custom user roles that restrict users from moving and modifying rules in these predefined categories. See [Managing Custom User Roles](#) on page 1908 and [Policies Menu](#) on page 1920 for more information.

You can add new custom categories between the predefined standard and root categories. 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. Any user who is allowed to modify access control policies can also add rules to these categories and modify rules in them without restriction.

The following procedure explains how to add a new category to an access control policy. See [Editing an Access Control Policy](#) on page 459 for the complete procedure for editing an access control policy.

To add a new category:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

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

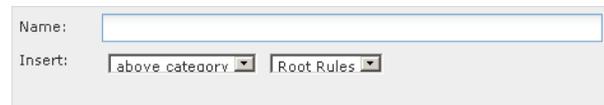
The policy Edit page appears.

3. Optionally, click a blank area in the row for an existing rule to set the default position of the new category.

4. Click **Add Category**.

Alternately, if you have added rules to your policy, you can right-click an existing rule and then click **Insert new category**.

The Add Category pop-up window appears.



5. Type a unique category **Name**.

You can enter an alphanumeric name, including spaces and special printable characters, with up to 30 characters.

6. 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.
Note that 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.
Note that this option is valid only when at least one rule exists in the policy.

7. Click **OK** to add your category, or click **Cancel** to discard it.

If you click **OK**, your category is added to the policy.

Note that you can click the edit icon () next to a category you add to edit the category name, or click the delete icon () to delete the category. Rules in a category you delete are added to the category above.

Searching for 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.

The following procedure explains how to search for rules in an access control policy. See [Editing an Access Control Policy](#) on page 459 for the complete procedure for editing an access control policy.

To search for rules:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon () next to the access control policy you want to search.
The policy Edit page appears.
3. Click the **Search Rules** prompt, type a search string, then press Enter.
Columns for rules with matching values are highlighted, with differentiated highlighting for the indicated (first) match.

TIP! You can also use the Tab key or click a blank page area to initiate the search.

4. You have the following options:
 - To navigate between matching rules, click the next-match (▼) or previous-match (▲) icon.
 - To clear the search string, click the clear icon (✕).
The page refreshes and highlighting clears.

Filtering Rules by Device

LICENSE: Any

You can filter the access control rules listed in your access control policy to display only the rules for one or more specified devices or device groups. The system uses the zone conditions in access control rules to associate rules with devices on your network. See [Working with Security Zones](#) on page 187 and [Adding Zone Conditions](#) on page 492 for more information.

Rules are hidden for devices and groups that you do not specify. Rules where you do not add zones are targeted for any zone, and therefore are targeted for all devices, so they are never hidden.

The following procedure explains how to filter rules by device or device group. See [Editing an Access Control Policy](#) on page 459 for the complete procedure for editing an access control policy.

To filter rules by device or device group:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon (✎) next to the access control policy you want to modify.
The policy Edit page appears.
3. Click **Filter by Device** above the list of rules.
The Filter by Device pop-up window appears. If you have added devices or device groups to your policy, a list of targeted devices and device groups appears.
4. Select one or more of the check boxes to display only the rules that apply to those devices or groups. Alternatively, select the **All** check box to reset and display all of the rules.
5. Click the **OK** button to update the list of rules.
The page updates to display rules for devices and device groups that you selected and hide rules for devices and device groups that you did not select.

TIP! Filters are cleared if you add a new rule, or if you edit and save an existing rule.

Working with Warnings and Errors

LICENSE: Any

Because of the number of configurable elements in an access control policy, policies can be very complex. Rules might be preempted by other rules. Functionality might be configured that depends on configuration outside of the access control policy. To help ensure that the policy you configure has the result you expect, the access control policy interface has a robust warning and error feedback system. If a rule or other element within a policy has a warning, the policy can be applied, but that piece of configuration will have no effect. If an element has an error, policy apply will fail unless the erroneous configuration is corrected.

To view the warning text for an object in the policy, hover your pointer over the warning icon (⚠) next to it. To view the error text for an object, hover your pointer over the error icon (❗) next to it.

If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue. If you disable a rule with a warning, note that the error icon remains and the policy will still not apply even with the rule disabled.

Understanding Invalid Configurations

Because outside settings that the access control policy depends on may change, an access control policy setting that was valid may become invalid.

For example, if you have a URL condition in a rule, the rule might be valid until you choose to target a device 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.

In another example, 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 not actually be applied to targeted devices.

Similarly, if you add a user to a rule, then change your LDAP user awareness settings to exclude that user, that rule will no longer apply because the user is no longer an access controlled user.

For all of these situations, warnings or errors appear in the access control policy or access control policy list to alert you to the issues.

Understanding Rule Pre-emption

The conditions of an access control rule may preempt a subsequent rule from matching traffic. For example:

```
Rule 1: allow Administrator users  
Rule 2: block Administrator users
```

The second rule above will never block traffic because the first rule will have already allowed the traffic.

Any type of rule condition can preempt a subsequent rule. For example, the VLAN range in the first rule below includes the VLAN in the second rule, so the first rule preempts the second rule:

```
Rule 1: allow VLAN 22-33
Rule 2: block VLAN 27
```

In the following example, Rule 1 matches any VLAN because no VLANs are configured, so Rule 1 preempts Rule 2, which attempts to match VLAN 2:

```
Rule 1: allow Source Network 10.4.0.0/16
Rule 2: allow Source Network 10.4.0.0/16, VLAN 2
```

A rule also preempts an identical subsequent rule where all configured conditions are the same. For example:

```
Rule 1: allow VLAN 1 URL www.example.com
Rule 2: allow VLAN 1 URL www.example.com
```

A subsequent rule would not be preempted if any condition is different. For example:

```
Rule 1: allow VLAN 1 URL www.example.com
Rule 2: allow VLAN 2 URL www.example.com
```

Managing Access Control Policies

LICENSE: Any

On the Access Control policy page (**Policies > Access Control**) you can view all your current access control policies by name with optional description and the following status information:

- when a policy is up to date on targeted devices, in green text
- when a policy is out of date on targeted devices, in red text

Options on this page allow you to compare policies, create a new policy, apply a policy to targeted devices, copy a policy, view a report that lists all of the most recently saved settings in each policy, and edit, or delete a policy.

TIP! You can export access control policies to, and import access control policies from, other Defense Centers in your deployment. See [Importing and Exporting Objects](#) on page 2225 for more information.

Depending on your choices when you add a device, either of two default access control policies might appear and already be applied to the device:

- The Default Access Control policy blocks all traffic from entering your network.
- The Default Intrusion Prevention policy allows all traffic and applies the Balanced Security and Connectivity intrusion policy to traffic on your network; see [Configuring Intrusion Policies](#) on page 667.

You can use either of these policies the same as you use policies you create.

The following table describes the actions you can take to manage your policies on the Access Control policy page:

Access Control Policy Management Actions

To...	You CAN...
create a new access control policy	click Create Policy . See Creating an Access Control Policy on page 458 for more information.
modify the settings in an existing access control policy	click the edit icon (). See Editing an Access Control Policy on page 459 for more information.
apply an access control policy to all devices targeted for the policy	click the policy apply icon (). See Applying an Access Control Policy on page 466 for more information.
determine what changed in policies to make them out of date on a device	click the red status message to see the detailed apply view, then click Out-of-date for the policy and device where you want to see what changed. See Applying Selected Policy Configurations on page 469 and Comparing Two Access Control Policies on page 462 for more information.
copy an access control policy	click the copy icon (). See Copying an Access Control Policy on page 460 for more information.
view a PDF report that lists the current configuration settings in an access control policy	click the report icon (). See Viewing an Access Control Policy Report on page 461 for more information.
compare access control policies	click Compare Policies . See Comparing Two Access Control Policies on page 462 for more information.
delete an access control policy	click the delete icon () , then click OK , or click Cancel if you decide not to delete the policy. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy.

Creating an Access Control Policy

LICENSE: Any

When you create a new access control policy you must, at minimum, give it a unique name and specify a default action. Although you are not required to identify the policy targets at policy creation time, you must perform this step before you can apply the policy; see [Managing Policy Targets](#) on page 434.

You have the following options when selecting a default action for a new policy:

- **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.
- **Network Discovery** creates a policy with the **Network Discovery Only** default action.

After you create the access control policy, you can modify the default action. For guidance on choosing a default action, see [Setting the Default Action](#) on page 428.

To create an access control policy:

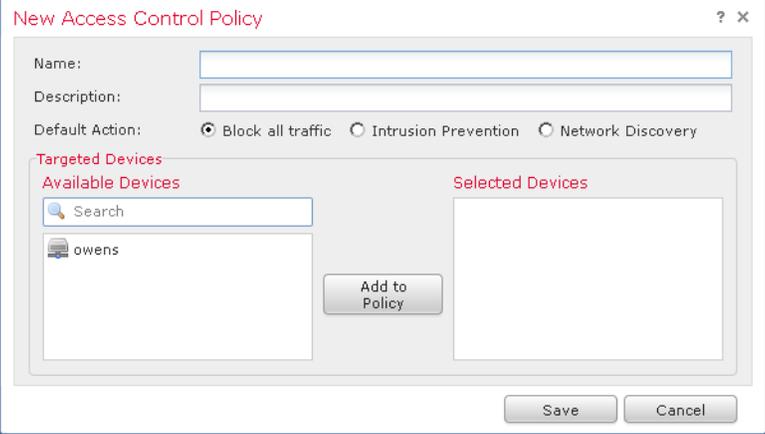
ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

2. Click **New Policy**.

The New Access Control Policy pop-up window appears.



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.

4. Specify the **Default Action**.

5. Select the **Available Devices** where you want to apply the policy.
Use Ctrl and Shift to select multiple devices, or right-click to **Select All**. To narrow the devices that appear, type a search string in the **Search** field. To clear the search, click the clear icon (✕).
6. Add the **Selected Devices**. You can click and drag, or you can click **Add to Policy**.
7. Click **Save**.
The access control policy Edit page appears. For information on configuring your new policy, including adding rules, see [Editing an Access Control Policy](#) on page 459. Note that you must apply the policy for it to take effect; see [Applying an Access Control Policy](#) on page 466.

Editing an Access Control Policy

LICENSE: Any

On the policy Edit page, you can configure your policy and organize access control rules. See [Configuring Policies](#) on page 426 and [Organizing Rules in a Policy](#) on page 449 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 Edit page. If you attempt to exit the policy Edit page 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 Edit page.

To protect the privacy of your session, after sixty minutes of inactivity on the policy Edit page, changes to your policy are discarded and you are returned to the Access Control 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 you attempt to edit the same policy in two browser windows, you are prompted whether to resume your edit in the new window, discard your changes in the original window and continue editing in the new window, or cancel the second window and return to the policy Edit page.

When multiple users edit the same policy concurrently, a message for each on the policy Edit page identifies other users who have unsaved changes. Any user who attempts to save their changes is cautioned that their changes will overwrite changes by other users. When the same policy is saved by multiple users, the last saved changes are retained.

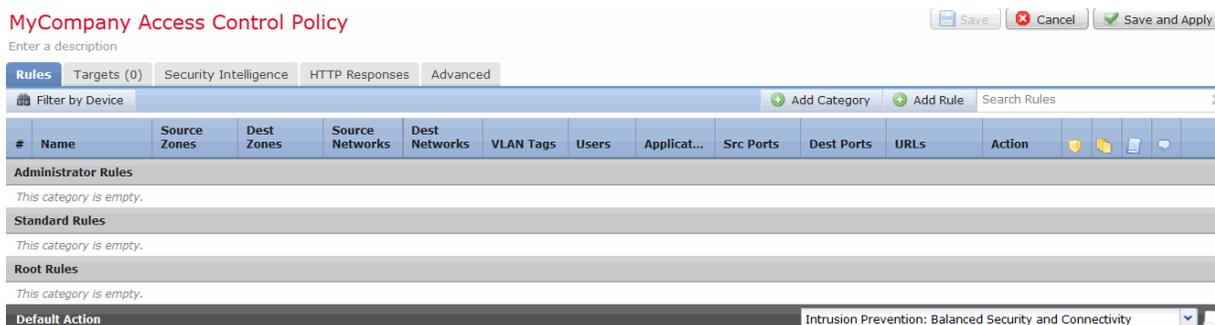
To edit an access control policy:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.

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

The policy Edit page appears.



3. You have the following choices:
 - To configure your policy, you can take any of the actions described in [Configuring Policies](#) on page 426 and summarized in the [Access Control Policy Configuration Actions](#) on page 427.
 - To organize rules in your policy, you can take any of the actions described in [Organizing Rules in a Policy](#) on page 449 and summarized in the [Access Control Rule Organization Actions](#) on page 450.
4. Save or discard your configuration. You have the following choices:
 - To save your changes and continue editing, click **Save**.
 - To save your changes and apply your policy, click **Save and Apply**. See [Applying an Access Control Policy](#) on page 466.
You must apply your policy to put your changes into effect.
 - To discard your changes, click **Cancel** and, if prompted, click **OK**.
Your changes are discarded and the Access Control page appears.

Copying an Access Control Policy

LICENSE: Any

You can copy and rename an access control policy. A policy you copy includes all policy rules and configurations.

To copy an access control policy:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.

2. Click the copy icon (📄) next to the access control policy you want to configure.

The Copy Access Control Policy pop-up window appears.



3. Enter a unique policy **Name**.
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.
4. Click **OK**.
Your copy appears on the Access Control page in alphabetical order by name.

Viewing an Access Control Policy Report

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 for auditing purposes or to inspect the current configuration.

TIP! 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 Two Access Control Policies](#) on page 462.

An access control policy report contains the sections described in the following table.

Access Control 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. See Editing an Access Control Policy on page 459.

Access Control Policy Report Sections (Continued)

SECTION	DESCRIPTION
Device Targets	Lists the managed devices targeted by the policy. See Managing Policy Targets on page 434.
HTTP Block Response HTTP Interactive Block Response	Provides details on the HTTP block response pages associated with the policy. See Adding an HTTP Response Page on page 436.
Security Intelligence	Provides details on the Security Intelligence whitelist and blacklist. See Filtering Traffic Based on Security Intelligence Data on page 438.
Default Action	Provides the default action. See Setting the Default Action on page 428.
Rules	Provides the rule action and conditions for each rule in the policy, by rule category. See Understanding and Writing Access Control Rules on page 472 and Working with Rule Categories on page 451.
Referenced Objects	Provides the name and configuration of all individual objects and group objects used in the policy, by type of condition (Networks, VLAN Tags, and so on) where the object is configured. See Understanding Rule Conditions on page 483 and Using Objects and Security Zones on page 158.

To view an access control policy report:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

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. 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 Two 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 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 Access Control Policy Comparison View](#) on page 463
- [Using the Access Control Policy Comparison Report](#) on page 465

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.

Default Intrusion Prevention (2012-05-17 10:36:37 by admin)	MyCompany's AC Policy (2012-05-24 16:26:06 by admin)
Policy Information	
Name	Default Intrusion Prevention
Description	
Modified	2012-05-17 10:36:37 by admin
HTTP Block Response	None
Advanced Settings	
Maximum URL Length	1024
Rules	
Category 1	
Name	Administrator Rules
Rule 1	
Name	monitor non-IE browsers
Action	monitor
Application Filters	
Chrome	
Firefox	
Konqueror	
Mobile Safari	

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.

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 on page 465 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 the [Access Control Policy Report Sections table](#) on page 461.

TIP! You can use a similar procedure to compare intrusion, file, system, or health policies.

To compare two access control policies:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

2. Click **Compare Policies**.

The Select Comparison window appears.



The screenshot shows a window titled "Select Comparison" with three dropdown menus. The first menu, "Compare Against", is set to "Running Configuration". The second menu, "Target/Running Configuration A", is set to "linden". The third menu, "Policy B", is set to "Default Intru".

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.

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.
5. Click **OK** to display the policy comparison view.
The comparison view appears.
6. Optionally, click **Comparison Report** to generate the access control policy comparison report.
The access control 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.

Applying an Access Control Policy

LICENSE: Any

After making any changes to an access control policy, you must apply the policy to one or more devices to implement the configuration changes on the networks monitored by the devices. You must target devices where you want to apply the policy before you can apply the policy. See [Managing Policy Targets](#) on page 434.

Keep the following points in mind when applying access control policies:

- In special cases, 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. This occurs when the Snort® process restarts; for example, the process restarts when you apply an access control policy that pushes a new version of Snort to a managed device following a Defense Center upgrade, when you apply a policy for the first time after a rule import that includes shared object rules, and, in some cases, when you install a VDB update. See [Performing Software Updates](#) on page 2056, [Updating the Vulnerability Database](#) on page 2070, and [Importing Rule Updates and Local Rule Files](#) on page 2072.
- On 3D7010, 3D7020, and 3D7030 managed devices, applying an access control policy takes up to five minutes. To minimize inconvenience, apply access control policies during a change window.
- If an access control policy requires licenses enabled through recently applied device configurations, the system queues the access control policy apply until the device configurations finish applying.
- Intrusion rules that are set to Drop and Generate Events in an associated intrusion policy where **Drop when Inline** is selected will generate events but will **not** drop any packets or block any attacks when you apply the intrusion policy to a device that uses a passive interface set or an inline interface set in tap mode. See [Setting Drop Behavior in an Inline Deployment](#) on

page 689 and [Tap Mode](#) on page 275 for more information.

- You cannot apply an access control policy to stacked devices running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). See [Managing Stacked Devices](#) on page 237 for more information.
- Some features require minimum versions of the Sourcefire 3D System, or specific device models. Managed devices must be running at least Version 5.1.1 to detect malware and perform file control, and Version 5.2 to block malware. See [Supported Capabilities by Appliance Model](#) on page 42 for a summary of features not supported on Series 2 appliances.
- The label for the apply button on the quick-apply pop-up window can differ depending on whether you are permitted to apply an access control policy, intrusion policy, or both; see [Using Custom User Roles with Access Control Policies](#) on page 433.
- At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all Sourcefire-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. See [Working with Application Conditions](#) on page 498 for more information.
- You can add an unlimited number of unique intrusion policies to an access control policy. However, when you apply the access control policy to a device, a pop-up window may warn that you have exceeded the maximum number of intrusion policies supported by the device. This maximum depends on a number of factors, including the physical memory and the number of processors on your device.

TIP! If you exceed the number of intrusion policies supported by your device, reevaluate your access control policy. You may want to consolidate intrusion policies so you can associate a single intrusion policy with multiple access control rules.

- Although you can apply any combination of an access control policy and its associated intrusion policies, applying an access control policy automatically applies all associated file policies. You cannot apply file policies independently.

See the following sections for more information:

- [Applying a Complete Policy](#) on page 468 explains how to use the quick-apply option to apply the access control policy along with any associated intrusion and file policies.
- [Applying Selected Policy Configurations](#) on page 469 explains how to select and apply any combination of the access control policy, any associated intrusion policies, or both.

Applying a Complete Policy

LICENSE: Any

You can apply an access control policy at any time. Applying an access control policy also applies any associated intrusion and file policies that are different from those currently running on devices targeted by the policy. A pop-up window allows you to apply all together as a single quick-apply action. Unchanged intrusion and file policies are not applied when you use the quick-apply option.

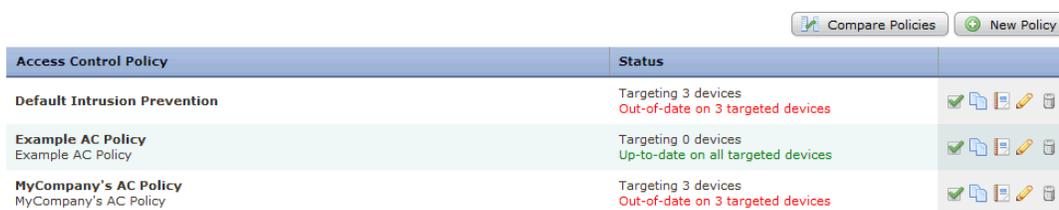
The label for the apply button on the quick-apply pop-up window can differ depending on whether you are permitted to apply an access control policy, intrusion policy, or both; see [Using Custom User Roles with Access Control Policies](#) on page 433.

To quick-apply a complete access control policy:

ACCESS: Admin/Security Approver

1. Select **Policies > Access Control**.

The Access Control page appears.



Access Control Policy	Status	
Default Intrusion Prevention	Targeting 3 devices Out-of-date on 3 targeted devices	   
Example AC Policy Example AC Policy	Targeting 0 devices Up-to-date on all targeted devices	   
MyCompany's AC Policy MyCompany's AC Policy	Targeting 3 devices Out-of-date on 3 targeted devices	   

2. Click the apply icon () next to the policy you want to apply.
The Apply Access Control Rules pop-up window appears.



Alternatively, you can click **Save and Apply** on the policy Edit page; see [Editing an Access Control Policy](#) on page 459.

3. Click **Apply All**.

Your policy apply task is queued. Click **OK** to return to the Access Control page.

TIP! You can monitor the progress of the policy apply task on the Task Status page (**System > Monitoring > Task Status**).

Applying Selected Policy Configurations

LICENSE: Any

You can use the detailed policy apply page to apply changes to your access control policy and to any associated intrusion policies. The detailed page lists each device targeted by the policy and provides a column for the access control policy by device, and a column for associated intrusion policies by device. You can specify whether to apply changes to an access control policy, to associated intrusion policies individually or in combination, or both for each targeted device.

You must apply both an access control policy and an associated intrusion policy in either of the following cases:

- when the access control policy is being applied to the device for the first time
- when an intrusion policy has been newly added to the access control policy

In both cases, the states of the access control policy and the intrusion policies are linked; that is, you must apply both or neither.

Note that regardless of the intrusion policies you apply, applying an access control policy automatically applies all associated file policies that are different from those currently running on devices targeted by the policy. You cannot apply file policies independently.

The Access Control Policy Column

The Access Control Policy column provides a check box for indicating whether to apply the access control policy.

TIP! Although you can reapply a policy while it is still in the task queue, that is, while the apply task has not yet completed, there is no benefit in doing this.

A status message indicates whether the policy is currently up to date or out of date. When the policy is out of date, you can conveniently display a comparison of the policy to the currently running policy in a new browser window. The comparison does not include differences in an intrusion policy associated with the access control policy.

The Intrusion Policies Column

The Intrusion Policies column provides one or more check boxes for indicating whether to apply intrusion policies associated with the access control policy to a device. A single grayed check box indicates that all associated intrusion policies are identical to currently running policies, in which case the check box is cleared and cannot be selected. You cannot apply an unchanged intrusion policy; only changed intrusion policies are listed, and can be selected individually. When the same intrusion policy is associated with multiple rules in a policy, the intrusion policy is listed only once for each device.

The check box for an intrusion policy is selected and the check box is grayed and cannot be changed when the access control policy and the intrusion policy must be applied together, as described above, in either of the following cases:

- when the access control policy is being applied to the device for the first time
- when an intrusion policy has been newly added to the access control policy

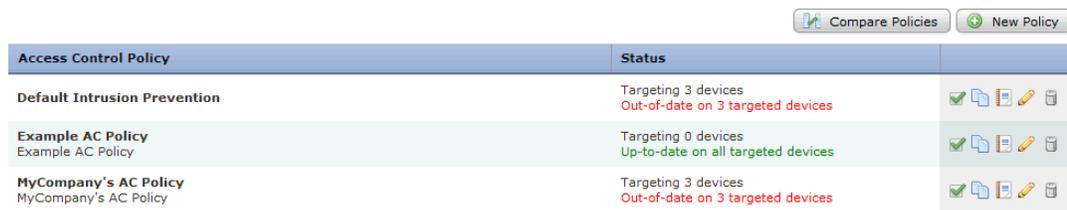
Status messages indicate whether intrusion policies are currently up to date or out of date. An intrusion policy is out of date when it is not identical to an intrusion policy currently running on the listed device. An identical intrusion policy on the device is up to date. When the policy is out of date, you can conveniently display a comparison of the policy to the currently running policy in a new browser window.

To apply selected access control policy configurations:

ACCESS: Admin/Security Approver

1. Select **Policies > Access Control**.

The Access Control page appears.



Access Control Policy	Status	
Default Intrusion Prevention	Targeting 3 devices Out-of-date on 3 targeted devices	   
Example AC Policy Example AC Policy	Targeting 0 devices Up-to-date on all targeted devices	   
MyCompany's AC Policy MyCompany's AC Policy	Targeting 3 devices Out-of-date on 3 targeted devices	   

2. Click the apply icon () next to the policy you want to apply.

The Apply Access Control Rules pop-up window appears.



Alternatively, you can click **Save and Apply** on the policy Edit page; see [Editing an Access Control Policy](#) on page 459.

3. Click **Details**.

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

Devices	Access Control Policy		Intrusion Policies	
	Select: All None		Select: All None	
linden	<input checked="" type="checkbox"/>	<input type="checkbox"/> Out-of-date	<input checked="" type="checkbox"/> Balanced Security anc	<input type="checkbox"/> Out-of-date
xiramat	<input checked="" type="checkbox"/>	<input type="checkbox"/> Out-of-date	<input checked="" type="checkbox"/> Balanced Security anc	<input type="checkbox"/> Out-of-date
tamarix	<input checked="" type="checkbox"/>	<input type="checkbox"/> Out-of-date	<input checked="" type="checkbox"/> Balanced Security anc	<input type="checkbox"/> Out-of-date

TIP! You can also open the pop-up window from the Access Control page (**Policies > Access Control**) by clicking on an out-of-date message in the **Status** column for the policy.

4. Select or clear the access control policy check box next to the device name to specify whether to apply the access control policy to a targeted device.
5. Select or clear the intrusion policy check box next to the device name to specify whether to apply an intrusion policy to a targeted device.
6. Click **Apply Selected Configurations**.

Your policy apply task is queued. Click **OK** to return to the Access Control page.

TIP! You can monitor the progress of the policy apply task on the Task Status page (**System > Monitoring > Task Status**).

IMPORTANT! A pop-up window may warn that you have exceeded the maximum number of intrusion policies supported by the device. To successfully apply the access control policy, remove applied intrusion policies from it until the number of remaining intrusion policies (including the default action) falls within the maximum.

CHAPTER 13

UNDERSTANDING AND WRITING ACCESS CONTROL RULES

A set of *access control rules* is a key component of an access control policy. Although you can create basic access control policies without them, access control rules allow you to manage, in a granular fashion, which traffic can enter your network, exit it, or cross from within without leaving it. For example, you could block some or all social networking traffic, prevent your sales department from accessing accounting records, monitor which users access which sites or networks, and so on.

IMPORTANT! Hardware-based fast-path rules and Security Intelligence-based traffic filtering (blacklisting) occur **before** network traffic is evaluated by access control rules.

Within an access control policy, the system matches traffic to rules in top-down order by rule number. In addition to its rule order and some other basic attributes, each rule has the following major components:

- a set of rule *conditions* that identifies the specific traffic you want to control
- a rule *action*, which determines how the system handles traffic that meets the rule's conditions
- file, malware, and intrusion *inspection* options, which allow you to examine (and optionally block) matching traffic that you would otherwise allow
- *logging* options, which allow you to keep a record of the matching traffic and how it was handled by the rule

The access control policy's *default action* handles traffic that is not blacklisted by Security Intelligence and does not meet the conditions of any non-Monitor rule in the policy. For more information on access control policies and the default action,

see [Using Access Control Policies](#) on page 424.

TIP! If you want to use the Sourcefire 3D System to perform intrusion detection and prevention but do not need to take advantage of discovery data, you can optimize performance by disabling new discovery. First, make sure that your applied access control policies do not contain rules with user, application, or URL conditions. Then, remove all rules from your network discovery policy and apply it to your managed devices. For more information on configuring discovery, see [Introduction to Network Discovery](#) on page 1250.

Although you can create access control rules with any license, certain rule conditions and inspection options require that you enable specific licensed capabilities on the access control policy’s targeted devices. You cannot apply a policy that uses licensed capabilities to unlicensed devices. The Defense Center uses warning icons (⚠) and confirmation dialogs to designate unlicensed features. For details, hover your pointer over a warning icon.

The following table explains the licenses you must have to use access control rules.

License Requirements for Access Control Rules

TO APPLY AN ACCESS CONTROL POLICY THAT INCLUDES RULES...	ADD THIS LICENSE...	TO ONE OF THESE DEFENSE CENTERS...	AND ENABLE IT ON ONE OF THESE DEVICES...
with zone, network, VLAN, or port conditions, or URL conditions that use literal URLs and URL objects only	Any	Any	Any, except Series 2 devices cannot perform URL filtering using literal URLs and URL objects
associated with intrusion policies, or file policies that do not perform malware detection or blocking	Protection	Any	Any, except Series 2 devices cannot perform Security Intelligence filtering
associated with file policies that perform malware detection or blocking	Malware	Any except DC500	Series 3, virtual
with application or user conditions	Control	Any, except the DC500 cannot perform user control	Series 3, virtual
with URL conditions that use URL category and reputation data	URL Filtering	Any except DC500	Series 3, virtual

See the following sections for more information on access control rules:

- [Creating and Editing Access Control Rules](#) on page 474
- [Understanding Rule Actions](#) on page 478
- [Understanding Rule Conditions and Condition Mechanics](#) on page 482
- [Working with Different Types of Conditions](#) on page 491
- [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511
- [Logging Connection, File, and Malware Information](#) on page 515
- [Adding Comments to a Rule](#) on page 521

Creating and Editing Access Control Rules

LICENSE: Any

An access control rule is simply a set of configurations and conditions that:

- qualifies network traffic
- specifies how and whether you further inspect and log traffic that matches those qualifications
- determines the traffic's eventual flow

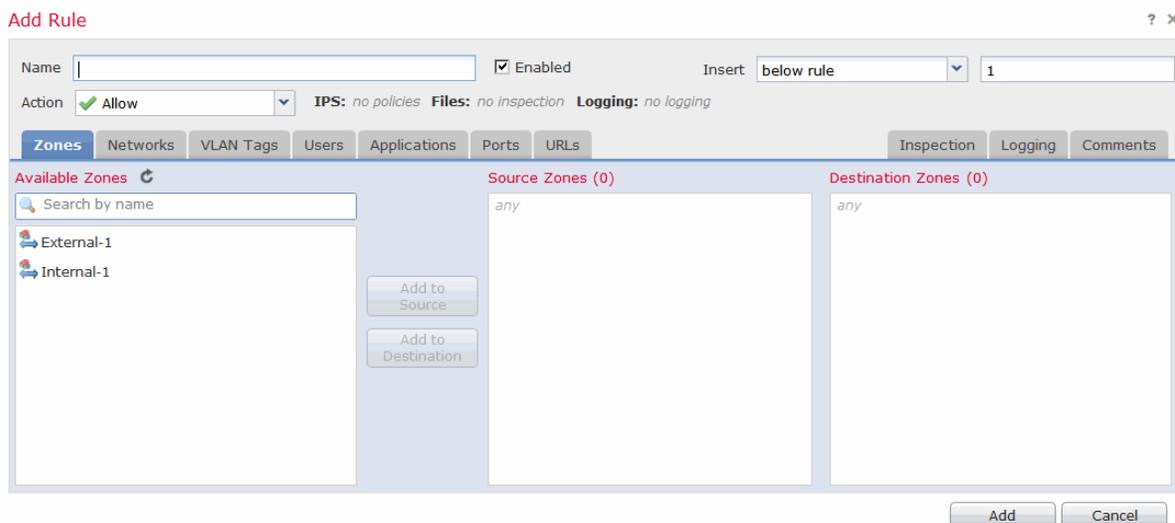
You create and edit access control rules from within an existing access control policy. Each rule belongs to only one policy.

When you apply an access control policy to a device, the Defense Center sends each rule defined in the policy to the device as a set of expanded rules, where each rule expresses one possible combination of conditions in the rule. For example, a rule with the Internal security zone as a source zone and LDAP and HTTPS source ports would be sent to the device as two rules: one to match traffic with a source zone of Internal over an LDAP source port, and one to match traffic with a source zone of Internal over an HTTPS source port.

Note that an access control policy with many complex rules may not apply to a managed device if the number of expanded rules exceeds the number allowed for that device. If this occurs, analyze the conditions in your rules to see if you can eliminate unnecessary settings.

The web interface for adding or editing a rule is similar. You specify the rule name, state, action, and position at the top of the page. You build conditions using the tabs on the left side of the page; each condition type has its own tab. You

configure inspection and logging options, as well as add comments to the rule, using the tabs on the right side of the page.



The following list summarizes the configurable components of an access control rule.

Name

Give each rule a unique name. You can use up to thirty printable characters, including spaces and special characters, with the exception of the colon (:).

Rule State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic. When viewing the list of rules in an access control policy, disabled rules are grayed out, although you can still modify them.

Action

A rule's action determines how the system handles traffic that matches the rule's conditions. You can trust, monitor, block, or allow (with or without further inspection) matching traffic. The access control policy's *default action* handles traffic that does not meet the conditions of any non-Monitor access control rule.

IMPORTANT! Access control rules actions, along with the policy's default action, determine the network traffic that you can examine using intrusion, file, or network discovery policies. The system does **not** perform inspection on trusted or blocked traffic.

For detailed information on rule actions and how they affect inspection and traffic flow, see [Understanding Rule Actions](#) on page 478.

Current Inspection and Logging Settings

The **IPS**, **Files**, and Logging options indicate the intrusion policy, file policy, and logging options currently selected in the rule. Click the **IPS** or **Files** setting to open the Inspection tab, or click the **Logging** setting to open the Logging tab.

Position (Order and Category)

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. Optionally, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add your own custom categories anywhere you like, but you cannot delete the Sourcefire-provided categories or change their order.

When you add a rule to a policy, you specify its position in one of two ways. First, you can **Insert** it in a category, which places it last (numerically) in that category. Or, you can place it **above** or **below** a specific rule, using rule numbers as a reference point. When editing an existing rule, you can **Move** the rule in a similar fashion. For more information, see [Organizing Rules in a Policy](#) on page 449.

Conditions

Rule conditions identify the specific traffic you want to control. Conditions can match traffic by any combination of multiple attributes, including security zone, network, VLAN, Active Directory LDAP user or group, application, transport protocol port, or URL information. Conditions can be simple or complex, and in some cases require that you apply a license to the access control policy's target devices.

For detailed information on adding conditions, see [Understanding Rule Conditions and Condition Mechanics](#) on page 482 and [Working with Different Types of Conditions](#) on page 491.

File and Intrusion Inspection Options

A rule's inspection options apply to traffic that you would normally allow. You configure the system to perform further inspection by associating an intrusion or file policy (or both) with a rule.

File policies perform *file control*, that is, they can detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. File policies can also use Sourcefire's advanced malware protection feature to determine if certain transmitted files represent a threat to your organization, then block them. Intrusion policies perform intrusion detection and prevention and can drop offending packets.

Both types of inspection require the Protection license. AMP requires a Malware license. For detailed information on associating an intrusion or file policy with a rule, see [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

Logging Options

An access control rule's logging options allow you to specify whether and how to keep a record of matching traffic, as well as log the detection of files or malware in that traffic.

In general, you can log a connection event at the beginning or end of a connection, or both. However, you can log only beginning-of-connection events for blocked traffic because matching traffic is denied without further inspection. Additionally, although the system automatically logs end-of-connection events for monitored traffic, beginning-of-connection logging for monitored traffic is determined by the first non-Monitor rule triggered by the traffic, or the default action.

If you choose to log at the end of connections, the system generates events when it detects the close of a connection, when it does not detect the end of a connection after a period of time, or when it can no longer track the session due to memory constraints.

You can log connections to the Defense Center database, as well as to the system log (syslog) or to an SNMP trap server. For detailed information, see [Logging Connection, File, and Malware Information](#) on page 515.

Comments

Each time you save changes to 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. For more information, see [Adding Comments to a Rule](#) on page 521.

To create or edit an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.

The Access Control page appears.

2. Click the edit icon (✎) next to the access control policy where you want to add a rule.

The policy Edit page appears.

3. Add a new rule or edit an existing rule:
 - 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.

Either the Add Rule or the Editing Rule page appears.

TIP! You can use the right-click context menu to perform many rule creation and management actions; see [Using the Context Menu](#) on page 63. You can also drag and drop rules to change their order.

4. Configure the rule components, as described earlier in this section. You can configure the following, or accept the defaults:
 - You must provide a unique rule **Name**.
 - Specify whether the rule is **Enabled**.
 - Select a rule **Action**.
 - Specify the rule position.
 - Configure the rule's conditions.
 - Configure the rule's **Inspection** options.
 - Specify **Logging** options.
 - Add **Comments**.

5. Click **Add** or **Save**.

Your changes are saved. You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Understanding Rule Actions

LICENSE: Any

Every access control rule has an associated action that determines:

- whether the system will trust, monitor, block, or allow (with or without further inspection) traffic that matches the rule's conditions
- for certain rule actions, whether the system further inspects matching traffic with intrusion, file, and network discovery policies before allowing it to pass
- 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 the Default Action](#) on page 428. For detailed information on rule actions and how they affect connection logging, see following sections, as well as [Logging Connection, File, and Malware Information](#) on page 515.

Allow

The **Allow** action allows matching traffic to pass. Optionally, you can associate an Allow rule with an intrusion or file policy, or both. These two types of policy further inspect and can block network traffic according to their own configurations:

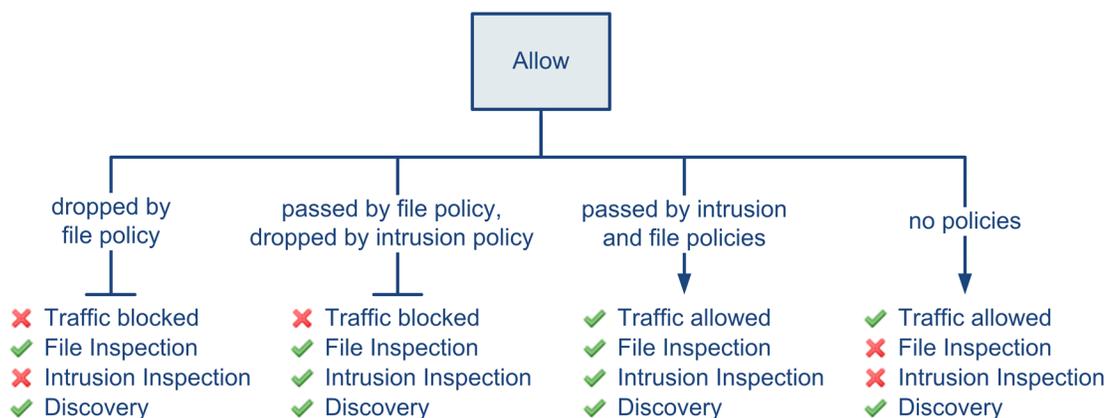
- Use an associated file policy to perform file control, that is, to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. File policies also allow you to inspect a restricted set of those files for malware, and optionally block detected malware.
- Use an associated intrusion policy to analyze network traffic according to intrusion detection and prevention configurations and, optionally, drop offending packets.

For instructions on how to associate an intrusion or file policy with an access control rule, see [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

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 Block and Interactive Block with Reset](#) on page 481). 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.

Regardless of whether the traffic is inspected or dropped by an intrusion or file policy, the system can inspect it using network discovery.

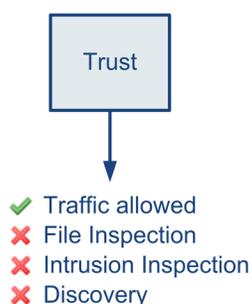


IMPORTANT! Selecting a rule action of **Allow** does not automatically guarantee discovery inspection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy. For more information, see [Introduction to Network Discovery](#) on page 1250.

You can log allowed network traffic at both the beginning and end of connections.

Trust

The Trust action allows traffic to pass without further inspection. You cannot inspect trusted traffic with a file, intrusion, or network discovery policy.



You can log trusted network traffic at both the beginning and end of connections. Note that TCP connections detected by a trust rule on the first packet will not generate a beginning-of-connection event. On virtual appliances, generated events will be reported one hour after the final session packet.

Monitor

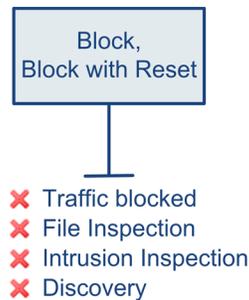
The **Monitor** action does not affect traffic flow; matching traffic is neither immediately permitted nor denied. Rather, traffic is matched against additional rules, if present, to determine whether to permit or deny it. The first non-Monitor rule matched determines traffic flow and any further inspection. 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. The action associated with a logged connection is either that of the first non-Monitor rule triggered by the connection, or the default action.

If locally-bound traffic matches a monitor rule in a Layer 3 deployment, that traffic may bypass inspection. To ensure inspection of the traffic, enable **Inspect Local Router Traffic** in the advanced device settings for the managed device routing the traffic.

Block and Block with Reset

The **Block** and **Block with reset** actions deny traffic without further inspection. Block with reset rules also reset the connection. You cannot inspect blocked traffic with a file, intrusion, or network discovery policy.



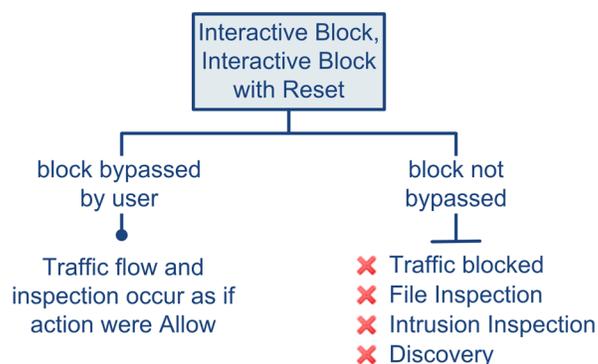
You can log blocked network traffic only at the beginning of connections.

Interactive Block and Interactive Block with Reset

For HTTP traffic, the **Interactive Block** and **Interactive Block with reset** actions give users a chance to bypass a website block, by clicking through a warning page. If a user does not bypass the block, matching traffic is denied without further inspection. Interactive Block with reset rules also reset the connection. For information on configuring the warning page, see [Adding an HTTP Response Page](#) on page 436.

On the other hand, if a user bypasses the block, matching network traffic is treated identically to allowed traffic; see [Allow](#) on page 479. When the system initially blocks a user's HTTP request using an Interactive Block rule, it marks the beginning-of-connection event with the Interactive Block or Interactive Block with

Reset action. 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. Therefore, as with Allow rules, you can associate either type of Interactive Block rule with a file and intrusion policy. The system can also use network discovery to inspect this user-allowed traffic.



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.

Understanding Rule Conditions and Condition Mechanics

LICENSE: Any

You can add conditions to access control rules to identify the type of traffic that matches the rule. You can add any of seven types of conditions to a rule, either alone or in any combination.

For each condition type, you select conditions you want to add to a rule from a list of available conditions. When applicable, condition filters allow you to constrain available conditions. Lists of available and selected conditions may be as short as a single condition or many pages long. You can search available conditions and display only those matching a typed name or value in a list that updates as you type.

Depending on the type of condition, lists of available conditions may be comprised of a combination of conditions provided directly by Sourcefire or configured using other Sourcefire 3D System features, including objects created using the object manager (**Objects > Object Management**), objects created directly from individual conditions pages, and literal conditions.

See the following sections for information on specifying rule conditions:

- [Understanding Rule Conditions](#) on page 483 defines the different types of rule conditions.
- [Adding Rule Conditions](#) on page 485 describes the controls used to select and add rule conditions.

- [Searching Condition Lists](#) on page 489 explains how to search available conditions and display only those matching a typed name or value in a list that updates as you type.
- [Adding Literal Conditions](#) on page 490 explains how to add literal conditions to a rule.
- [Using Objects in Conditions](#) on page 491 explains how to add individual objects to the system from the configuration pages for relevant condition types.

Understanding Rule Conditions

LICENSE: Any

You can set an access control rule to match traffic meeting any of the conditions described in the following table:

Access Control Rule Condition Types

CONDITION	DESCRIPTION	SUPPORTED DEFENSE CENTERS	SUPPORTED DEVICES
Zones	A configuration of one or more interfaces where you can apply policies. Zones provide a mechanism for classifying traffic on source and destination interfaces, and you can add source and destination zone conditions to rules. See Working with Security Zones on page 187 for information on creating zones using the object manager. See Adding Zone Conditions on page 492 for more information on adding these conditions.	Any	Any
Networks	Any combination of individual IP addresses, CIDR blocks, and prefix lengths, either specified explicitly or using network objects and groups (see Working with Network Objects on page 161). You can add source and destination network conditions to rules. See Adding Network Conditions on page 493 for more information on adding these conditions.	Any	Any
VLAN Tags	A number from 0 to 4094 that identifies traffic on your network by VLAN. See Working with VLAN Tag Objects on page 174 for information on creating individual and group VLAN Tag objects using the object manager. See Adding VLAN Tag Conditions on page 495 for more information on adding these conditions.	Any	Any

Access Control Rule Condition Types (Continued)

CONDITION	DESCRIPTION	SUPPORTED DEFENSE CENTERS	SUPPORTED DEVICES
Users	Individual LDAP users and user groups retrieved from a Microsoft Active Directory Server. See Understanding LDAP Authentication on page 1854 for information on specifying and retrieving the users and groups you want to use for user control. See Adding User Conditions on page 497 for more information on adding these conditions.	Any except DC500	Series 3, virtual
Applications	Applications provided by Sourcefire, user-defined applications, and application filters you create using the object manager. See Working with Application Detectors on page 1663 and Working with Application Filters on page 176 for more information. See Working with Application Conditions on page 498 for more information on adding these conditions.	Any	Series 3, virtual
Ports	Transport protocol ports, including individual and group port objects you create based on transport protocols. See Working with Port Objects on page 173 for information on creating individual and group transport protocol objects using the object manager. See Adding Port Conditions on page 504 for more information on adding these conditions.	Any	Any
URLs	Sourcefire-provided URLs grouped by category and reputation, literal URLs, and any individual and group URL objects you create using the object manager. See Enabling Sourcefire Cloud Communications on page 2032 and Working with URL Objects on page 175 for more information. See Adding URL Conditions on page 506 for more information on adding these conditions.	Any except DC500 (DC500 does support literal URLs, URL objects, and URL object groups)	Series 3, virtual

You can filter traffic that matches one or more types of access control rule conditions, in any combination. The system links multiple conditions of the same type with an OR operation, and links different condition types with an AND operation. For example, if your rule conditions are:

Destination Networks: 10.4.0.0/16, 10.5.0.0/16
VLAN Tags: 11

then, the rule would match traffic going to a host that is in either 10.4.0.0/16 or 10.5.0.0/16 on VLAN 11, such as:

10.1.1.1 to 10.4.12.1 on VLAN 11

or:

192.168.2.1 to 10.5.15.23 on VLAN 11

When you specify multiple conditions in a rule, traffic that matches all conditions of the rule matches the rule. When you do not add any conditions of a particular type to a rule, the system uses the default setting of **any** traffic of that type, which means that the system does not filter traffic based on that condition type.

Note that you can add source and destination zone conditions and source and destination network conditions. For other conditions, where you do not specify source or destination, all traffic is compared to the rule to identify matching traffic.

Adding Rule Conditions

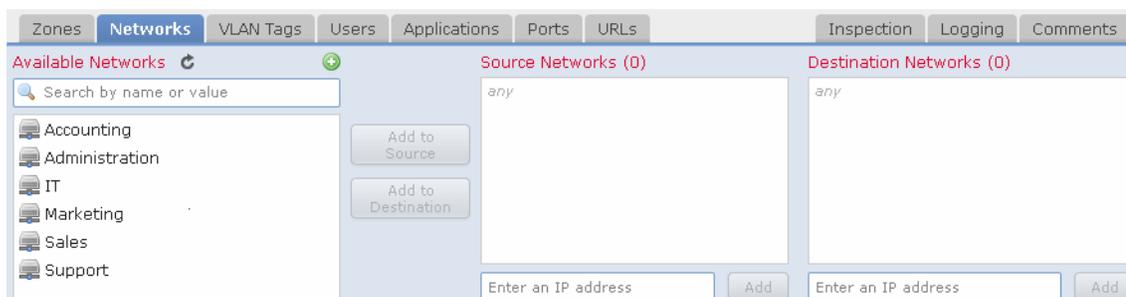
LICENSE: Any

Adding conditions to access control rules is essentially the same for each type of condition. You select from one or two lists of available conditions on the left, and add the selected conditions to one or two lists of selected conditions on the right.

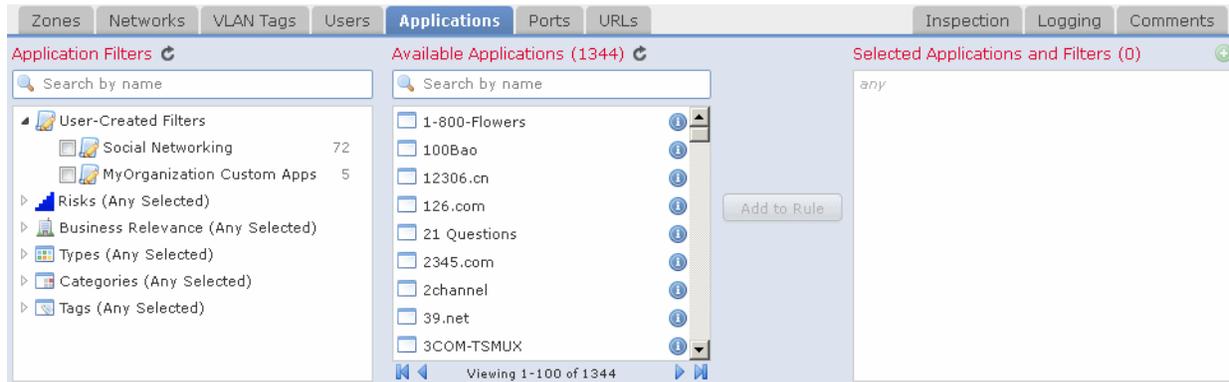
For all condition types, you select one or more individual available conditions by clicking on them to highlight them. For application conditions, you can also select or clear check boxes to constrain the list of available applications using Sourcefire-provided or user-defined filters.

In all cases, you can either click a button between the two types of lists to add available conditions that you select to your lists of selected conditions, or drag and drop available conditions that you select into the list of selected conditions.

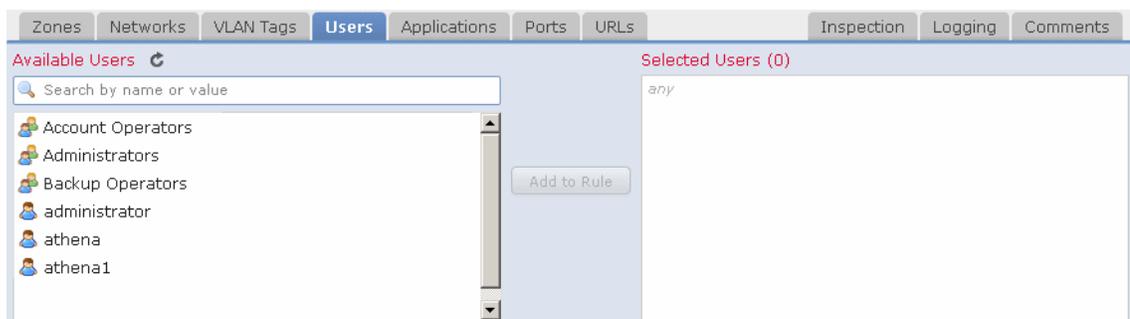
Some pages (Zones, Networks, and Ports) have one list of available conditions on the left, which can be added to either of two lists of selected conditions on the right.



Other pages (Applications and URLs) have two lists of available conditions on the left, which can be used together to select available conditions to add to a single list of selected conditions on the right.



Other pages (VLAN Tags and Users) have one list of available conditions on the left, which can be added to a single list of selected conditions on the right.



You can add up to 50 conditions of each type to a list of selected conditions. For example, you can add up to 50 source zone conditions, up to 50 destination zone filters, up to 50 user conditions, and so on, until you reach the upper limit for the appliance.

Note that when you apply an access control policy to a device, the Defense Center sends each rule defined in the policy to the device as a set of expanded rules, where each rule expresses one possible combination of conditions in the rule. For example, a rule with the Internal security zone as a source zone and LDAP and HTTPS source ports would be sent to the device as two rules: one to match traffic with a source zone of Internal over an LDAP source port, and one to match traffic with a source zone of Internal over an HTTPS source port.

An access control policy with many complex rules may not apply to a managed device if the number of expanded rules exceeds the number allowed for that device. If this occurs, analyze the conditions in your rules to see if you can eliminate unnecessary settings.

When a list of available conditions contains more conditions than can be displayed on a single page, you can use navigation links under the list to switch between pages.

The following table describes the actions you can take to select and add conditions to a rule.

Adding Conditions

To...	You CAN...
select available conditions to add to a list of selected conditions	click the available condition; use the Ctrl and Shift keys to select multiple conditions.
select all listed available conditions	right-click the row for any available condition, then click Select All .
search a list of available conditions or filters	click inside the search field and type a search string. See Searching Condition Lists on page 489 for more information.
clear a search when searching available conditions or filters	click the reload icon () above the search field or the clear icon () in the search field.
add selected conditions from a list of available conditions to a list of selected source or destination conditions	click Add to Source or Add to Destination . You can add zone and network conditions to lists of source and destination conditions. See Adding Zone Conditions on page 492, Adding Network Conditions on page 493, or Adding Port Conditions on page 504 for more information.
add selected conditions from a list of available conditions to a single list of selected conditions	click Add to Rule . VLAN tag, user, application, and URL conditions use single lists of selected conditions.
drag and drop selected available conditions into a list of selected conditions	right-click a selected condition, then drag and drop into the list of selected conditions.
add a literal condition to a list of selected conditions using a literal field	click to remove the prompt from the literal field, type the literal condition, then click Add . Network, VLAN tag, and URL conditions provide a field for adding literal conditions.

Adding Conditions (Continued)

To...	You CAN...
add a literal condition to a list of selected conditions using a drop-down list	select a condition from the drop-down list, then click Add . Port conditions provide a drop-down list for adding literal conditions. See Adding Port Conditions on page 504 for more information.
add an individual object or condition filter so you can then select it from the list of available conditions	click the add icon (). See Using Objects and Security Zones on page 158 for information on adding objects using the object manager.
delete a single condition from a list of selected conditions	click the delete icon () next to the condition
delete a condition from a list of selected conditions	right-click to highlight the row for a selected condition, then click Delete .
delete multiple conditions from a list of selected conditions	use the Shift and Ctrl keys to select multiple conditions, or right-click and Select All ; next, right-click to highlight the row for a selected condition, then click Delete Selected .

Conditions you select are grayed out and can no longer be added to the same list of selected conditions. When you select conditions you have added, the add button is also grayed out. When you select conditions you have not previously added, the add button activates and can be used. In the following example, the 100Bao and 2channel applications have been added and are currently selected. Both the selected applications and the **Add to Rule** button are grayed out.



Similarly, if conditions cannot be used in combination, such as mixed transport protocols for source and destination ports, conditions that are invalid based on previous selections are grayed out.

On the relevant condition page, and also on the policy Edit page, you can hover your pointer over an individual object to display the contents of the object, and over a group object to display the number of individual objects in the group.

The following basic procedure explains how to add conditions to a new rule. See [Creating and Editing Access Control Rules](#) on page 474 for complete instructions on adding and modifying rules.

To add available conditions to a list of selected conditions:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon (✎) next to the access control policy you want to modify.
The policy Edit page appears.
3. Click **Add Rule**.
The Add Rule page appears.



4. Click the tab for the type of condition you want to add to the rule.
The conditions page appears for the type of condition you selected.
5. Take any of the available actions in the [Adding Conditions](#) on page 487.
6. Click **Add** to save your configuration.
Your rule is added and the policy Edit page appears.

Searching Condition Lists

LICENSE: Any

You can filter a list of available access control rule conditions and condition categories to limit the number of items displayed in the list. The list updates as you type to display matching items.

Optionally, you can search on object names and on the values configured for 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*.

The following basic procedure explains how to filter a list in a new rule. See [Creating and Editing Access Control Rules](#) on page 474 for complete instructions on adding and modifying rules.

To search a list of available conditions or condition categories:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.

2. Click the edit icon (✎) next to the access control policy you want to modify. The policy Edit page appears.
3. Click **Add Rule**.
The Add Rule page appears. The following graphic shows the search field on the Networks conditions page.

4. To search a list, click inside the search field to clear the prompt, then type a search string.
The list updates as you type to display matching items and a clear list icon (✕) appears in the search field. The list updates and no items are listed when none match the search string.
5. Optionally, click the reload icon (↻) above the search field or click the clear icon (✕) in the search field to clear the search string.
The complete list appears.
6. Click **Add** to save your configuration.
Your rule is added and the policy Edit page appears.

Adding Literal Conditions

LICENSE: Any

You can add a literal value to the list of selected conditions for the following condition types:

- Networks
- VLAN Tags
- Ports
- URLs

For all but port conditions, you type the literal value in a configuration field below the list of selected conditions.

In the case of port conditions, you select a protocol from a drop-down list. When the protocol is **A11** (for destination ports) and, optionally, when the protocol is **TCP** or **UDP**, you type a port number in a configuration field. When the protocol is **ICMP** or **IPV6-ICMP**, you select a type and, if appropriate, a related code. When you add a source port, the protocol defaults to **TCP**. You must specify a protocol when setting a literal port.

Each relevant conditions page provides the controls needed to add literal values. Values you type in a configuration field appear as red text if the value is invalid, or until it is recognized as valid. Typed values change to black text as you type when they are recognized as valid. A grayed **Add** button activates when a valid value is

recognized. Literal values you add appear immediately in the list of selected conditions.

See the following sections for specific details on adding each type of literal value:

- [Adding Network Conditions](#) on page 493
- [Adding VLAN Tag Conditions](#) on page 495
- [Adding Port Conditions](#) on page 504
- [Adding URL Conditions](#) on page 506

Using Objects in Conditions

LICENSE: Any

Application filters and objects that you create in the object manager (**Objects > Object Management**) are immediately available for you to select from relevant lists of available access control rule conditions. See [Using Objects and Security Zones](#) on page 158 for information.

You can also create many objects on-the-fly from the access control policy. A control on relevant conditions pages provides access to the same configuration controls that you use in the object manager.

Individual objects created on-the-fly appear immediately in the list of available objects, and you can add them to the current rule, and to other existing and future rules. On the relevant conditions page, and also on the policy Edit page, you can hover your pointer over an individual object to display the contents of the object, and over a group object to display the number of individual objects in the group.

Working with Different Types of Conditions

LICENSE: Any

You can filter traffic by one or more of seven types of rule conditions, in any combination. See the following sections for more information:

- [Adding Zone Conditions](#) on page 492 explains how to filter traffic by security zones that you create using the object manager.
- [Adding Network Conditions](#) on page 493 explains how to filter traffic by IP address or address block.
- [Adding VLAN Tag Conditions](#) on page 495 explains how to filter traffic by VLAN tag.
- [Adding User Conditions](#) on page 497 explains how to filter traffic by users and user groups retrieved from a Microsoft Active Directory Server.
- [Working with Application Conditions](#) on page 498 explains how to filter traffic based on a predefined list of applications provided by Sourcefire, custom applications, and application filters you create using the object manager.

- [Adding Port Conditions](#) on page 504 explains how to filter traffic by specified transport protocol ports.
- [Adding URL Conditions](#) on page 506 explains how to filter traffic by URL, including by statistics such as reputation and category.

Adding Zone Conditions

LICENSE: Any

The security zones on your system are comprised of interfaces on your managed devices. Zones that you add to an access control rule target the rule to devices on your network that have interfaces in those zones. You can add security zones as conditions for access control rules. See [Working with Security Zones](#) on page 187 for information on creating security zones using the object manager.

Keep the following important points in mind when you filter traffic by zone:

- All zones in a rule must be of the same type (switched, routed, and so on).
- You can add a passive zone only as a source zone.
- The warning icon (⚠) next to a zone in the list of available zones indicates that the zone does not include an interface. When you hover your pointer over the icon, a message explains that the zone must include at least one interface for the rule to take effect. See [Working with Security Zones](#) on page 187.

IMPORTANT! In a Layer 2 deployment, you cannot block egress traffic based on destination network or destination security zone. You must instead write access control rules that block ingress traffic based on blocking source network or source security zone. For more information on Layer 2 deployments, see [Setting Up Virtual Switches](#) on page 281.

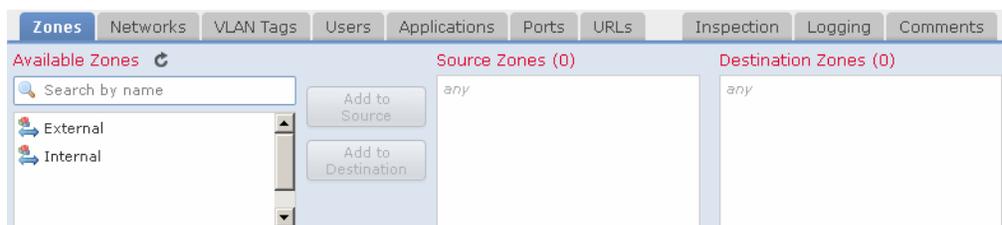
The following procedure explains how to add source and destination zone conditions while adding or editing an access control rule. See [Understanding Rule Conditions and Condition Mechanics](#) on page 482 for more detailed information.

To add zone conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **Zones** tab on the rule Edit page.

The Zones page appears.



2. Optionally, click the **Search by name** prompt above the Available Zones list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.

3. Click a condition in the Available Zones list. Use the Shift and Ctrl keys to select multiple conditions, or right-click and then click **Select All**.

Conditions you select are highlighted.

The warning icon (⚠) next to a zone indicates that the rule will not take effect because the zone does not include an interface. See [Working with Security Zones](#) on page 187.

4. You have the following choices:

- To filter traffic by source zone, click **Add to Source**.
- To filter traffic by destination zone, click **Add to Destination**.

Optionally, you can drag and drop selected conditions into the Source Zones or Destination Zones list.

Selected conditions are added. Note that you can add the same condition as both a source zone and a destination zone.

5. Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Adding Network Conditions

LICENSE: Any

You can add any of the following kinds of network conditions to an access control rule:

- individual and group network objects that you have created using the object manager

See [Working with Network Objects](#) on page 161 for information on creating individual and group network objects using the object manager.

- individual network objects that you add from the Network conditions page, and can then add to your rule and to other existing and future rules
See [Using Objects in Conditions](#) on page 491 for more information.
- literal, single IP addresses or address blocks
See [Adding Literal Conditions](#) on page 490 for more information.

IMPORTANT! In a Layer 2 deployment, you cannot block egress traffic based on destination network or destination security zone. You must instead write access control rules that block ingress traffic based on source network or source security zone. For more information on Layer 2 deployments, see [Setting Up Virtual Switches](#) on page 281.

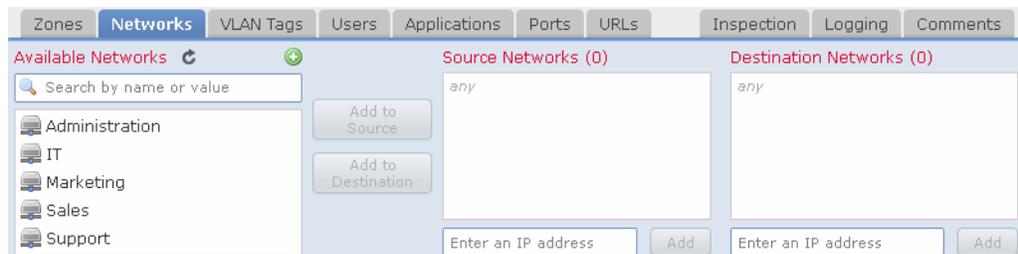
If you add rules to an access control policy that contain conditions matching destination IPv6 traffic, add an Allow rule with port conditions specifying traffic using the IPv6 Neighbor Discovery Protocol (ICMPv6 types 135 and 136) before those rules. For more information on port conditions, see [Adding Port Conditions](#) on page 504.

The following procedure explains how to add source and destination network conditions while adding or editing an access control rule. See [Understanding Rule Conditions and Condition Mechanics](#) on page 482 for more detailed information.

To add network conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **Networks** tab on the rule Edit page.
The Networks page appears.



2. Optionally, click the **Search by name or value** prompt above the **Available Networks** list, then type a name or value.
The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.
3. Click a condition in the **Available Networks** list. Use the Shift and Ctrl keys to select multiple conditions, or right-click and then click **Select All**.
Conditions you select are highlighted.

4. You have the following choices:
 - To filter traffic by source network, click **Add to Source**.
 - To filter traffic by destination network, click **Add to Destination**.

Alternatively, you can drag and drop selected conditions into the **Source Networks** or **Destination Networks** list.

Conditions you selected are added. Note that you can add the same condition as both a source network and a destination network.

5. Optionally, click the add icon () above the Available Networks list to add an individual network object.

You can add multiple IP addresses, CIDR blocks, and prefix lengths to each network object. Optionally, you can then select the object you added. See [Working with Network Objects](#) on page 161 and [Using Objects in Conditions](#) on page 491 for more information.

6. Optionally, 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**.

The list updates to display your entry. See [Adding Literal Conditions](#) on page 490 for more information.

7. Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Adding VLAN Tag Conditions

LICENSE: Any

You can add any of the following kinds of VLAN tag conditions to an access control rule:

- individual and group VLAN tag objects that you have created using the object manager
See [Working with Network Objects](#) on page 161 for information creating individual and group VLAN tag objects using the object manager.
- individual VLAN tag objects that you add from the VLAN Tags conditions page, and can then add to your rule and to other existing and future rules
See [Using Objects in Conditions](#) on page 491 for more information.
- literal VLAN tag conditions
See [Adding Literal Conditions](#) on page 490 for more information.

The system inspects all traffic on your network for VLAN tags you specify, and uses the innermost VLAN tag to identify a packet by VLAN.

The following procedure explains how to add VLAN conditions while adding or editing an access control rule. See [Understanding Rule Conditions and Condition](#)

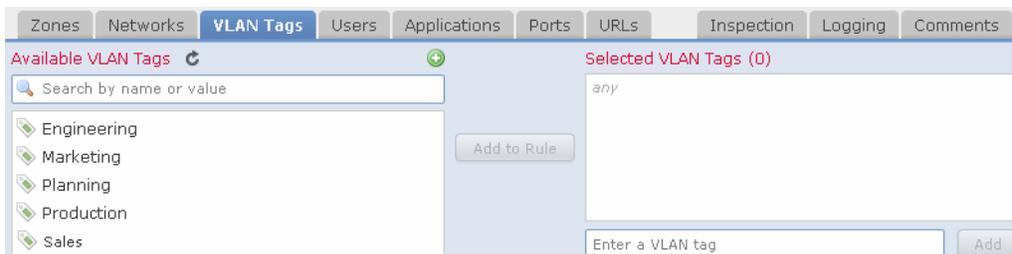
[Mechanics](#) on page 482 for more detailed information.

To add VLAN tag conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **VLAN Tags** tab on the rule Edit page.

The VLAN Tags page appears.



2. Optionally, click the **Search by name or value** prompt above the Available VLAN Tags list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.

3. Click a condition in the Available VLAN Tags list. Use the Shift and Ctrl keys to select multiple conditions, or right-click, then click **Select All**.

Conditions you select are highlighted.

4. You have the following choices:

- Click **Add to Rule**.
- Drag and drop selected conditions into the Selected VLAN Tags list.

Conditions you selected are added.

5. Optionally, click the add icon (+) above the Available VLAN Tags list to add a VLAN tag object.

In each VLAN tag object you add, you can specify any VLAN tag from 1 to 4094; use a hyphen to specify a range of VLAN tags. You can then select the object you added. See [Working with VLAN Tag Objects](#) on page 174 and [Using Objects in Conditions](#) on page 491 for more information.

6. Optionally, click the **Enter a VLAN Tag** prompt beneath the Selected VLAN Tags list, type a VLAN tag or range, then click **Add**.

You can specify any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

The list updates to display your entry. See [Adding Literal Conditions](#) on page 490 for more information.

7. Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Adding User Conditions

LICENSE: Control

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

You can configure access control rules to match traffic for users and user groups retrieved from a Microsoft Active Directory Server.

Before you can write access control rules with user conditions, you must configure a connection between the Defense Center 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 and groups you can use in user conditions. For more information, see [Creating LDAP Connections with the Defense Center](#) on page 1298.

In addition, you must install Sourcefire User Agents, which you install on your Microsoft Active Directory Servers. The agents monitor users when they authenticate against Active Directory credentials, and send records of those logins to the Defense Center. These records associate users with IP addresses, which is what allows access control rules with user conditions to trigger. For more information, see [Configuring Defense Center-User Agent Connections](#) on page 1307.

Keep in mind that if you specify a group, that automatically includes all of the group's members, including members of any sub-groups, with the exception of individually excluded users and members of excluded sub-groups.

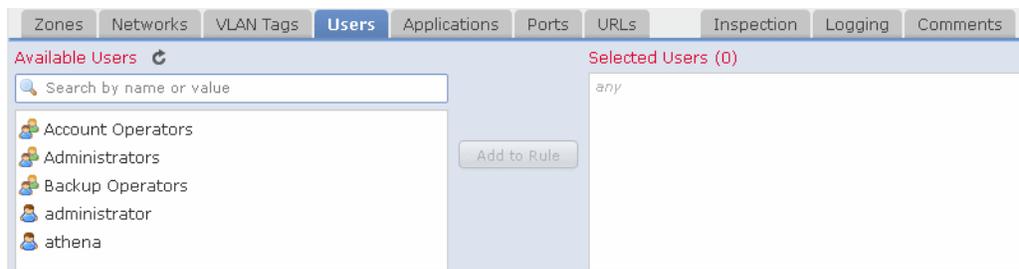
WARNING! If you configure user awareness parameters that include a very large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user mappings based on groups, due to memory limitations. As a result, access control rules based on user groups may not fire as expected.

The following procedure explains how to add user conditions while adding or editing an access control rule. See [Understanding Rule Conditions and Condition Mechanics](#) on page 482 for more detailed information.

To add user conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **Users** tab on the rule Edit page.
The Users page appears.



2. Optionally, click the **Search by name or value** prompt above the Available Users list, then type a name or value.
The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.
3. Click a condition in the Available Users list. Use the Shift and Ctrl keys to select multiple conditions, or right-click and then click **Select All**.
Conditions you select are highlighted.
4. You have the following choices:
 - Click **Add to Rule**.
 - Drag and drop selected conditions into the Selected Users list.Conditions you selected are added.
5. Save or continue editing the rule.
You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Working with Application Conditions

LICENSE: Control

SUPPORTED DEVICES: Series 3, virtual

You can configure access control rules to match application traffic. You can use either individual applications or application filters, either Sourcefire-provided and user-defined, as conditions for an access control rule. You can add applications and filters, in any combination, as long as the total number of items does not exceed 50. If the existing filters do not meet your needs, you can create an application filter on the fly while creating an application condition; you can then

use the new filter in your rule and in other existing and future rules. See the following sections for more details:

- For information on Sourcefire-provided and user-defined applications, see [Understanding Application Detection](#) on page 1263 and [Working with Application Detectors](#) on page 1663.
- For information on Sourcefire-provided and user-defined application filters, see [Working with Application Filters](#) on page 176.
- For information on adding an application filter on the fly, see [Using Objects in Conditions](#) on page 491.

Note the following when adding applications:

- The system applies the default policy action to packets that do not have a payload in a connection where an application is identified; this would be the case, for example, when a TCP connection is being established.
- It is not possible to identify applications or filter URLs before a connection is established between the client and the server. Therefore, when a packet matches all the other conditions in a rule containing an application or a URL, if application identification has not been completed, the packet is allowed to pass. This behavior allows a connection to be established so that applications can be identified.
- At least one detector must be enabled (see [Activating and Deactivating Detectors](#) on page 1678) for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all Sourcefire-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. See [Applying an Access Control Policy](#) on page 466.

See the following sections for more information:

- [Understanding Application Condition Lists](#) on page 499
- [Adding Application Conditions](#) on page 501

Understanding Application Condition Lists

LICENSE: Control

SUPPORTED DEVICES: Series 3, virtual

The Applications conditions page displays three lists:

- The *Application Filters* list on the left displays filters that you can select to constrain the applications listed in the Available Applications list.
- The *Available Applications* list in the middle provides applications from which you can select those you want to add as conditions to your rule.
- The *Selected Applications* list on the right displays the applications that you have added to your rule.

Note the following when selecting the filters in the Application Filters list whose applications you want to display in the Available Applications list:

- You can select multiple filters in the Application Filters list under any combination of filter types provided by Sourcefire.

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, for example, the Medium filter contained 110 applications and the High filter contained 82 applications, the system would display 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 would display only those applications that are included in both the Medium or High Risk type AND the Medium or High Business Relevance type.

- You cannot select a custom filter in the Application Filters list in combination with another filter, including another custom filter; this is because you cannot add a filter to a custom filter.
- Selecting one or more filters in the Application Filters list adds an **All apps matching the filter** condition to the Available Applications list. Likewise, searching the Available Applications list when you have not selected any filter in the Application Filters list also adds an **All apps matching the filter** condition to the Available Applications list. 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.

Adding the **All apps matching the filter** condition to the Selected Applications and Filters list counts as one condition against the maximum of 50 conditions, regardless of the number of individual conditions that comprise it.

When you add **All apps matching the filter**, the name of the filter you add 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 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**. For example, add **All apps matching the filter** for the first filter (for example, Risks, High), clear all your selections and make new selections for a different filter type (for example, Business Relevance, High) then add **All apps matching the filter** again.

- When you apply your access control policy, the system generates a single list of unique applications that you have added to the Selected Applications list. This eliminates any duplicate application conditions that you might add.

Adding Application Conditions

LICENSE: Control

SUPPORTED DEVICES: Series 3, virtual

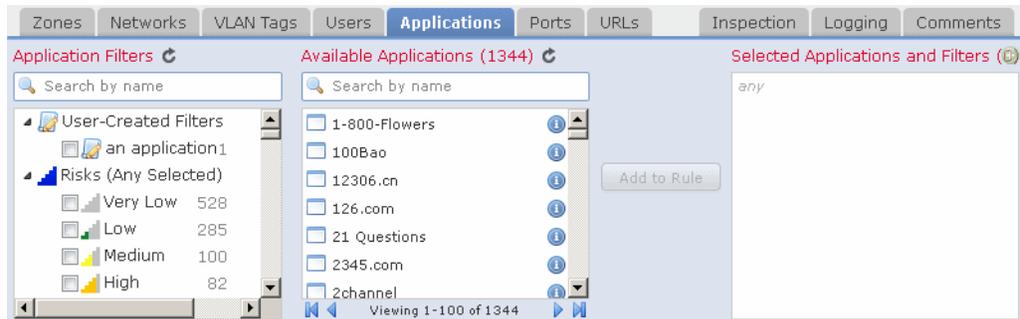
The following procedure explains how to add application conditions while adding or editing an access control rule. See [Understanding Rule Conditions and Condition Mechanics](#) on page 482 for more detailed information.

To add application conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **Applications** tab.

The Applications page appears.



2. Optionally, click the **Search by name** prompt above the Applications Filters list or the Available Applications list, then type a name.

The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.

Note that selecting a custom application filter disables the search fields; this is because you cannot add a filter to a selected custom filter.

3. Optionally, constrain the list of applications displayed in the Available Applications list. You have the following choices:
 - 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.

Note that a number beside each filter indicates the number of applications in the filter.
 - Right-click a Sourcefire-provided filter type (Risks, Business Relevance, Types, Categories, or Tags) and click **Check All** or **Uncheck All**. For more information on Sourcefire-provided filter types, see [the Application Characteristics table](#) on page 1264.

The Available Applications list updates in the following ways:

- The list includes the applications in currently selected filters.
 - An **All apps matching the filter** selection appears; this available condition includes all applications and filters currently displayed in the Available Applications list.

Note that you cannot select or add individual applications in combination with **All apps matching the filter**; note also that, although you can add each separately, doing so will result in duplicate rule conditions. The system reduces duplicate conditions to a single condition when you apply your access control policy.
 - A number above the Available Applications list indicates the number of applications in the currently displayed list.
4. Optionally, click the information icon () next to an application in the Available Applications list.

A pop-up window appears with summary information about the application. You have the following choices:

- To display additional information, click any of the Internet search links provided.
 - To exit the pop-up window and return to the Applications page, click the close icon () or click another location within the Available Applications list.
5. Click an application in the Available Applications list. Use the Shift and Ctrl keys to select multiple applications, or right-click to select all currently displayed applications. Note that you can add a maximum of 50 conditions.
 6. You have the following choices:
 - Click **Add to Rule**.
 - Drag and drop selected conditions into the Selected Applications and Filters list.

Selected conditions are added. Filters appear under the heading *Filters*, and applications appear under the heading *Applications*.

7. Optionally, click the add icon (⊕) above the Selected Applications and Filters list to add a custom filter comprised of all the individual applications and filters currently in the Selected Applications and Filters list.

Custom filters you create from the Applications conditions page or using the object manager appear under the *User-Created Filters* heading in the Application Filters list.

See [Working with Application Filters](#) on page 176 for information on adding application filters using the object manager. See [Using Objects in Conditions](#) on page 491 for information on adding filters from conditions pages.

8. Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Adding Port Conditions

LICENSE: Any

Add a port condition to a rule to match network traffic based on the source and destination port and transport protocol. You can add any of the following kinds of port conditions to an access control rule:

- individual and group port objects that you have created using the object manager
See [Working with Port Objects](#) on page 173 for information on creating individual and group port objects using the object manager.
- individual port objects that you add from the Ports conditions page, and can then add to your rule and to other existing and future rules
See [Using Objects in Conditions](#) on page 491 for more information.
- literal port values, consisting of a transport protocol, a port, or both (for some transport protocol selections)
See [Adding Literal Conditions](#) on page 490 for more information.

The following procedure explains how to add port conditions while adding or editing an access control rule. See [Understanding Rule Conditions and Condition Mechanics](#) on page 482 for more detailed information.

Note that 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 select an ICMP or ICMPv6 type for a port, you can only select a relevant code for the port. 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>.

When you add both source and destination ports to a rule, you can only add port objects or port literals that share a single transport protocol (TCP or UDP) for all ports in the rule. After you add a port to the Selected Source Ports list, you can only add subsequent ports using the same protocol (TCP or UDP) to either port list. Similarly, after you add a destination port, any additional source or destination port you add must have the same protocol. For example, after 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).

If you add only source ports to a rule, you can add ports that use different transport protocols. For example, if a rule has no destination ports, you can add both DNS over TCP and DNS over UDP to the rule. Similarly, if you add only destination ports, you can add destination port literals or port objects using different transport protocols. After you add ports using both protocols to the Selected Source Ports list, you cannot add any ports to the Selected Destination Ports list, and vice versa.

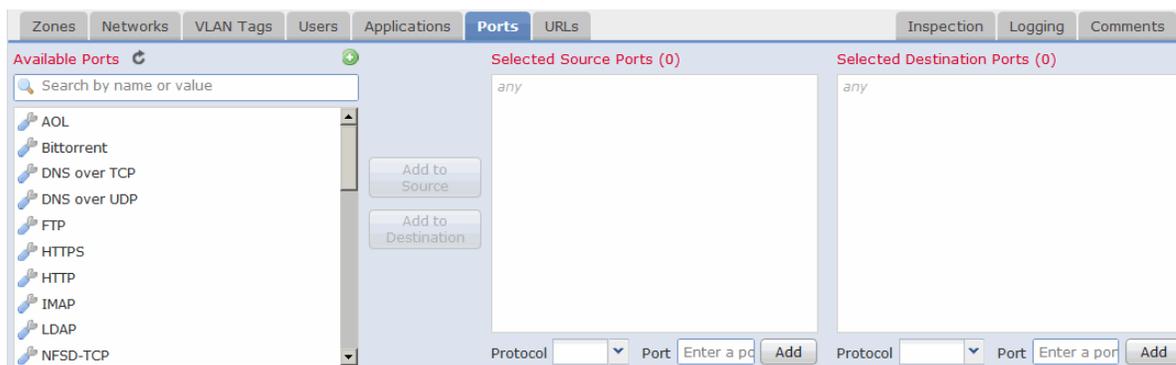
Note that you cannot add a port object or port object group containing a port with a protocol that is invalid for the context. For example, you cannot add an ICMP port object as a source port. If you add a port with an invalid protocol to a port object group already in a rule, a warning displays next to the rule. If you add both source and destination ports, the rule editor requires that all port objects and groups match the protocol specified in the first literal port created in the rule. See [Working with Warnings and Errors](#) on page 455 for more information on how the system omits invalid configurations from the access control policies applied to target devices.

To add port conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **Ports** tab on the rule Edit page.

The Ports page appears.



2. Optionally, click the **Search by name or value** prompt above the Available Ports list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.

3. Click a condition in the Available Ports list. Use the Shift and Ctrl keys to select multiple conditions, or right-click to select all conditions. Note that you can add a maximum of 50 conditions.
Conditions you select are highlighted.
4. You have the following choices:
 - Click **Add to Source** to add the selected port to the Source Ports list.
 - Click **Add to Destination** to add the selected port to the Destination Ports list.
 - Drag and drop available ports into a list.
5. Optionally, to create and add an individual port object click the add icon () above the Available Ports list.
You can identify a single port in each port object that you add. You can then select objects you added as conditions for your rule. See [Working with Port Objects](#) on page 173 and [Using Objects in Conditions](#) on page 491 for more information.
6. Optionally, to add a literal port select an entry from the **Protocol** drop-down list beneath the Selected Source Ports or Selected Destination Ports list.
If you select **TCP**, **UDP** or, for destination ports, **All**, enter a port, if needed; then click **Add**. For destination ports, if you select **ICMP** or **IPv6-ICMP**, a pop-up window appears where you select a type and a related code, if needed, then click **Add**. You can specify a single port with a value from 0 to 65535.
Conditions you selected are added, as long as you add ports with protocols that do not conflict with ports already added.
7. Save or continue editing the rule.
You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Adding URL Conditions

LICENSE: URL Filtering

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: feature dependent

The Sourcefire 3D System 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 Sourcefire cloud by the Defense Center. This feature is called *URL filtering*.

Without a URL Filtering license, you can specify individual URLs or groups of URLs to allow or block. However, before you can perform category and reputation-based URL filtering, you must not only enable URL Filtering licenses on your target devices, but also explicitly allow the Defense Center to contact the Sourcefire cloud. Note that neither the DC500 Defense Center nor Series 2

devices support URL filtering using category and reputation data, and Series 2 devices do not support specifying individual URLs or URL groups. For more information on URL filtering prerequisites, see [Enabling Sourcefire Cloud Communications](#) on page 2032.

Using URL Category and Reputation Data

When the Defense Center contacts the Sourcefire cloud, it retrieves data on many commonly visited URLs, and saves that data to local databases on your appliances. Each of these URLs has an associated category and reputation:

- 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 allow you to quickly create URL conditions for access control rules; these rules can group and combine URL categories and risks. For example, you could create a URL condition that blocks all **High risk** URLs in the **Abused Drugs** category. Then, if someone on your monitored network attempts to browse to any URL with that category and reputation, the session is blocked.

IMPORTANT! To display URL category and reputation data for URLs in connection logs, intrusion events, and application details, you must create at least one access control rule with a URL condition.

Note that if the cloud does not know the category or reputation of a URL, or if the Defense Center cannot contact the cloud, the URL does **not** trigger access control rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.

Also, the system cannot filter URLs before a connection is established between the client and the server. Therefore, when a packet matches all the other conditions in a rule containing a URL condition, if URL identification has not been completed, the packet is allowed to pass. This behavior allows a connection to be established so that URLs can be identified.

When creating a URL condition, selecting a reputation level behaves differently depending on the rule action:

- If you want the rule to block traffic (the rule action is **Block**, **Block with reset**, **Interactive Block**, or **Interactive Block with reset**) 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 you want the rule to allow traffic (the rule action is **Allow**, **Trust**, or **Monitor**), selecting a reputation level also selects all reputations better 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. If you do not specify a reputation level, the system defaults to **any**, meaning all levels. For example, you could block all malware sites, regardless of reputation. You can also use **any** to represent any category. For example, you could block all URLs that match **any** category, but have a **High risk** reputation.

Using Literal URLs, URL Objects, and URL Groups

You are not limited to creating URL conditions using categories and reputations; you can specify individual URLs or groups of URLs to achieve more granular, custom control over allowed and blocked URLs. This is also the only kind of URL filtering you can perform without a special license. You can add either of the following kinds of URL conditions to an access control rule:

- individual URL objects and URL object groups, which represent individual URLs and groups of URLs, respectively; see [Working with URL Objects](#) on page 175

Unlike URL categories, you cannot qualify URL objects and groups with reputations. If the existing URL objects do not meet your needs, you can create a URL object on the fly while creating a URL condition. You can then use the new object in your rule and in other existing and future rules. For more information, see [Using Objects in Conditions](#) on page 491.

- literal URLs; see [Adding Literal Conditions](#) on page 490 for more information

Unlike creating a URL object on the fly, adding a literal URL to an access control rule does not allow you to reuse the URL in other rules.

To determine whether network traffic matches a URL condition, the system performs a simple substring match. If the value of a URL object or literal URL matches any part of a URL requested by a monitored host, the URL condition of

the access control rule is satisfied. 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/

This means that when specifying individual URLs in URL conditions, you must carefully consider other traffic that might be affected. For 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 *verisign.com*, which might not be your intent.

Note that to block HTTPS traffic, you can enter the common name 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.

Choosing a URL Filtering Strategy

When deciding how to build a URL condition, keep in mind that although using URL literals, objects, and groups gives you precise control over allowed and blocked URLs, you must be careful to make sure that your rules do not have unintended consequences.

Alternately, relying on category and reputation data from the Sourcefire cloud gives you less precise control, but simplifies policy creation and administration. It also grants you more assurance that the system will filter URLs as expected. More important, because the cloud is continually updated with new URLs, as well as new categories and risks for existing URLs, using the cloud ensures 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 apply new policies.

For example:

- 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, 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.

Search Query Parameters in URLs

Note that 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.

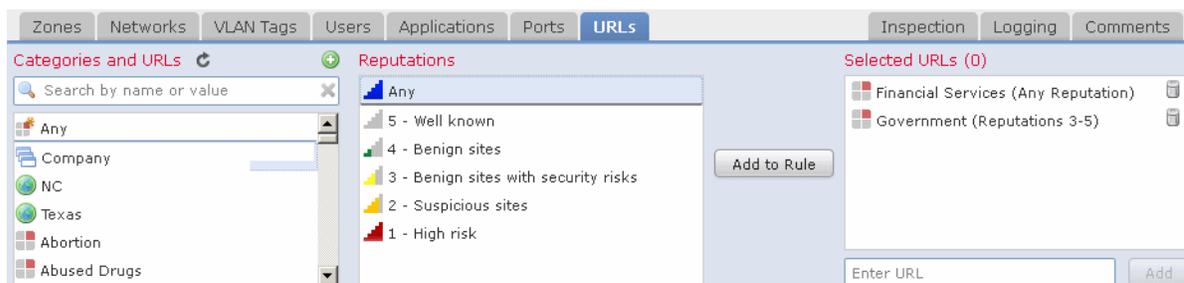
The following procedure explains how to add URL conditions to an access control rule while adding or editing the rule. See [Understanding Rule Conditions and Condition Mechanics](#) on page 482 for more detailed information.

To add URL conditions to an access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select the **URLs** tab.

The URLs page appears.



2. Optionally, click the **Search by name or value** prompt above the Available Users list, then type a name or value.

The list updates as you type to display matching conditions. See [Searching Condition Lists](#) on page 489 for more information.

3. Click a condition in the Categories and URLs list to select the condition. Use the Shift and Ctrl keys to select multiple conditions. To clear selected conditions, click any condition in the list.

Note that selecting all conditions in the Categories and URLs list exceeds the maximum of 50 items you can add to the Selected URLs list.

Conditions you select are highlighted.

4. Optionally, click a reputation level in the Reputations window. Note that you can select only a single reputation level even though you can right-click and then click **Select All** to select **Any**.

The level you selected is highlighted.

5. You have the following choices:

- Click **Add to Rule**.
- Drag and drop selected conditions into the Selected URLs list.

Conditions you selected are added with selected reputation levels appended.

6. Optionally, click the add icon () above the Categories and URLs list to add an individual URL object.
You can specify a single URL in each individual URL object you add. You can then select objects you added as conditions for your rule. See [Working with URL Objects](#) on page 175 and [Using Objects in Conditions](#) on page 491 for more information.
7. Optionally, click the **Enter URL** prompt beneath the Selected URLs list, type a literal URL, then click **Add**.
The list updates to display your entry. See [Adding Literal Conditions](#) on page 490 for more information.
Note that you cannot specify a reputation level for a literal URL.
8. Save or continue editing the rule.
You must apply the access control policy for your changes to take effect; see [Applying an Access Control Policy](#) on page 466.

Performing File and Intrusion Inspection on Allowed Traffic

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

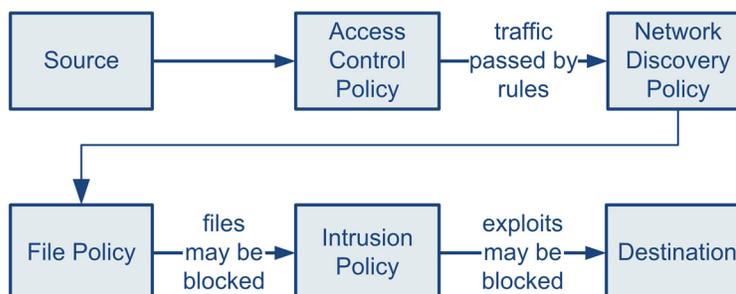
In addition to handling traffic matching the conditions in an access control rule, you can perform further inspection on allowed traffic by associating the rule with an intrusion or file policy.

When you make this association, 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. Depending on your deployment and on policy configurations, both intrusion and file policies can prevent network traffic from reaching its intended destination.

As shown in the diagram below, for traffic that matches an Allow or user-bypassed Interactive Block rule:

- the system automatically performs discovery on the networks listed in the currently applied network discovery policy,
- an optional file policy performs file control and AMP, and
- an optional intrusion policy performs detection and prevention.

Because file inspection occurs before any intrusion policy inspection, blocked files (including malware) are not inspected for intrusion-related exploits.



For more information on Allow and Interactive Block rules, and why only access control rules with those actions can trigger additional inspection, see [Understanding Rule Actions](#) on page 478. Also note that you can associate an intrusion policy, but not a file policy, with the access control default action.

TIP! The system does not perform any kind of inspection on trusted traffic. Although configuring an Allow rule with neither an intrusion nor file policy passes traffic like a Trust rule, Allow rules let you perform discovery on matching traffic.

An access control policy can have multiple access control rules associated with file and intrusion policies, which allows you to match different inspection profiles against different types of traffic on your network.

Note that the number of unique intrusion policies you can use in a single access control policy depends on the model of the target devices; more powerful devices can handle more. The system does not allow you to apply an access control policy if the target devices have insufficient resources to perform inspection. If you attempt to apply an access control policy with more intrusion policies than your device can support, a pop-up window warns that you have exceeded the maximum number of intrusion policies supported by the device.

TIP! If you exceed the number of intrusion policies supported by your device, reevaluate your access control policy. You may want to consolidate intrusion policies so you can associate a single intrusion policy with multiple access control rules.

File Policies and Access Control Rules

A *file policy* is a set of configurations that the system uses to perform file control — that is, to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. With a Malware license, file policies also allow you to inspect a restricted set of those files for malware, and optionally block detected malware. For detailed information on file policies, see [Understanding and Creating File Policies](#) on page 1206.

When you associate a file policy with an access control rule, the Defense Center automatically enables file and malware event logging for that file policy. Sourcefire recommends that you leave this logging setting enabled.

Also, when a file policy generates an event, the system automatically logs the end of the associated connection to the Defense Center database, regardless of any other logging configurations in the invoking access control rule. For more information, see [Logging Connection, File, and Malware Information](#) on page 515.

Note that because you cannot use a Malware license with a DC500, you cannot use that appliance to apply file policies that include rules with the Block Malware or Malware Cloud Lookup action. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy that includes rules with these actions to those appliances.

WARNING! Managed devices **must** be running at least Version 5.1.1 of the Sourcefire 3D System to perform file control and malware detection and at least Version 5.2.0 to block malware.

Intrusion Policies and Access Control Rules

An *intrusion policy* is a set of intrusion detection and prevention configurations that the system uses to analyze network traffic and, optionally, drop offending packets. The system logs intrusion policy violations as intrusion events. For detailed information on intrusion policies, including how to create custom policies, see [Introduction to Sourcefire Intrusion Prevention](#) on page 581 and [Configuring Intrusion Policies](#) on page 667.

When an intrusion policy associated with an access control rule generates an event, the system automatically logs the end of the associated connection to the Defense Center database, regardless of any other logging configurations in the rule. To disable this connection logging on Series 3 or virtual appliances, use the CLI. For more information, see [Logging Connection, File, and Malware Information](#) on page 515.

In contrast, 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 in intrusion detection and prevention-only deployments, where you do not want to log any connection data.

Note, however, if you enable beginning-of-connection logging for the default action, the system **does** log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection. For more information, see [Logging Connections for the Default Action](#) on page 431.

You can associate any of the following intrusion policies with an access control rule.

Sourcefire Authored Policies

Each of these non-modifiable *default* intrusion policies is tuned for a specific balance of security and connectivity. By using a default policy either out-of-the-box or as the basis for a custom policy, you can take advantage of the experience of the Sourcefire Vulnerability Research Team (VRT). For more information, see [Using Default Intrusion Policies](#) on page 692.

WARNING! Do **not** use **Experimental Policy 1** unless instructed to do so by a Sourcefire representative. Sourcefire uses this policy for testing.

User Created Policies

You can select a *custom* intrusion policy that is tailored to inspect the traffic that traverses your network and improve performance in your environment.

In addition to custom policies that you create, Sourcefire provides two custom policies: Initial Inline Policy and Initial Passive Policy. These two policies use the Balanced Security and Connectivity default policy as the base policy. The only difference between them is their **Drop When Inline** setting, which is enabled in the inline policy and disabled in the passive policy. For more information, see [Using a Custom Base Policy](#) on page 693.

The following basic procedure explains how to associate an intrusion or file policy with a new access control rule. See [Creating and Editing Access Control Rules](#) on page 474 for complete instructions on adding and modifying rules.

To associate an intrusion or file policy with a new access control rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon () next to the access control policy you want to modify.
The policy Edit page appears.
3. Click **Add Rule**.
The Add Rule page appears.
4. Ensure the **Action** is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.

5. Select the **Inspection** tab.

The Inspection page appears.



TIP! You can edit file policies and user-created intrusion policies by clicking the edit icon () next to the appropriate drop-down list.

6. Select an **Intrusion Policy**.

Select **None** to disable intrusion inspection for traffic that matches the access control rule.

WARNING! Do **not** select **Experimental Policy 1** unless instructed to by a Sourcefire representative. Sourcefire uses this policy for testing.

7. Select a **File Policy**.

Select **None** to disable file inspection for traffic that matches the access control rule.

8. Click **Add** to save your changes.

The rule is added and the policy Edit page appears.

Logging Connection, File, and Malware Information

LICENSE: Any

For each access control rule in your policies, you must decide whether you want to log connection data for the traffic that matches the conditions in the rule. Tying connection logging to individual rules gives you granular control over the connections you want to log. An access control rule's logging configuration also determines whether you log file and malware events associated with the connection.

TIP! You can log two other types of connection data, outside of access control rules. First, you can log connections handled by the default action. You can also log the decision made by the system to either deny (blacklist) or inspect (blacklist set to monitor-only) a connection based on Security Intelligence data.

Deciding Which Connections to Log

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, only enable logging for the rules critical to your analysis. However, if you want a broad view of your network traffic, you can enable logging for additional access control rules or for the default action.

To optimize performance, Sourcefire recommends that you log either the beginning or the end of the connection, but not both. Note that for a single connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information that was gathered over the duration of the session.

Also, keep in mind that the Sourcefire 3D System uses connection data to display the Connection Summary dashboard, create traffic profiles, trigger correlation rules based on connection data or traffic profile changes, and add connection trackers to correlation rules. For connections that you do not log to the Defense Center database, you cannot take advantage of these features.

You can log connection events to the Defense Center database, as well as to the system log (syslog) or to an SNMP trap server. When and how you can log connections depends on the rule action (see [Understanding Rule Actions](#) on page 478), as summarized in the following table.

Connection Logging Options

RULE ACTION OR LOGGING OPTION	LOG AT:		SEND TO:	
	BEGINNING	END	DEFENSE CENTER	SYSLOG/SNMP
Trust Default Action: Trust	X	X	X	X
Allow Default Action: Intrusion Default Action: Discovery	X	X	X	X
Monitor		X (required)	X (required)	X
Block Block with reset Default Action: Block	X		X	X
Interactive Block Interactive Block with reset	X	X (if bypassed, events show Allow action)	X	X
Security Intelligence	X		X	X

Note that regardless of an access control rule's logging configuration, the system may automatically log connections that contain file or intrusion events; see [Logging Connections Associated with File and Malware Events](#) on page 519 and [Logging Connections Associated with Intrusions](#) on page 519.

Deciding Where to Log or Send Connection Events

When you log a connection event, you can save it to the Defense Center database. The Sourcefire 3D System uses connection data to display the Connection Summary dashboard, create traffic profiles, trigger correlation rules based on connection data or traffic profile changes, and add connection trackers to correlation rules. If you want to take advantage of these features, you **must** log connections to the Defense Center database. For information on database limits, see [Configuring Database Event Limits](#) on page 1978.

You can also log connection events to the syslog or to an SNMP trap server using alert responses. For information on setting up alert responses, see [Working with Alert Responses](#) on page 526.

Logging the Beginning or End of a Connection

Depending on the rule action, you can log a connection event at the beginning or end of a connection, or both. Because matching traffic is denied without further inspection, the system can log only beginning-of-connection events for blocked or Security Intelligence blacklisted traffic.

In general, if you want to perform any kind of detailed analysis on connection data, you should log connection events at the end of connections. This is because beginning-of-connection events do not have 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.

For this reason, the system uses only end-of-connection data to populate connection summaries (see [Understanding Connection Summaries](#) on page 542), which the system then uses to create connection graphs and traffic profiles. Therefore, if you want to use the view connection summaries in custom workflows, view connection data in graphical format, or create and use traffic profiles, you **must** log connection events at the end of connections.

If, however, you simply want to log an event each time the system detects a new connection, beginning-of-connection events are sufficient. You can trigger correlation rules based on either beginning- or end-of-connection events.

Logging Block Rules

Because matching traffic is denied without further inspection, blocking rules can only log beginning-of-connection events. You can, however, configure end-of-connection logging for interactive blocking rules. This is because when the user clicks through the warning page displayed by the system, (see [Adding an HTTP Response Page](#) on page 436), the connection is considered a new, allowed connection which the system can monitor until it terminates.

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 HTTP request is initially blocked; this event has an associated action of **Interactive Block** or **Interactive Block with reset** in the connection log
- 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**

Logging Monitor Rules

For packets that match a Monitor rule, the system always generates a connection event at 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.

Because traffic matching a Monitor rule 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 is one of:

- the action for the first non-Monitor rule triggered by the connection, or
- the default action

The system does **not** generate a separate event each time a single connection matches a Monitor rule. Because a single connection can match multiple Monitor rules, each connection event logged to the Defense Center database can include and display information on up to eight matched monitor rules — the first eight Monitor rules that the connection matches.

Similarly, if you send connection events to the syslog or an 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 first eight Monitor rules that the connection matched.

TIP! Even though the rule action in the connection log can never be **Monitor**, you can still trigger correlation policy violations on connections that match Monitor rules. For more information, see [Specifying Correlation Rule Trigger Criteria](#) on page 1464.

Logging File and Malware Events

When a file policy associated with an access control rule generates a file or malware event, the rule's logging configuration determines whether that event is logged to the database. This setting is automatically enabled, although you can disable it.

File policies can generate the following types of event:

- *file events*, which represent detected or blocked files, including malware files
- *malware events*, which represent the detection or blocking of malware in files evaluated by Malware Cloud Lookup or Block Malware rules
- *retrospective malware events*, which are generated when the malware disposition for a previously detected file changes

When a file policy generates a file or malware event, the system automatically logs the end of the associated connection to the Defense Center database, regardless of the logging configuration of the invoking access control rule.

For more information on performing file inspection, see [Understanding and Creating File Policies](#) on page 1206 and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

Logging Connections Associated with File and Malware Events

Each connection event logged to the Defense Center database can include and display information on the files detected or blocked in a connection. When a file policy generates a file or malware event, the system automatically logs the end of the associated connection to the Defense Center database, regardless of the logging configuration of the invoking access control rule. You cannot disable this connection logging.

For connections where a file was blocked, the associated action in the connection log is **Block** even though you associated the file policy with 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 Connections Associated with Intrusions

Each connection event logged to the Defense Center database can include and display information on the intrusions detected or blocked in a connection. When an intrusion policy associated with an access control rule generates an intrusion event, the system automatically logs the end of the associated connection to the Defense Center database, regardless of the logging configuration of the rule.

TIP! To disable this connection logging on Series 3 or virtual appliances, use the CLI; see [log-ips-connections](#) on page 2270.

For connections where an intrusion was blocked, the associated action in the connection log is **Block**, with a reason of **Intrusion Block**, even though you associated the intrusion policy with an Allow rule.

Note that 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.

For more information, see [Logging Connections for the Default Action](#) on page 431.

Logging the Default Action

The options for logging traffic handled by the policy default action largely parallel the options for logging traffic handled by individual access control rules. For example, if your default action blocks all traffic, you cannot log end-of-connection events for the default action. For more information, see [Logging Connections for the Default Action](#) on page 431.

Logging Security Intelligence Filtering Decisions

Logging blacklisted connections allows you to generate a connection event when the system detects network traffic to or from a blacklisted IP address.

Events generated by Security Intelligence filtering represent the beginning of a connection and the decision made by the system to either deny (blacklist) or inspect (blacklist set to monitor-only) the connection. For these inspected connections, the system may generate additional connection events depending on the logging settings in the access control rule or default action that later handles the connection.

The options for logging Security Intelligence filtering decisions are similar to the options for logging traffic handled by individual access control rules. For detailed information, see [Logging Blacklisted Connections](#) on page 445.

Understanding the Connection Log

The information available for any individual connection event depends on several factors, including the options you set when configuring connection logging. For details, see [Information Available in Connection Events](#) on page 550.

The following procedure explains how to set a new rule to log a connection in traffic that matches the conditions of an access control rule. See [Creating and Editing Access Control Rules](#) on page 474 for complete instructions on adding and modifying rules.

To configure an access control rule to log connection, file, and malware information:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon () next to the access control policy you want to modify.
The policy Edit page appears.
3. Click **Add Rule**.
The Add Rule page appears.

4. Select the **Logging** tab.

The Logging tab appears. The following graphic shows the Logging page for a rule associated with a file policy.



5. Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

You cannot log end-of-connection events for blocked traffic.

6. Use the **Log Files** check box to specify whether the system should log any file and malware events associated with the connection.

Associating a file policy with the rule automatically enables the check box. Sourcefire recommends that you leave this setting enabled.

7. Specify where to send connection events. You have the following choices:

- To send connection events to the Defense Center, select **Defense Center**. When your rule action is **Monitor**, you must log connections to the Defense Center.
- To send connection events to 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](#) on page 530.
- 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](#) on page 528.

8. Click **Add** to save your changes.

The rule is added and the policy Edit page appears.

Adding Comments to a Rule

LICENSE: Any

You can add comments to an access control rule. 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 edit or delete a comment until you save your rule, then you can no longer edit or delete the comment.

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. You can display comments while creating or editing a rule.

Note that you will be prompted to comment when you save changes to a rule if adding a comment is optional or required and you have not already added a comment during the current edit session. See [Configuring Access Control Policy Preferences](#) on page 1970 for more information.

The following basic procedure explains how to add comments to a new rule. See [Creating and Editing Access Control Rules](#) on page 474 for complete instructions on adding and modifying rules.

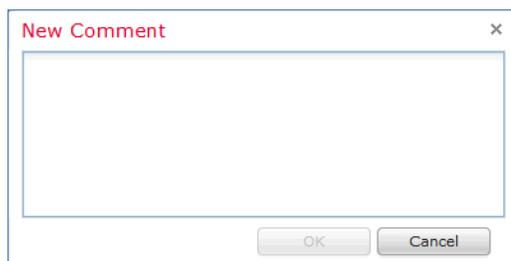
To add a comment to a rule:

ACCESS: Admin/Access Admin/Network Admin

1. Select **Policies > Access Control**.
The Access Control page appears.
2. Click the edit icon (✎) next to the access control policy you want to modify.
The policy Edit page appears.
3. Click **Add Rule**.
The Add Rule page appears.
4. Select the **Comments** tab.
The Comments page appears.



5. Optionally, to add a comment, click **New Comment**.
The New Comment pop-up window appears.



6. Type your comment and click **OK** to add your comment, or click **Cancel** to abandon the comment and return to the Comments page.

Note that you will be prompted to comment when you save changes to a rule if adding a comment is optional or required and you have not already added a comment during the current edit session. See [Configuring Access Control Policy Preferences](#) on page 1970 for more information.

7. Click **Add** to save your changes.
The rule is added and the policy Edit page appears.

CHAPTER 14

CONFIGURING EXTERNAL ALERTING

While the Sourcefire 3D System provides various views of events within the web interface, you may want to configure external event notification to facilitate constant monitoring of critical systems. You can configure the Sourcefire 3D System to generate alerts that notify you via email, SNMP trap, or syslog when one of the following is generated:

- an intrusion event with a specific impact flag
- a specific type of discovery event
- a network-based malware event or retrospective malware event
- a correlation event, triggered by a specific correlation policy violation
- a connection event, triggered by a specific access control rule
- a specific status change for a module in a health policy

To have the system send these alerts, you must first create an *alert response*, which is a set of configurations that allows the Sourcefire 3D System to interact with the external system where you plan to send the alert. Those configurations may specify, for example, an email relay host, 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 impact flags, discovery events, and malware events using their own configuration pages.
- You associate correlation events with alert responses (and remediation responses; see [Creating Remediations](#) on page 1607) in your correlation policies.

- You associate SNMP and syslog alert responses with logged connections using access control rules and policies. Email alerting is not supported for logged connections.
- You associate alert responses with health module status changes using the health monitor.

There is another type of alerting you can perform in the Sourcefire 3D System, which is to configure email, SNMP, and syslog intrusion event notifications for individual intrusion events, regardless of impact flag. You configure these notifications in intrusion policies; see [Configuring External Responses to Intrusion Events](#) on page 1033 and [Adding Alerts](#) on page 761.

The following table explains the licenses you must have to generate alerts.

License Requirements for Generating Alerts

TO GENERATE AN ALERT BASED ON...	YOU NEED THIS LICENSE...
an intrusion event with a specific impact flag	FireSIGHT + Protection
a specific type of discovery event	FireSIGHT
a network-based malware event	Malware
a correlation policy violation	the license that was required to trigger the policy violation
a connection event	the license that was required to log the connection
health module status changes	Any

For more information, see:

- [Working with Alert Responses](#) on page 526
- [Configuring Impact Flag Alerting](#) on page 535
- [Configuring Discovery Event Alerting](#) on page 536
- [Configuring Advanced Malware Protection Alerting](#) on page 537
- [Adding Responses to Rules and White Lists](#) on page 1517
- [Logging Connection, File, and Malware Information](#) on page 515
- [Logging Connections for the Default Action](#) on page 431
- [Configuring Health Monitor Alerts](#) on page 2159

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 Sourcefire 3D System 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. For example, an impact flag alert contains timestamp, intrusion rule, impact flag, and event description information. As another example, discovery event alerts also contain timestamp and description information, as well as discovery event type information.

If you are using an alert response in a correlation policy, the information in the alert depends on the type of event that triggered the correlation policy violation.

IMPORTANT! If you configure an alert as a response to a correlation rule that contains a connection tracker, the alert information you receive is the same as that for alerts on traffic profile changes, even if the correlation rule itself is based on a different kind of event.

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 (**Policies > Actions > Alerts**).

Alerts				
Impact Flag Alerts				
Discovery Event Alerts				
Advanced Malware Protection Alerts				
				 Create Alert
Name	Type	In Use	Enabled	
Sample Email Alert Response	Email	In Use	<input checked="" type="checkbox"/>	 
Sample SNMP Alert Response	SNMP	Not Used	<input checked="" type="checkbox"/>	 
Sample Syslog Alert Response	Syslog	Not Used	<input type="checkbox"/>	 

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 Email Alert Response](#) on page 527
- [Creating an SNMP Alert Response](#) on page 528
- [Creating a Syslog Alert Response](#) on page 530

- [Modifying an Alert Response](#) on page 534
- [Deleting an Alert Response](#) on page 534
- [Enabling and Disabling Alert Responses](#) on page 534

Creating an Email Alert Response

LICENSE: Any

Note that you **cannot** perform email alerting on logged connections in an access control policy.

Before you create an email alert response, you should make sure that the Defense Center can reverse-resolve its own IP address. You should also configure your mail relay host as described in [Configuring a Mail Relay Host and Notification Address](#) on page 1983.

To create an email alert response:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**.

The Alerts page appears.

2. From the **Create Alert** drop-down menu, select **Create Email Alert**.

The Create Email Alert Configuration pop-up window appears.



The screenshot shows a dialog box titled "Create Email Alert Configuration". It has a title bar with a question mark and a close button. The dialog contains four input fields: "Name", "To", "From", and "Relay Host". The "Relay Host" field is pre-filled with "mail.example.com" and has a pencil icon next to it. At the bottom right, there are "Save" and "Cancel" buttons.

3. In the **Name** field, type the name you want to use to identify the alert response.
4. In the **To** field, type the email addresses where you want to send alerts. Separate email addresses with commas.
5. In the **From** field, type the email address that you want to appear as the sender of the alert.

6. Next to **Relay Host**, verify the listed mail server is the one that you want to use to send the alert.

To change the server, or if you have not yet configured a relay host, click the edit icon (✎) to display the System Policy page in a pop-up window, then follow the directions in [Configuring a Mail Relay Host and Notification Address](#) on page 1983. You must apply the system policy after you edit it for your changes to take effect.

7. Click **Save**.
The alert response is saved and is automatically enabled.

Creating an SNMP Alert Response

LICENSE: Any

You can create SNMP alert responses using SNMPv1, SNMPv2, or SNMPv3.

IMPORTANT! If you want to monitor 64-bit values with SNMP, you must use SNMPv2 or SNMPv3. SNMPv1 does not support 64-bit monitoring.

If your network management system requires the Defense Center's management information base (MIB) file, you can obtain it at `/etc/sf/DCEALERT.MIB`.

To create an SNMP alert response:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**.
The Alerts page appears.

- From the **Create Alert** drop-down menu, select **Create SNMP Alert**.
The Create SNMP Alert Configuration pop-up window appears.

The screenshot shows a dialog box titled "Create SNMP Alert Configuration". It contains the following fields and options:

- Name: [Text input field]
- Trap Server: [Text input field]
- Version: [Dropdown menu, currently showing v3]
- User Name: [Text input field]
- Authentication:
 - Protocol: [Dropdown menu, currently showing None]
 - Password: [Text input field]
- Privacy:
 - Protocol: [Dropdown menu, currently showing None]
 - Password: [Text input field]
- Engine ID: [Text input field, with placeholder text "Example: 123456789a"]

Buttons: Save, Cancel

- In the **Name** field, type the name that you want to use to identify the SNMP response.
- 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.
- 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.

The screenshot shows the same dialog box, but with the Version dropdown menu set to v1. The Community String field is now visible, and the Authentication and Privacy sections are hidden. The buttons remain the same.

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.
7. From the **Authentication Protocol** drop-down list, select the protocol you want to use for authentication.
8. In the **Authentication Password** field, type the password required for authentication with the SNMP server.
9. From the **Privacy Protocol** list, select **None** to use no privacy protocol or **DES** to use Data Encryption Standard as the privacy protocol.
10. In the **Privacy Password** field, type the privacy password required by the SNMP server.
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.

Sourcefire recommends that you use the hexadecimal version of the Defense Center's IP address. For example, if the Defense Center has an IP address of 10.1.1.77, use 0a01014D0.
12. Click **Save**.

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.

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.

Available Syslog Facilities (Continued)

FACILITY	DESCRIPTION
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.

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.

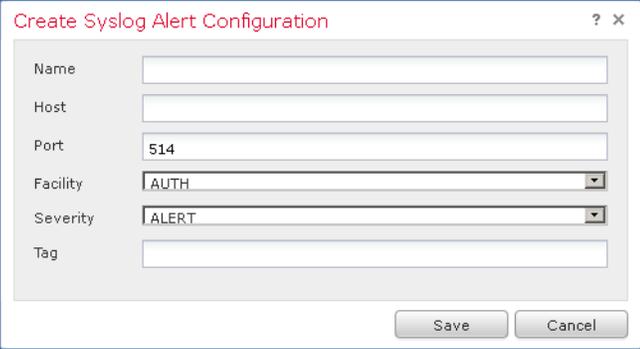
To create a syslog alert:

ACCESS: Admin

1. Select **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.



The screenshot shows a dialog box titled "Create Syslog Alert Configuration". It has a title bar with a question mark and a close button. The dialog contains the following fields and controls:

- Name:** A text input field.
- Host:** A text input field.
- Port:** A text input field containing the value "514".
- Facility:** A dropdown menu with "AUTH" selected.
- Severity:** A dropdown menu with "ALERT" selected.
- Tag:** A text input field.
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

2. In the **Name** field, type the name you want to use to identify the saved response.
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.
4. In the **Port** field, type the port the server uses for syslog messages.
By default, this value is 514.
5. From the **Facility** list, select a facility.
See the [Available Syslog Facilities table](#) on page 531 for a list of the available facilities.
6. From the **Severity** list, select a severity.
See the [Syslog Severity Levels table](#) on page 532 for a list of the available severities.
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 FromDC, type FromDC in the field.
8. Click **Save**.
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:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**.
The Alerts page appears.
2. Next to the alert response you want to edit, click the edit icon (✎).
A configuration pop-up window for that alert response appears.
3. Make changes as needed.
4. Click **Save**.
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:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**.
The Alerts page appears.
2. Next to the alert response you want to delete, click the delete icon (🗑️).
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:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**.

The Alerts page appears.

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 Impact Flag Alerting

LICENSE: Protection

You can configure the system to alert you whenever an intrusion event with a specific impact flag occurs. Impact flags help you evaluate the impact an intrusion has on your network by correlating intrusion data, network discovery data, and vulnerability information. For more information, see [Using Impact Flags to Evaluate Events](#) on page 641.

To configure impact flag alerting:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**, then select the **Impact Flag Alerts** tab.

The Impact Flag Alerts page appears.

The screenshot shows the 'Alerts' configuration page. At the top, there are three dropdown menus for 'Syslog', 'Email', and 'SNMP', all set to 'None'. Below these is the 'Impact Flag Configuration' section, which includes a table with columns for 'Impact Flag', 'Syslog Notification', 'Email Notification', and 'SNMP Notification'. Each row in the table has a checkbox for each notification type. The rows are: 'Unknown', 'Unknown Target', 'Currently Not Vulnerable', 'Potentially Vulnerable', and 'Vulnerable'. A 'Save' button is located at the bottom right of the table.

Impact Flag	Syslog Notification	Email Notification	SNMP Notification
0 Unknown	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4 Unknown Target	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 Currently Not Vulnerable	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2 Potentially Vulnerable	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
1 Vulnerable	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

2. In the Alerts section, select the alert response you want to use for each alert type.
To create a new alert response, select **New** from any drop-down list. For more information, see [Working with Alert Responses](#) on page 526.
3. In the Impact Configuration section, select the check boxes that correspond to the alerts you want to receive for each impact flag.
4. Click **Save**.
Your impact flag alerting settings are saved.

Configuring Discovery Event Alerting

LICENSE: FireSIGHT

You can configure the system to alert you whenever a specific type of discovery event occurs. For information about the different event types, see [Understanding Discovery Event Types](#) on page 1389 and [Understanding Host Input Event Types](#) on page 1394.

Note that to generate an alert based on a discovery event type, you must configure your network discovery policy to log that event type; see [Configuring Discovery Event Logging](#) on page 1295. By default, logging is enabled for all event types.

To configure discovery event alerting:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**, then select the **Discovery Event Alerts** tab.
The Discovery Event Alerts page appears.

Alerts

Select alert responses to use for discovery event alerts. You must click Save after you make your selections.

Syslog

Email

SNMP

Events Configuration

Select the event types for which you want to generate alerts.

Event	Syslog Notification	Email Notification	SNMP Notification
Add Client	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Add Host	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Vulnerability Set Invalid	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Vulnerability Set Valid	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Save

2. In the **Alerts** section, select the alert response you want to use for each alert type.
To create a new alert response, select **New** from any drop-down list. For more information, see [Working with Alert Responses](#) on page 526.
3. In the **Events Configuration** section, select the check boxes that correspond to the alerts you want to receive for each discovery event type.
4. Click **Save**.
Your discovery event alerting settings are saved.

Configuring Advanced Malware Protection Alerting

LICENSE: Malware

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

You can configure the system to alert you whenever any network-based malware event, including a retrospective event, is generated. You cannot, however, alert on endpoint-based (FireAMP) malware events. For information on malware events, see [Working with Malware Events](#) on page 1229.

To generate alerts based on malware events, you must create a file policy that performs malware cloud lookups, then associate that policy with an access control rule. For more information, see [Understanding and Creating File Policies](#) on page 1206 and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

To configure malware event alerting:

ACCESS: Admin

1. Select **Policies > Actions > Alerts**, then select the **Advanced Malware Protections Alerts** tab.

The Advanced Malware Protection Alerts page appears.

Alerts

Select alert responses to use for advanced malware detection event alerts. You must click Save after you make your selections.

Syslog:
Email:
SNMP:

Event Configuration

Select the even types for which you want to generate alerts.

Event	Syslog	Email	SNMP
Retrospective Events	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
All network-based malware events	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

2. In the **Alerts** section, select the alert response you want to use for each alert type.

To create a new alert response, select **New** from any drop-down list. For more information, see [Working with Alert Responses](#) on page 526.

3. In the **Event Configuration** section, select the check boxes that correspond to the alerts you want to receive for each malware event type.

Keep in mind that **All network-based malware events** includes **Retrospective Events**.

4. Click **Save**.

Your malware event alerting settings are saved.

CHAPTER 15

WORKING WITH CONNECTION DATA

Sourcefire managed devices continuously monitor traffic generated by the hosts on your network. You can use the access control feature to generate *connection events* when network traffic matches specific conditions. Connection events contain data about the detected sessions, including timestamps, IP addresses, geolocation, applications, and so on.

Logging connection events to the Defense Center database allows you to take advantage of the analysis, reporting, and correlation features in the Sourcefire 3D System. Optionally, you can send most connection events to the syslog or an SNMP trap server.

To supplement the connection data gathered by your managed devices, you can use records generated by NetFlow-enabled devices to generate connection events. This is especially useful if you have NetFlow-enabled devices deployed on networks that your Sourcefire managed devices cannot monitor.

To further enhance the geolocation information provided with many connection events, you can configure geolocation updates for your system. For more information on geolocation, see [Using Geolocation](#) on page 1816.

For more information, see:

- [Understanding Connection Data](#) on page 540
- [Viewing Connection Data Graphs and Tables](#) on page 555
- [Working with Connection Graphs](#) on page 555
- [Working with Connection Data Tables](#) on page 570
- [Searching for Connection Data](#) on page 575

- [Viewing the Connection Summary Page](#) on page 578
- [Understanding NetFlow](#) on page 1271

Understanding Connection Data

LICENSE: Any

For networks monitored by Sourcefire managed devices, you can configure and apply access control policies to log connection events when:

- network traffic is blacklisted or monitored by Security Intelligence
- network traffic meets the conditions of a non-Monitor access control rule
- network traffic is handled by an access control policy's default action
- network traffic meets the conditions of at least one Monitor rule (automatically enabled)
- an intrusion policy associated with an access control rule generates an event (automatically enabled)
- a file policy associated with an access control rule detects or blocks a file, or discovers or blocks malware (automatically enabled)

Tying connection logging to individual access control rules, policies, and configurations gives you granular control over the connections you want to log.

Note that because NetFlow data collection is not linked to access control rules, you do not have granular control over which NetFlow connections you want to log. Sourcefire managed devices detect records exported by NetFlow-enabled devices, generate unidirectional end-of-connection events based on the data in those records, and finally send those events to the Defense Center to be logged in the database. You cannot send NetFlow events to the system log or an SNMP trap server.

For more information on connection logging, see the following sections:

- [Logging Connection, File, and Malware Information](#) on page 515 explains how to log traffic that meets the conditions of an access control rule, and also contains general guidance on when and how to log those connections. This section also explains how connection logging is affected by the rule action, and how connection event logging relates to intrusion, file, and malware event logging.
- [Logging Blacklisted Connections](#) on page 445 explains how to use the Security Intelligence feature to log the decision to deny (blacklist) or inspect (blacklist set to monitor-only) connections.
- [Logging Connections for the Default Action](#) on page 431 explains how to log connections handled by an access control policy's default action.

- [Understanding NetFlow](#) on page 1271 provides more information on NetFlow, and compares NetFlow connection events with connection events based on traffic monitored by the Sourcefire 3D System.
- [Creating a Network Discovery Policy](#) on page 1275 explains how to create and manage your discovery policy, which is also where you configure NetFlow data collection.

The following table explains the licenses you must have to log connection data.

License Requirements for Logging Connection Data

To...	YOU NEED THIS LICENSE...
perform basic connection logging, including NetFlow connection logging	Any
add data to the network map, including host and user data, based on the information in connection logs; view geolocation information associated with connection events	FireSIGHT
log connections: <ul style="list-style-type: none"> • that represent Security Intelligence filtering decisions • in an access control rule that performs intrusion detection and prevention • in an access control rule that performs file control, but not advanced malware protection 	Protection
log connections in an access control rule that performs advanced malware protection	Malware
log connections in an access control rule that performs application or user control	Control
log connections in an access control rule with URL conditions that use URL category and reputation data display URL category and URL reputation information for URLs requested by monitored hosts	URL Filtering

The following sections provide additional details on the kinds of information available about detected connections, as well as how you log, aggregate, and use connection data as part of your analysis.

- [Understanding Connection Summaries](#) on page 542
- [Connection Data Fields](#) on page 543
- [Information Available in Connection Events](#) on page 550
- [Uses for Connection Data in the Sourcefire 3D System](#) on page 554

Understanding Connection Summaries

LICENSE: Any

The Sourcefire 3D System aggregates connection data collected over five-minute intervals into connection summaries, which the system uses to generate connection graphs and traffic profiles. Optionally, you can create custom workflows based on connection summary data, which you use in the same way as you use workflows based on individual connection events.

To be aggregated, multiple connections must:

- represent the end of connections
- have the same source and destination IP addresses, and use the same port on the responder (destination) host
- use the same protocol (TCP or UDP)
- use the same application protocol
- either be detected by the same Sourcefire managed device, or be exported by the same NetFlow-enabled device

Each connection summary includes total traffic statistics, as well as the number of connections in the summary. Because NetFlow-enabled devices generate unidirectional connections, a summary's connection count is incremented by two for every connection based on NetFlow data.

Note that connection summaries do not contain all of the information associated with its aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

For more information, see the following sections:

- [Long-Running Connections](#) on page 542
- [Combined Connection Summaries from External Responders](#) on page 543
- [Information Available in Connection Events](#) on page 550

Long-Running Connections

LICENSE: Any

If a monitored session spans two or more five-minute intervals over which connection data is aggregated, the connection is considered a *long-running connection*. When calculating the number of connections in a connection summary, the system increments the count only for the five-minute interval in which a long-running connection was initiated.

Also, when calculating the number of packets and bytes transmitted by the initiator and responder in a long-running connection, the system does not report the number of packets and bytes that were actually transmitted during each five-minute interval. Instead, the system assumes a constant rate of transmission and calculates estimated figures based on the total number of packets and bytes

transmitted, the length of the connection, and what portion of the connection occurred during each five-minute interval.

Combined Connection Summaries from External Responders

LICENSE: Any

To reduce the space required to store connection data and speed up the rendering of connection graphs, the system combines connection summaries when:

- one of the hosts involved in the connection is not on your monitored network
- other than the IP address of the external host, the connections in the summaries meet the aggregation criteria listed in [Understanding Connection Summaries](#) on page 542: protocol, application protocol, detecting device, and so on

When viewing connection summaries in the event viewer and when working with connection graphs, the system displays `external` instead of an IP address for the non-monitored hosts.

As a consequence of this aggregation, if you attempt to drill down to the table view of connection data (that is, access data on individual connections) from a connection summary or graph that involves an external responder, the table view contains no information.

Connection Data Fields

LICENSE: Any

Each connection table view or connection graph contains information about the connections or connection summaries you are viewing, including timestamps, IP addresses, geolocation information, applications, and so on.

The following list details the connection data logged by the Sourcefire 3D System. For a discussion of the factors that determine the information logged in any individual connection event, see the next section: [Information Available in Connection Events](#) on page 550.

Time

The ending time of the five-minute interval that the system used to aggregate connections in a connection summary.

First Packet or Last Packet

The date and time the first or last packet of the session was seen.

Action

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. Note that TCP connections detected by a trust rule on the first packet will not generate a beginning-of-connection event. On virtual appliances, generated events will be reported 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**.

For Security Intelligence-monitored connections, the action is that of the first non-Monitor access control rule triggered by the connection. 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**.

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 Resume Allow** indicates that file transmission was originally blocked by a Block Files or Block Malware file rule, then the HTTP session automatically resumed after a new access control policy was applied that allowed the file.
- **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**.

Initiator IP or Responder IP

The host IP address (and host name, if DNS resolution is enabled) that initiated, or responded to, the session responder. So that you can identify the blacklisted IP address in a blacklisted connection, host icons next to blacklisted IP addresses look slightly different.

Initiator Country or Responder Country

When a routable IP is detected, the country associated with the host IP address that initiated the session, or with the session responder. An icon of the country's flag is displayed, as well as the country's ISO 3166-1 alpha-3 country code. Hover your pointer over the flag icon to view the country's full name.

Note that the DC500 Defense Center does not support this feature.

Initiator User

The user logged into the session initiator.

Security Intelligence 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 Sourcefire Intelligence Feed. Note that this field is only populated if the **Reason** is **IP Block** or **IP Monitor**. For more information, see [Filtering Traffic Based on Security Intelligence Data](#) on page 438.

Note that neither the DC500 Defense Center nor Series 2 devices support this feature.

Ingress Security Zone or Egress Security Zone

The ingress or egress security zone associated with the connection.

Source Port/ICMP Type or Destination Port/ICMP Code

The port, ICMP type, or ICMP code used by the session initiator or session responder.

Application Protocol

The application protocol, which represents communications between hosts, detected in the connection.

Client and Client Version

The client application and version of that client 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`.

Web Application

The web application, which represents the content or requested URL for HTTP traffic detected in the connection. If the system cannot identify the specific web application in HTTP traffic, this field displays `Web Browsing`.

Category, Tag (Application Protocol, Client, Web Application)

Criteria that characterize the application to help you understand the application's function. For more information, see the [Application Characteristics table](#) on page 1264.

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. For more information, see the [Application Characteristics table](#) on page 1264.

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. For more information, see the [Application Characteristics table](#) on page 1264.

URL, URL Category, and URL Reputation

The URL requested by the monitored host during the session and its associated category and reputation, if available.

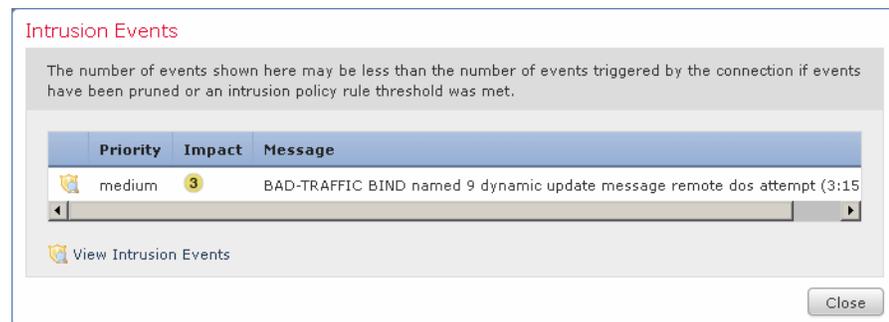
If the system identifies or blocks an SSL application, the requested URL is in encrypted traffic, so the system identifies the traffic based on an SSL certificate. For SSL applications, therefore, this field indicates the common name contained in the certificate. For more information see [Configuring Advanced Access Control Policy Settings](#) on page 447 and the [Enabling Sourcefire Cloud Communications table](#) on page 2032.

Note that neither the DC500 Defense Center nor Series 2 devices support URL category or reputation data.

Intrusion Events

The intrusion events, if any, associated with the connection. Instead of a list of events, the Defense Center displays the view intrusion events icon () in this field.

Click the icon to display a pop-up window with a list of intrusion events associated with the connection, as well as their priority and impact.

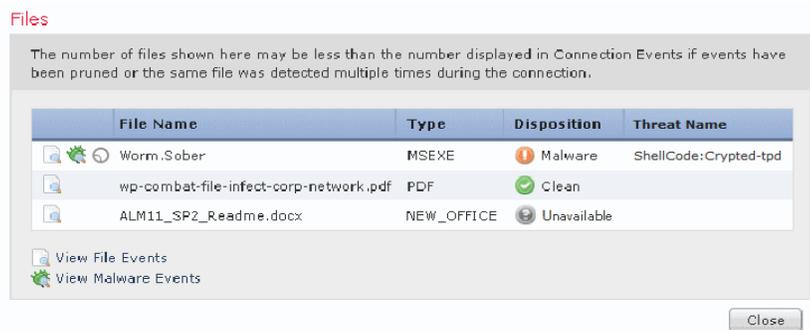


For more information, see [Viewing Intrusion Events Associated with a Connection](#) on page 574.

Files

The file events, if any, associated with the connection. Instead of a list of files, the Defense Center displays the view files icon () in this field. The number on the icon indicates the number of files (including malware files) detected or blocked in that connection.

Click the icon to display a pop-up window with a list of the files detected in the connection, as well as their types and if applicable, their malware lookup dispositions.



Note that neither the DC500 Defense Center nor Series 2 devices support network-based malware file detection.

For more information, see [Viewing Files Detected in a Connection](#) on page 572.

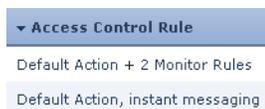
Access Control Policy

The access control policy that contains the access control rule or default action that logged the connection.

Access Control Rule

The access control rule or default action that handled the connection, as well as up to eight Monitor rules matched by that connection.

If the connection matched one Monitor rule, the Defense Center displays the name of the rule that handled the connection, followed by the Monitor rule name. If the connection matched more than one Monitor rule, the event viewer displays how many Monitor rules it matched, for example, **Default Action + 2 Monitor Rules**.



To display a pop-up window with a list of the first eight Monitor rules matched by the connection, click **N Monitor Rules**.



Device

The managed device that detected the connection, or for connections exported by NetFlow-enabled devices, the managed device that processed the NetFlow data.

Ingress Interface or Egress Interface

The ingress or egress interface associated with the connection.

TCP Flags

The TCP flags detected in the connection.

Source Device

The IP address of the NetFlow-enabled device that exported the data for the connection. If the connection was detected by a managed device, this field contains a value of `FireSIGHT`.

NetBIOS Domain

The NetBIOS domain used in the session.

Initiator Packets or Responder Packets

The total number of packets transmitted by the session initiator or the session responder.

Initiator Bytes or Responder Bytes

The total number of bytes transmitted by the session initiator or the session responder.

Connections

The number of connections in the summary. For long-running connections, that is, connections that span multiple connection summary intervals, only the first connection summary interval is incremented.

Count

The number of connections that match the information that appears in each row. Note that the **Count** field appears only after you apply a constraint that creates two or more identical rows.

IMPORTANT! If you create a custom workflow and do not add the **Count** column to a drill-down page, each connection is listed individually and packets and bytes are not summed.

Information Available in Connection Events

LICENSE: Any

The information available for any individual connection or connection summary event depends on several factors.

Detection Method

With one exception (TCP flags), the information available in NetFlow records is more limited than the information generated by monitoring network traffic using managed devices. For more information, see the [Differences Between NetFlow and FireSIGHT Data table](#) on page 1272.

Logging Method

For connections detected directly by Sourcefire managed devices, you can log a connection event at the beginning or end of a connection, or both — depending on the access control rule action, default action, or Security Intelligence blacklist. NetFlow-based connections are considered end-of-connection.

Beginning-of-connection events do not have information that must be 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). Beginning-of-connection events are also not guaranteed to have information about application or URL traffic in the session.

Associated File and Intrusion Policies

Only connections logged by access control rules with associated file policies contain file information. Similarly, you must associate intrusion policies with either access control rules or the default action to view intrusion information in the connection log.

Note that connections detected by pre-5.1.1 devices are not associated with file or intrusion information. Managed devices must be running at least Version 5.1.1 of the Sourcefire 3D System to perform file control and network-based malware detection and at least Version 5.2 to perform malware blocking. For more information, see [Performing File and Intrusion Inspection on Allowed Traffic](#).

Connection Event Type

Connection summaries do not contain all of the information associated with their aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Keep in mind that connection graphs are based on connection summary data, which use only end-of-connection logs. If you logged only beginning-of-connection data, connection graphs and connection summary event views contain no data.

Traffic Type

The system only reports information present in the traffic. For example, non-HTTP traffic does not contain information on URLs or web applications. Or, there could be no user associated with the initiator host.

Other Configurations

An advanced setting in the access control policy controls the number of characters the system stores in the connection log for each URL requested by monitored hosts in HTTP sessions. If you use this setting to disable URL logging, the system does not display individual URLs in the connection log, although you can still view category and reputation data, if it exists.

Also, not all connection events have a **Reason**, which is a field populated only in specific situations, such as when a user bypasses an Interactive Block configuration; see [Reason](#) on page 544.

Appliance Model

Because Series 2 devices and the DC500 Defense Center support feature subsets, the DC500 does not display and Series 2 devices do not detect or provide the following connection data:

- Security Intelligence data
- URL category or reputation data
- File data associated with network-based malware detection

Additionally, because the DC500 Defense Center does not support geolocation data, it does not display the event initiator or responder country.

See [Series 2 Appliances](#) on page 40 for a summary of Series 2 appliance features.

The following table lists each connection event field and whether the system displays information in that field, depending on the detection method, logging method, and connection event type.

TIP! In the table view of connection events, the **Source Device** field, as well as the **Category** and **Tag** fields for each type of application are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under Disabled Columns.

Connection Data Based on Logging and Detection Methods

FIELD	DETECTION METHOD:		LOGGING METHOD:		CONNECTION EVENT:	
	FIRE SIGHT	NETFLOW	START	END	SINGLE	SUMMARY
Time	X	X		X		X
First Packet	X	X	X	X	X	
Last Packet	X	X		X	X	
Action	X		X	X	X	
Reason	X		X	X	X	
Initiator IP	X	X	X	X	X	X
Initiator Country	X		X	X	X	X
Initiator User	X	X	X	X	X	X
Responder IP	X	X	X	X	X	X
Responder Country	X		X	X	X	X
Security Intelligence Category	X		X		X	
Ingress Security Zone	X		X	X	X	X
Egress Security Zone	X		X	X	X	X
Source Port/ICMP Code	X	X	X	X	X	
Destination Port/ICMP Type	X	X	X	X	X	X

Connection Data Based on Logging and Detection Methods (Continued)

FIELD	DETECTION METHOD:		LOGGING METHOD:		CONNECTION EVENT:	
	FIRE SIGHT	NETFLOW	START	END	SINGLE	SUMMARY
Application Protocol	X	X	if available	X	X	X
Client	X		if available	X	X	
Client Version	X		if available	X	X	
Web Application	X		if available	X	X	
Category, Tag (Application Protocol, Client, Web Application)	X		if available	X	X	
Application Risk	X		if available	X	X	
Business Relevance	X		if available	X	X	
URL	X		if available	X	X	
URL Category	X		if available	X	X	
URL Reputation	X		if available	X	X	
Intrusion Events	X			X	X	
Files	X			X	X	
Access Control Policy	X		X	X	X	
Access Control Rule	X		X	X	X	
Device	X	X	X	X	X	X
Ingress Interface	X		X	X	X	X
Egress Interface	X		X	X	X	X
TCP Flags		X		X	X	
Source Device	X	X	FireSIGHT	X	X	X
NetBIOS Domain	X		X	X	X	

Connection Data Based on Logging and Detection Methods (Continued)

FIELD	DETECTION METHOD:		LOGGING METHOD:		CONNECTION EVENT:	
	FIRE SIGHT	NETFLOW	START	END	SINGLE	SUMMARY
Initiator Packets	X	X	not useful	X	X	X
Responder Packets	X	X	not useful	X	X	X
Initiator Bytes	X	X	not useful	X	X	X
Responder Bytes	X	X	not useful	X	X	X
Connections	X	X		X		X
Count	X	X	X	X	X	

Uses for Connection Data in the Sourcefire 3D System

LICENSE: Any

Logging connection events to the Defense Center database allows you to take advantage of many features in the Sourcefire 3D System, including:

- viewing the Connection Summary dashboard, which provides you with an at-a-glance view of the connections logged by the system; see [Using Dashboards](#) on page 66
- viewing detailed information on the connections logged by the system, which you can display in a graphical or tabular format; see [Viewing Connection Data Graphs and Tables](#) on page 555
- creating reports based on the connections logged by the system; see [Working with Reports](#) on page 1724
- using connection data to create and view a profile of your normal network traffic, called a traffic profile; see [Creating Traffic Profiles](#) on page 1585
- creating correlation rules that trigger and generate correlation events when the system detects certain connection data, or when a traffic profile changes; see [Creating Rules for Correlation Policies](#) on page 1461
- adding connection trackers to correlation rules, so that after the rule's initial criteria are met, the system begins tracking certain connections and only generates a correlation event if the tracked connections meet additional criteria; see [Adding a Connection Tracker](#) on page 1486

Viewing Connection Data Graphs and Tables

LICENSE: Any

To help you gain in-depth insight to connection data, the system can present connection data both graphically and in a tabular format. The page you see when you access connection data differs depending on the workflow you use. You can use one of the predefined workflows or create a custom workflow that displays only the information that matches your specific needs.

Each connection table view or connection graph contains information about the connections or connection summaries you are viewing, including timestamps, IP addresses, applications, and so on. The information available for any individual connection detected by the Sourcefire 3D System depends on several factors, including detection method and logging options. For more information, see [Connection Data Fields](#) on page 543 and [Information Available in Connection Events](#) on page 550.

TIP! The Connection Summary dashboard can provide you with an at-a-glance view of the connections logged by the system. For more information, see [Using Dashboards](#) on page 66.

To view connection data:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Connections**.

The first page of the default connection data workflow appears. There are two possibilities:

- The workflow page displays a **graph**. See [Working with Connection Graphs](#) on page 555 for information on the actions you can perform.
- The workflow page displays a **table**. See [Working with Connection Data Tables](#) on page 570 for information on the actions you can perform.

To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

Working with Connection Graphs

LICENSE: Any

One of the ways the system can present connection data is graphically. There are three different types of connection graphs: line graphs, bar graphs, and pie charts. Bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point.

You can manipulate connection graphs in various ways, including:

- changing the type of data that the graph displays
- switching between graph types
- constraining the graph so it shows data for specific time ranges, hosts, applications, ports, and devices

Because traffic profiles are based on connection data (see [Creating Traffic Profiles](#) on page 1585), you can view traffic profiles as line graphs. You can manipulate these graphs in the same way as you would any other connection graph, with some restrictions.

IMPORTANT! To view traffic profiles, you must have Administrator access. Compare this with other connection graphs, which you can view with any Security Analyst or Administrator access.

When you view a connection graph, as described in [Viewing Connection Data Graphs and Tables](#) on page 555, you can perform the basic actions described in the following table.

ACCESS: Admin/Any Security Analyst

Basic Connection Graph Functions

To...	You CAN...
learn more about the data that appears	find more information in Connection Data Fields on page 543.
modify the time and date range	find more information in Setting Event Time Constraints on page 1820.
view a host's profile	on a graph displaying connection data by initiator or responder, click either a bar on a bar graph or a wedge on a pie chart and select View Host Profile .
use a different workflow, including a custom workflow	click (switch workflow) by the workflow title.
navigate between pages in the current workflow	find more information in Using Workflow Pages on page 1812.
navigate to other event views to view associated events	find more information in Navigating Between Workflows on page 1836.

There are many other ways you can manipulate connection graphs as you perform in-depth analysis of connection data. For more information, see:

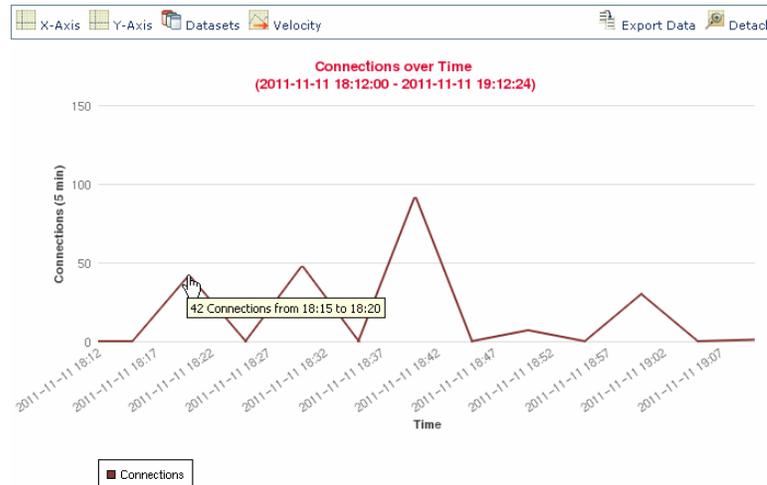
- [Changing the Graph Type](#) on page 557 explains how to change between bar graphs and pie chart, and between standard line graphs and velocity graphs.
- [Selecting Datasets](#) on page 560 explains how to display several values on the y-axis for each x-axis data point on line graphs and bar graphs.
- [Viewing Information About Aggregated Connection Data](#) on page 563 explains how to get more information about the data points on a graph, or to display the host profile of a host whose statistics are being graphed.
- [Manipulating a Connection Graph on a Workflow Page](#) on page 563 explains how to constrain the data that appears on a connection graph without advancing the workflow to the next page.
- [Drilling Down Through Connection Data Graphs](#) on page 564 explains how to constrain the data that appears on a connection graph while advancing the workflow to the next page.
- [Recentering and Zooming on Line Graphs](#) on page 565 explains how to recenter a line graph around any point in time.
- [Selecting Data to Graph](#) on page 565 explains how to change the data displayed on a connection graph by changing its x- or y-axis.
- [Detaching Connection Graphs](#) on page 569 explains how you can detach a connection graph into a new browser window and perform further analysis without affecting the default time range for the Defense Center.
- [Exporting Connection Data](#) on page 569 explains how to export the connection data used to construct a graph as a CSV (comma-separated values) file.

Changing the Graph Type

LICENSE: Any

There are three different connection graphs: line graphs, bar graphs, and pie charts. *Line graphs* plot data over time. For example, the following line graph

displays the total number of connections detected on a monitored network over a one-hour time span. Traffic profiles are always displayed as line graphs.

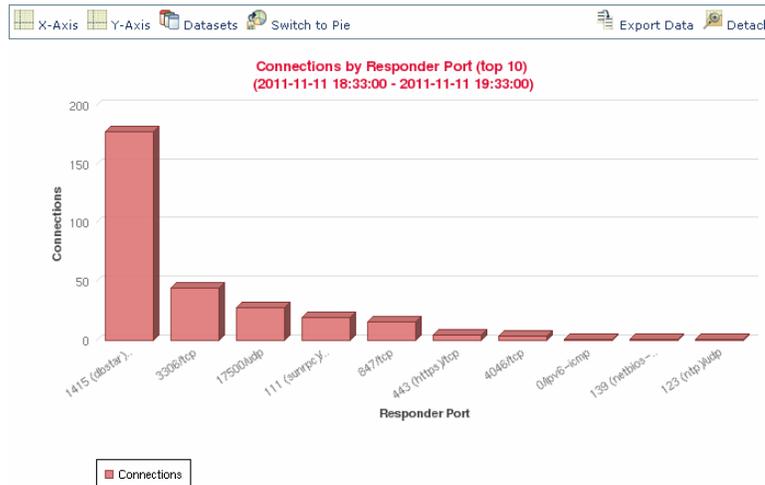


By default, line graphs appear in *standard view*. A standard line graph aggregates data over five minute intervals, plots the aggregated data points, and connects the points.

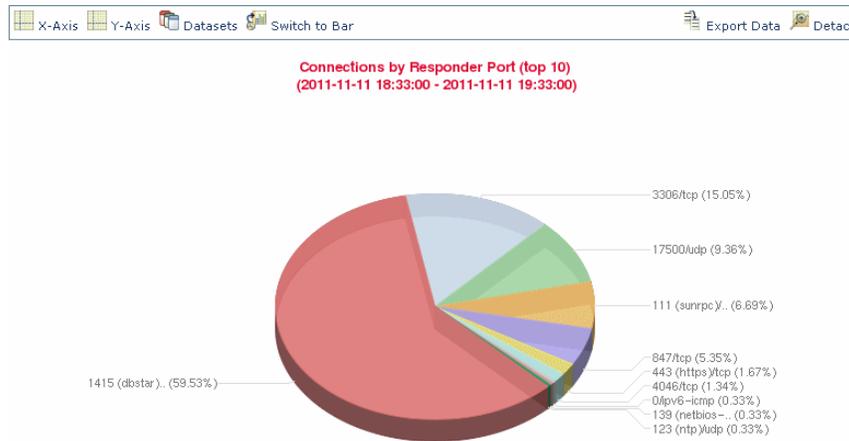
However, you can change a line graph from standard view to *velocity view*. A velocity line graph shows the rate of change between those data points. If you change the above graph to a velocity graph, the y-axis changes from indicating the number of connections to indicating the change in the number of connections over time.



Bar graphs display data grouped into discrete categories. For example, a bar graph could show the number of connections detected on a monitored network for the 10 most active ports over a one-hour time span.



Pie charts, like bar graphs, also display data grouped into discrete categories. The following pie chart shows the same information as the bar graph above.



Follow the directions in the following table to switch between standard and velocity line graphs, and to switch between bar graphs and pie charts.

ACCESS: Admin/Any Security Analyst

Changing Graph Types

TO CHANGE...	YOU CAN...
a bar graph to a pie chart	click Switch to Pie . Note that pie charts cannot display multiple datasets; see Selecting Datasets on page 560.
a pie chart to a bar graph	click Switch to Bar .
a line graph from a standard graph to a velocity graph	click Velocity and select Velocity .
a line graph from a velocity graph to a standard graph	click Velocity and select Standard .

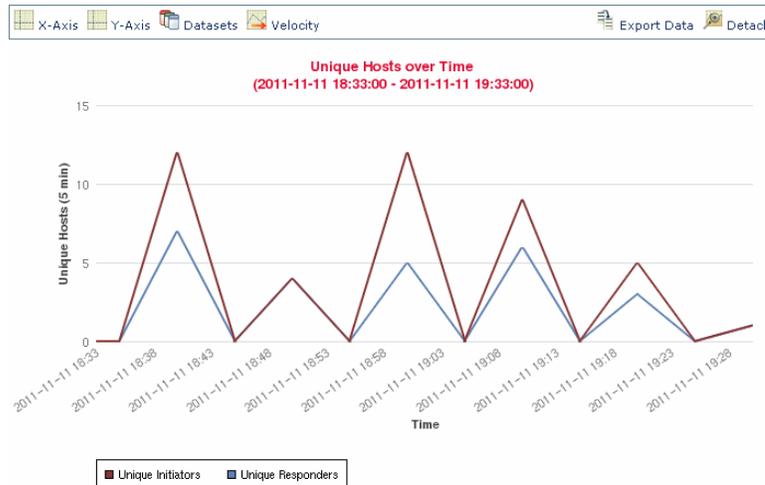
Selecting Datasets

LICENSE: Any

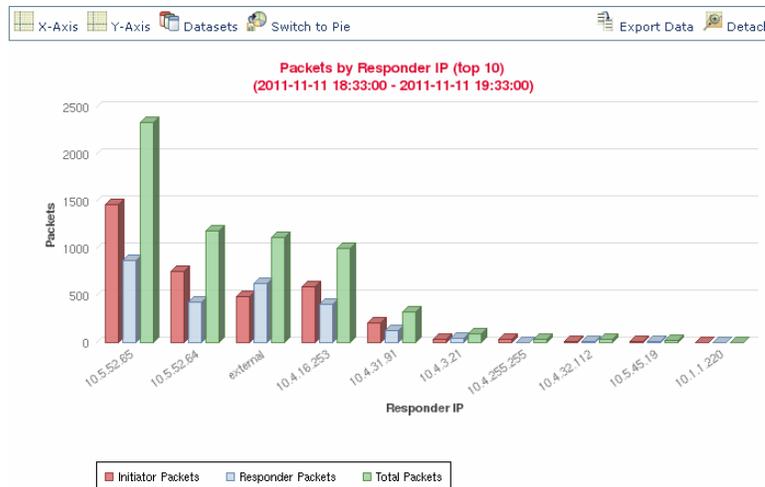
Both bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point. For example, you could display the total number of unique initiators and the total number of unique
Pie charts can only display one dataset.

On line graphs, multiple datasets appear as multiple lines, each with a different color. For example, the following graphic displays the total number of unique

initiators and the total number of unique responders detected on a monitored network over a one hour interval.



On bar graphs, multiple datasets appear as a set of colored bars for each x-axis data point. For example, the following bar graph displays the total packets transmitted on a monitored network, packets transmitted by initiators, and packets transmitted by responders.



You **cannot** display multiple datasets on a pie chart. If you switch to a pie chart from a bar graph that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When selecting which dataset to display, the Defense Center favors total statistics over initiator and responder statistics, and

favors initiator statistics over responder statistics. The [Dataset Options](#) table describes the datasets you can display on the x-axis of a connection graph.

Dataset Options

IF THE Y-AXIS DISPLAYS...	YOU CAN SELECT AS DATASETS...
Connections	<p>the default only, which is the number of connections detected on the monitored network (Connections)</p> <p>This is the only option for traffic profile graphs.</p>
KBytes	<p>combinations of:</p> <ul style="list-style-type: none"> • the total kilobytes transmitted on the monitored network (Total KBytes) • the number of kilobytes transmitted from host IP addresses on the monitored network (Initiator KBytes) • the number of kilobytes received by host IP addresses on the monitored network (Responder KBytes)
KBytes Per Second	<p>the default only, which is the total kilobytes per second transmitted on the monitored network (Total KBytes Per Second)</p>
Packets	<p>combinations of:</p> <ul style="list-style-type: none"> • the total packets transmitted on the monitored network (Total Packets) • the number of packets transmitted from host IP addresses on the monitored network (Initiator Packets) • the number of packets received by host IP addresses on the monitored network (Responder Packets)
Unique Hosts	<p>combinations of:</p> <ul style="list-style-type: none"> • the number of unique session initiators on the monitored network (Unique Initiators) • the number of unique session responders on the monitored network (Unique Responders)
Unique Application Protocols	<p>the default only, which is the number of unique application protocols on the monitored network (Unique Application Protocols)</p>
Unique Users	<p>the default only, which is the number of unique users logged into session initiators on the monitored network (Unique Initiator Users)</p>

To select the datasets displayed on a connection graph:

ACCESS: Admin/Any Security Analyst

- ▶ Click **Datasets** and select the datasets you want to graph.
The datasets you can select are described in the [Dataset Options](#) table.

Viewing Information About Aggregated Connection Data

LICENSE: Any

Connection graphs are based on aggregated data over five-minute intervals, also called *connection summaries*. You can get more information about the specific connection summaries used to construct a connection graph. For example, on a graph of connections over time, you may want to know exactly how many connections were detected over a specific interval.

To get detailed information on aggregated connection data:

ACCESS: Admin/Any Security Analyst

- ▶ Position your cursor over a point on a line graph a bar in a bar graph, or a wedge in a pie chart. A tooltip appears with detailed information about the data used to construct that portion of the graph.

Manipulating a Connection Graph on a Workflow Page

LICENSE: Any

When you open a connection data workflow, the data is initially constrained only by a time range. You can constrain connection graphs with additional criteria without advancing the workflow to the next page.

TIP! Constraining connection data in this manner changes the x-axis (also called the independent variable when viewing a pie chart) of the graph. To change the independent variable without constraining the connection data, use the **X-Axis** and **Y-Axis** menus. For more information, see [Selecting Data to Graph](#) on page 565.

To constrain connection data:

ACCESS: Admin/Any Security Analyst

1. Click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart.
2. Select a **View by...** option.

You can constrain connection data based on any of the criteria listed in the [X-Axis Functions table](#) on page 567.

For example, consider a graph of connections over time. If you constrain a point on the graph by port, a bar graph appears, showing the 10 most active ports based on the number of detected connection events, but constrained by the ten-minute time span that is centered on the point you clicked.

If you further constrain the graph by clicking on one of the bars and selecting **View by Initiator IP**, a new bar graph appears, constrained by not only the same ten-minute time span as before, but also by the port represented by the bar you clicked.

IMPORTANT! Unless you are working with a detached graph, constraining connection data in this manner changes the time range. For more information on detached graphs, see [Detaching Connection Graphs](#) on page 569.

Drilling Down Through Connection Data Graphs

LICENSE: Any

When you open a connection data workflow, the data is initially constrained only by a time range. You can constrain connection graphs while advancing the workflow to the next page.

To drill down in a connection data workflow:

ACCESS: Admin/Any Security Analyst

1. Click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart.
2. Select **Drill-down**.

You drill down to the next workflow page, constraining using the item you clicked:

- Clicking a point on a line graph constrains the time range on the next page to a 10-minute span, centered on the point you clicked.
- Clicking a bar on a bar graph or a wedge on a pie chart constrains the next page based on the criterion represented by the bar or wedge. For example, clicking on a bar that represents port use drills down to the next page in the workflow, which is constrained by the port represented by the bar you clicked.

Recentering and Zooming on Line Graphs

LICENSE: Any

You can recenter line graphs around any point in time. You can recenter using either the default time range, or you can choose a different time range.

IMPORTANT! Unless you are working with a detached graph, recentering changes the default time range. For more information on detached graphs, see [Detaching Connection Graphs](#) on page 569.

To recenter using the default time range:

ACCESS: Admin/Any Security Analyst

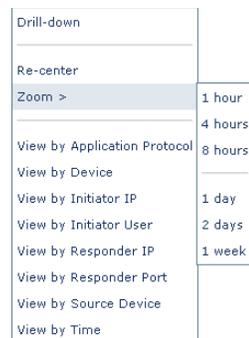
- ▶ Click the point on the line graph where you want to recenter the graph, and click **recenter**.

The graph is redrawn, centered on the point you clicked, with a time span that is the same length as your default time range.

To recenter using a different time range:

ACCESS: Admin/Any Security Analyst

1. Click the point where you want to recenter the graph and click **Zoom**.



2. Select the time span for the new graph, which can be as short as one hour or as long as one week.

The graph is redrawn, centered on the point you clicked, with the time span you selected.

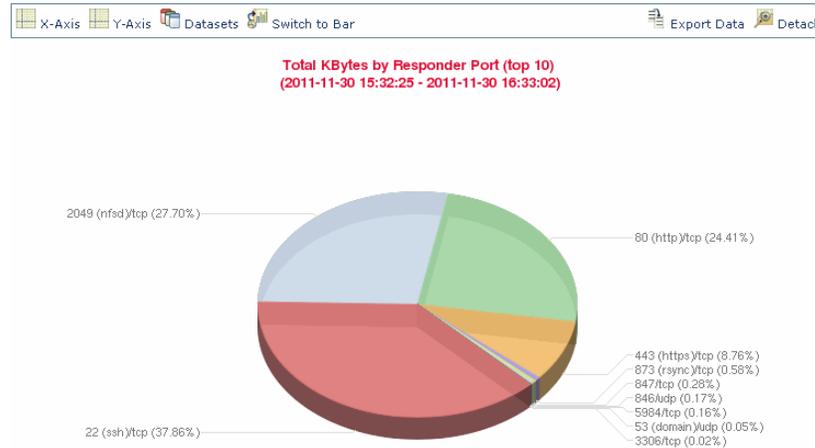
Selecting Data to Graph

LICENSE: Any

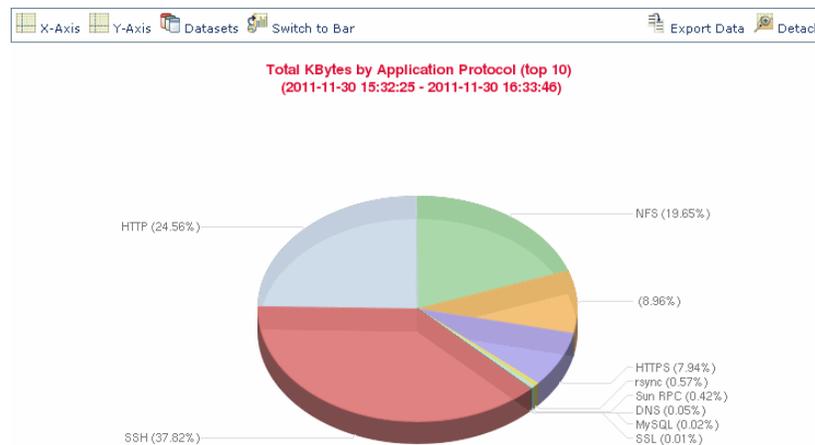
You can display different data on a connection graph by changing either the x-axis, the y-axis, or both.

Note that on a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable. For example, consider a

pie chart that graphs kilobytes per port. In this case, the x-axis is **Responder Port** and the y-axis is **KBytes**. This pie chart represents the total kilobytes of data transmitted over a monitored network during a certain interval. The wedges of the pie represent the percent of the data that was detected on each port.

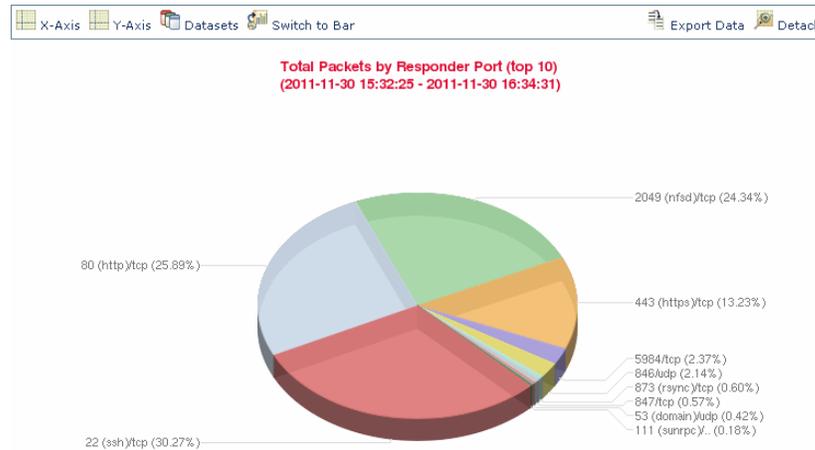


If you change the x-axis of the chart to **Application Protocol**, the pie chart still represents the total kilobytes of data transmitted, but the wedges of the pie represent the percentage of the data transmitted for each detected application protocol.



However, if you change the y-axis of the first pie chart to **Packets**, the pie chart represents the total number of packets transmitted over the monitored network

during a certain interval, and the wedges of the pie represent the percentage of the total number of packets that was detected on each port.



Follow the directions in the [X-Axis Functions](#) table to change the x-axis of a connection graph.

X-Axis Functions

TO GRAPH CONNECTION DATA...	YOU CAN...
by the 10 most active application protocols on the monitored network based on the number of detected connection events	click X-Axis and select Application Protocol .
by the 10 most active managed devices on the monitored network based on the number of detected connection events	click X-Axis and select Device .
by the 10 most active host IP addresses on the monitored network based on the number of connection events where that host IP address initiated the connection transaction	click X-Axis and select Initiator IP .
by the 10 most active users on the monitored network based on the number of connection events where the host where the user is logged in initiated the connection transaction	click X-Axis and select Initiator User .
by the 10 most active host IP addresses on the monitored network based on the number of connection events where that address was the responder in the connection transaction	click X-Axis and select Responder IP .

X-Axis Functions (Continued)

TO GRAPH CONNECTION DATA...	YOU CAN...
by the 10 most active ports on the monitored network based on the number of detected connection events where the host was the responder in the connection transaction	click X-Axis and select Responder Port .
by the 10 most active source devices, which include NetFlow-enabled devices that exported the connection data for the connections, plus a source device named FireSIGHT for all connections detected by Sourcefire managed devices	click X-Axis and select Source Device .
over time	click X-Axis and select Time .

Follow the directions in the [Y-Axis Functions](#) table to change the y-axis of a connection graph.

Y-Axis Functions

TO...	YOU CAN...
graph the number of connections on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select Connections .
graph the total kilobytes transmitted on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select KBytes .
graph the total kilobytes per second transmitted on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select KBytes Per Second .
graph the total number of packets transmitted on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select Packets .
graph the total number of unique hosts detected on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select Unique Hosts .

Y-Axis Functions (Continued)

To...	You CAN...
graph the total number of unique application protocols detected on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select Unique Application Protocols .
graph the total number of unique users detected on the monitored network by the criterion you chose for the x-axis	click Y-Axis and select Unique Users .

Detaching Connection Graphs

LICENSE: Any

If you want to perform further analysis on a connection graph, without affecting the default time range, you can detach the graph into a new browser window. You can perform all the same actions on detached connection graphs that you can on embedded connection graphs. You can also print a detached graph by clicking **Print**. Note that traffic profile graphs are, by default, detached graphs.

TIP! If you are viewing a detached graph, click **New Window** to create another copy of the detached graph in a new browser window. You can then perform different analyses on each of the detached graphs.

To detach a graph:

ACCESS: Admin/Any Security Analyst

- ▶ Click **Detach**.

Exporting Connection Data

LICENSE: Any

You can easily share connection data with others by exporting it as a CSV (comma-separated values) file.

TIP! You can also save a connection graph as an image by right-clicking on the graph and following your browser's prompts.

To export connection data:

ACCESS: Admin/Any Security Analyst

1. Click **Export Data**.

A pop-up window appears, displaying a table view of the data on your graph.

Time	Connections
2011-11-30 15:35:00	30346
2011-11-30 15:45:00	1180
2011-11-30 16:45:00	1755
2011-11-30 16:50:00	11671

2. Click **Download CSV File** and save the file.

Working with Connection Data Tables

LICENSE: Any

The Sourcefire 3D System's event viewer allows you to view connection data in a table, as well as manipulate the event view depending on the information relevant to your analysis. The page you see when you access connection data differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view.

The Sourcefire-provided *Connection Events* workflow provides a summary view of basic connection and detected application information, which you can then use to drill down to the table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

Using the event viewer, you can:

- search for, sort, and constrain events, as well as change the time range for displayed events
- specify the columns that appear (table view only)
- view the host profile associated with an IP address, or the user details and host history associated with a user identity
- view files (including malware files) and intrusions detected in connections
- view geolocation information associated with an IP address
- view the full text of a URL in a connection event
- view events using different workflow pages within the same workflow
- view events using a different workflow altogether
- drill down page-to-page within a workflow, constraining on specific values
- bookmark the current page and constraints so you can return to the same data (assuming the data still exists) at a later time

- create a report template using the current constraints
- delete events from the database
- use the IP address context menu to whitelist, blacklist, or obtain additional information about a host or IP address associated with a connection

Note that when you constrain connection events on a drill-down page, the packets and bytes from identical events are summed. However, if you are using a custom workflow and did not add a **Count** column to a drill-down page, the events are listed individually and packets and bytes are not summed.

The following sections contain information on viewing and analyzing connection event tables:

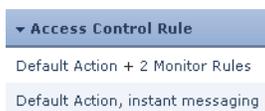
- [Understanding and Using Workflows](#) on page 1791 provides detailed instructions on using the event viewer.
- [Using Geolocation](#) on page 1816 provides information on how to view and interpret geolocation information associated with connection events.
- [Configuring Event View Settings](#) on page 2218 explains how to change the default workflow for viewing connection data.
- [Connection Data Fields](#) on page 543 and [Information Available in Connection Events](#) on page 550 provide details on the data in connection events.
- [Working with Connection Events Associated with Monitor Rules](#) on page 571 explains how to constrain connection events using Monitor rule criteria.
- [Viewing Files Detected in a Connection](#) on page 572 explains how to view the files, including malware files, detected or blocked in a connection.
- [Viewing Intrusion Events Associated with a Connection](#) on page 574 explains how to view the intrusion events associated with a connection.

Working with Connection Events Associated with Monitor Rules

LICENSE: Any

When you view logged connections using the event viewer, the Defense Center displays the access control rule or default action that handled each connection, as well as up to eight Monitor rules matched by each of those connections.

If a connection matched one Monitor rule, the Defense Center displays the name of the rule that handled the connection, followed by the Monitor rule name. If the connection matched more than one Monitor rule, the event viewer displays how many Monitor rules it matched, for example, **Default Action + 2 Monitor Rules**.



You can constrain connection event views using matched Monitor rules, using either of the following:

- the access control rule or default action that handled the connection
- any individual Monitor rule matched by a connection

To constrain connection events using Monitor rule matching:

ACCESS: Admin/Any Security Analyst

1. Select Analysis > Connections.

The first page of the default connection data workflow appears.

2. Display the workflow you want to use for your analysis. Make sure the drill-down page or table view you are using shows the Access Control Rule field.

3. How do you want to constrain the events?

- to constrain on the access control rule or default action that handled the connection, click the rule name or **Default Action**.
- to constrain on the only Monitor rule that matched a logged connection, click the Monitor rule name.
- to constrain on one of several Monitor rules that matched a logged connection, click an **N Monitor Rules** value. For example, click **2 Monitor Rules**.

The Monitor Rules pop-up window for that connection event appears, listing the first eight Monitor rules matched by the connection. Click the Monitor rule name you want to use to constrain connection events.



Your events are constrained. If you were using a drill-down page, the event view advances to the next page in the workflow.

Viewing Files Detected in a Connection

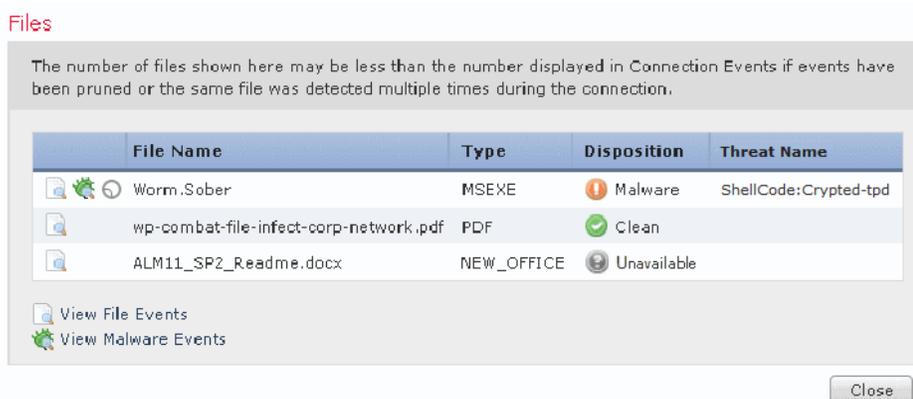
LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

If you associate a file policy with one or more access control rules, the system can detect files (including malware) in matching traffic. Using the event viewer, you can see the file events, if any, associated with the connections logged by those rules.

Instead of a list of files, the Defense Center displays the view files icon () in the **Files** column. The number on the icon indicates the number of files (including malware files) detected or blocked in that connection. Clicking on the icon does not drill down to the next workflow page or constrain connection events. Instead, it displays a pop-up window with a list of the files detected in the connection as well as their types, and if applicable, their malware dispositions.



In the pop-up window, you can click:

- a file's view icon () to view details in a table view of file events
- a malware file's view icon () to view details in a table view of malware events
- a file's trajectory icon () to track the file's transmission through your network
- **View File Events** or **View Malware Events** to view details on all of the connection's detected file or network-based malware events

TIP! To quickly view file or malware events associated with one or more connections, select the connections using the check boxes in the event viewer, then click either the **Files** or **Malware** tab above the column headings. You can view the connections used to transmit files in a similar way. For more information, see [Navigating Between Workflows](#) on page 1836.

When you view associated events, the Defense Center uses your default workflow for that event type. For more information on file and malware events, see [Working with File Events](#) on page 1222 and [Working with Malware Events](#) on page 1229. For more information on using the network file trajectory feature, see [Working with Network File Trajectory](#) on page 1241.

Note that not all file and malware events are associated with connections, as follows:

- Connections detected by pre-5.1.1 devices are not associated with files. Managed devices must be running at least Version 5.1.1 of the Sourcefire 3D System to perform file control and network-based malware detection.
- Endpoint-based malware events are not associated with connections. Those events are generated by FireAMP Connectors, instead of by the system inspecting network traffic.
- Many IMAP-capable email clients use a single IMAP session, which ends only when the user exits the application. Although long-running connections are logged by the system (see [Long-Running Connections](#) on page 542), files downloaded in the session are not associated with the connection until the session ends.

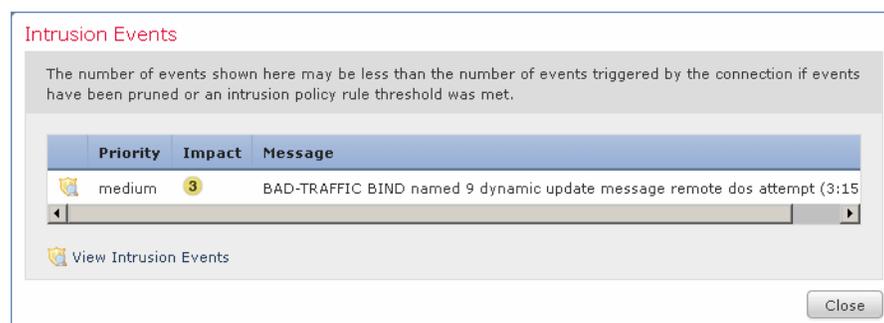
Note also that neither Series 2 devices nor the DC500 Defense Center support network-based advanced malware protection.

Viewing Intrusion Events Associated with a Connection

LICENSE: Protection

If you associate an intrusion policy with an access control rule or default action, the system can detect exploits in matching traffic. Using the event viewer, you can see the intrusion events, if any, associated with logged connections.

Instead of a list of events, the Defense Center displays the view intrusion events icon (🛡️) in the **Intrusion Events** column. Clicking on the icon does not drill down to the next workflow page or constrain connection events. Instead, it displays a pop-up window with a list of the intrusion events associated with the connection, as well as their priority and impact.



In the pop-up window, you can click a listed event's view icon (🔍) to view details in the packet view. You can also click **View Intrusion Events** to view details on all of the connection's associated intrusion events.

TIP! To quickly view intrusion events associated with one or more connections, select the connections using the check boxes in the event viewer, then click the **Intrusion** tab above the column headings. You can view the connections associated with intrusion events in a similar way. For more information, see [Navigating Between Workflows](#) on page 1836.

When you view associated events, the Defense Center uses your default intrusion events workflow. For more information on intrusion events, see [Working with Intrusion Events](#) on page 593.

Note that managed devices **must** be running at least Version 5.1.1 of the Sourcefire 3D System to associate intrusion events with connection events. Connections detected by pre-5.1.1 devices are not associated with intrusion events.

Searching for Connection Data

LICENSE: Any

Using the Defense Center's Search page, you can search for specific connection events or connection summaries, display the results in the event viewer, and save your search criteria to reuse later. Custom Analysis dashboard widgets, report templates, and custom user roles can also use saved searches.

The screenshot shows the 'Search Information' form. At the top, it says 'Note: If a search name is not specified, an automatically generated name will be used.' Below this, there are several fields and options:

- Table:** A dropdown menu set to 'Connection Events'.
- Name:** An empty text input field.
- Save As Private:** A checked checkbox.
- Constraint:** A section with several criteria:
 - First Packet:** > 2009-07-16 13:00:31, < today at 4:30pm
 - Last Packet:** > 2009-07-16 13:00:31, < today at 4:30pm
 - Active:** Allow, Deny, or IP
 - Response Bytes:** > 0
 - Traffic (Bytes):** > 50
 - Connections:** > 10

At the bottom of the form are two buttons: 'Search' and 'Save As New Search'.

Searches delivered with the system, labeled with (sourcefire) in the Saved Searches list, serve as examples.

Because connection graphs are based on connection summaries, the same criteria that constrain connection summaries also constrain connection graphs.

Fields marked with an asterisk (*) constrain connection graphs and connection summaries, as well as individual connection events.

If you search connection summaries using invalid search constraints and view your results using a connection summary page in a custom workflow, the invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough, as shown in the following graphic.



Also, keep in mind that your search results depend on the available data in the events you are searching. In other words, depending on the available data, your search constraints may not apply. See [Information Available in Connection Events](#) on page 550 for information on when data is available for each connection data field.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Special Search Syntax for Connection Data

To supplement the general search syntax listed above, the following table describes some special search syntax for connection data.

Connection Data Special Search Syntax

SEARCH CRITERION	SPECIAL SYNTAX
a Monitor rule matched by the connection	Use the Access Control Rule criterion to search for connections that matched individual Monitor rules. Because traffic matching a Monitor rule is always later handled by another rule or by the default action, you cannot search for a connection with an action of Monitor . Searching for the name of a Monitor rule returns all connections that matched that Monitor rule, regardless of the rule or default action that later handled the connection.
a criterion with a numerical value (Bytes, Packets, Connections)	You can precede the number with greater than (>), greater than or equal to (>=), less than (<), less than or equal to (<=), or equal to (=). TIP! To view meaningful results for searches using the Connections criterion, you must use a custom workflow that has a connection summary page.
Files or Intrusion Events associated with the connection	You cannot use the connection events Search page to search for file, malware, and intrusion events associated with a connection. For information on viewing these associated events, see Viewing Files Detected in a Connection on page 572 and Viewing Intrusion Events Associated with a Connection on page 574.
the Initiator User or URL for a connection	The system performs a partial match, that is, you can search for all or part of the field contents without using asterisks.
the total Traffic (in bytes) or transport Protocol used in the connection	These columns do not appear in table views. To determine if there is a protocol or traffic constraint on a connection table view, expand the search constraints. To search for a specific protocol, use the name or number protocol as listed in http://www.iana.org/assignments/protocol-numbers .
TCP Flags in a NetFlow connection	Type a list of comma-separated TCP flags to view all connections that have <i>at least</i> one of those flags (instead of all). You can also select the Only check box to search for connections that have any of the flags you specify as their only TCP flag.

To search for connection data:

ACCESS: Admin/Any Security Analyst

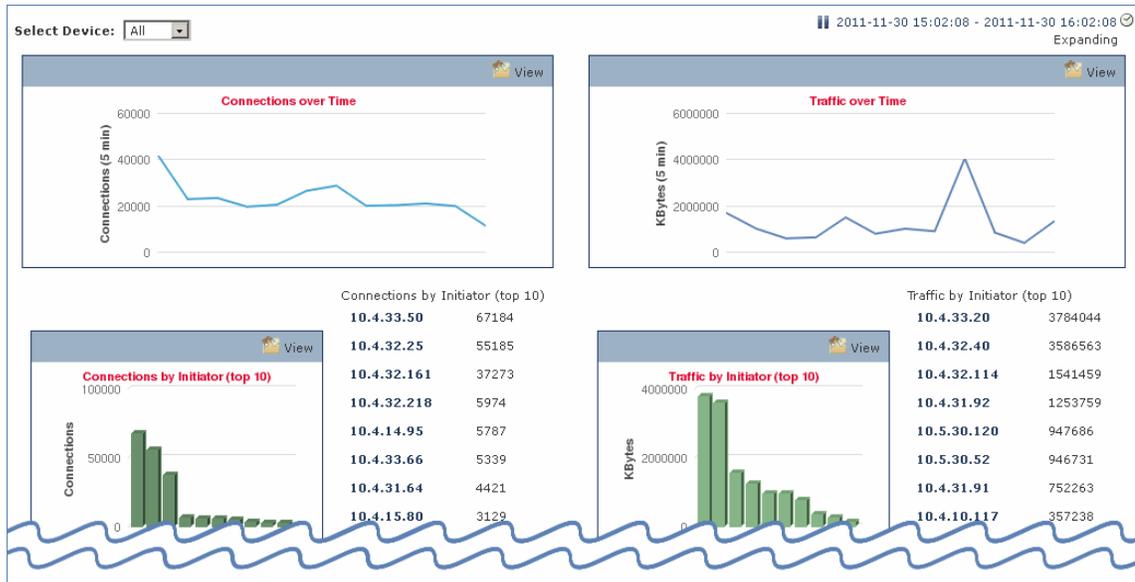
1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Connection Events**.
The page reloads with the appropriate constraints.
3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields.
See [Connection Data Fields](#) on page 543 for information on the fields in the connection events table.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.
6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in your default malware events workflow, constrained by the current time range.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**).

Viewing the Connection Summary Page

LICENSE: Any

The Connection Summary page provides graphs of the activity on your monitored network organized by different criteria. For example, the Connections over Time

graph displays the total number of connections on your monitored network over the interval that you select.



IMPORTANT! The Connection Summary page is visible only to users who have custom roles that are restricted by searches on connection events and who have been granted explicit access to the Connection Summary page. For more information, see [Understanding Restricted User Access Properties](#) on page 1912 and [Managing Custom User Roles](#) on page 1908.

The following table describes the different actions you can perform on the Connection Summary page.

Connection Summary Page Actions

To...	You CAN...
modify the time and date range for the Connection Summary page	find more information in Setting Event Time Constraints on page 1820.
manipulate connection graphs	find more information in Working with Connection Graphs on page 555.
detach a connection graph from the page	click View on the graph you want to detach. For more information on detached graphs, see Detaching Connection Graphs on page 569.

You can perform almost all the same actions on connection summary graphs that you can perform on connection graphs. However, because the graphs on the Connection Summary page are based on aggregated data, you cannot examine the individual connection events on which the graphs are based. In other words, you cannot drill down to a connection data table view from a connection summary graph.

To view the Connection Summary page:

ACCESS: Custom

1. Select **Overview > Summary > Connection Summary**.

The Connection Summary page appears for the current time range on your Defense Center.

2. From the **Select Device** list, select the device whose summary you want to view, or select **All** to view a summary of all devices.

CHAPTER 16

INTRODUCTION TO SOURCEFIRE INTRUSION PREVENTION

You can deploy your Sourcefire 3D System both to detect and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data. 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.

The Sourcefire 3D System can perform as both an intrusion detection system and an intrusion prevention system depending on:

- how you attach managed devices to your network: inline or out of band
- how you configure the devices' interface sets: passive, inline, switched, or routed
- the drop behavior of rules set to Drop and Generate Events: enabled or disabled

After you deploy your devices and configure them according to your needs, the Sourcefire 3D System uses several mechanisms to look for the broad range of exploits that attackers have developed. You can then use a broad range of tools to analyze and respond to the intrusion events.

Sensing Intrusions

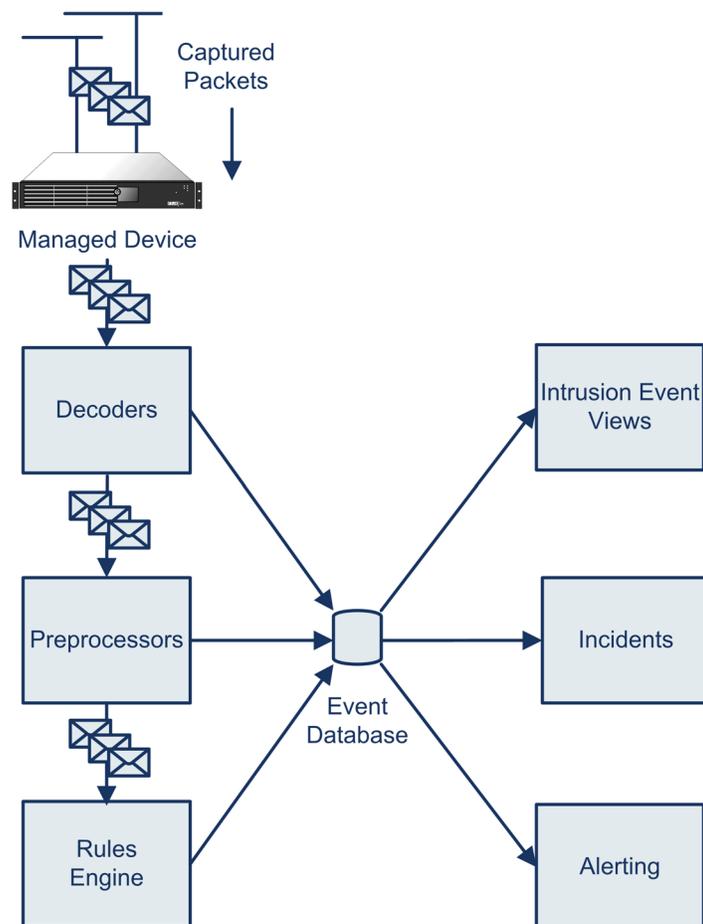
Packet decoders and preprocessors detect anomalous traffic that might signal an intrusion attempt and, when you have enabled accompanying decoder and preprocessor rules, report on detected anomalies. Next, intrusion rules examine the decoded packets for attacks based on patterns. Used together, intrusion rules

and preprocessors provide broader and deeper packet inspection than a signature-based system and help to identify intrusions more effectively.

The Sourcefire Vulnerability Research Team (VRT) regularly sends out updates, called Sourcefire rule updates, that may contain new intrusion rules, so you can be sure that you are detecting the most recently released attacks.

Responding to Intrusions

When a packet travels over a segment, the managed device captures and analyzes it using a series of decoders and preprocessors and then a rules engine. When the device identifies a possible intrusion, it generates an *intrusion event*, which is a record indicating the date, time, the type of exploit, and contextual information about the source of the attack and its target. Unless your device is deployed passively, the system can block possible intrusions or replace harmful content in a packet. For packet-based events, a copy of the packet or packets that triggered the event is also recorded.



To learn more about how a Sourcefire 3D System deployment can help protect your network, see the following sections:

- [Understanding How Traffic Is Analyzed](#) on page 583
- [Analyzing Intrusion Event Data](#) on page 587
- [Using Intrusion Event Responses](#) on page 588
- [Understanding Intrusion Prevention Deployments](#) on page 588
- [The Benefits of Custom Intrusion Policies](#) on page 591

Understanding How Traffic Is Analyzed

LICENSE: Protection

The system uses award-winning Snort® technology to analyze network traffic and generate intrusion events, which are records of the traffic that violates the intrusion policy applied to the device that is monitoring a specific network segment. Event analysts can review the events and determine whether they are important in the context of your network.

Intrusion events can be generated by:

- a link layer decoder, such as the Ethernet II decoder
- a network layer decoder, such as the IP decoder
- a transport layer decoder such, as the TCP decoder
- an application layer decoder or preprocessor, such as the HTTP Inspect preprocessor
- the rules engine

Events include such information as:

- the date and time the event was generated
- the event priority
- when you use network discovery, the impact flag associated with the event
- whether the packet that caused the event was dropped or would have been dropped in an inline, switched, or routed deployment
- the name of the device that generated the event
- the protocol of the packet that caused the event
- the source IP address and port for the event
- the destination IP address and port for the event
- the name of the user logged into the source host
- the ICMP type and code (for ICMP traffic)
- the Sourcefire 3D System component that generated the event (for example, the rule, decoder, or preprocessor)
- a brief description of the event

- the classification of the rule that generated the event
- the VLAN where the host is a member

For a complete list and descriptions of the information included in intrusion events, see [Understanding Intrusion Events](#) on page 604.

IMPORTANT! For events generated by shared object rules, the rule itself is not available.

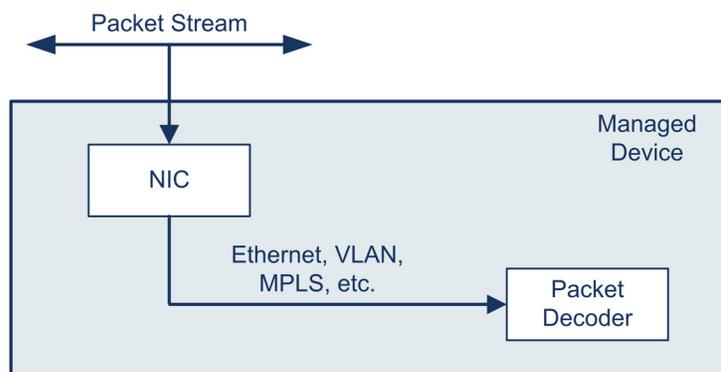
The following sections describe more about how the system acquires and processes information:

- [Capturing and Decoding Packets](#) on page 584
- [Processing Packets](#) on page 585
- [Generating Events](#) on page 586

Capturing and Decoding Packets

LICENSE: Protection

Before packets can be inspected, the packets must be captured from the network. The following illustration shows how the system sniffs packets, then decodes them before any further analysis.



As the system captures packets, it sends them to the packet decoder. The packet decoder converts the packet headers and payloads into a format that can be easily used by the preprocessors and the rules engine. 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, as described in the following table.

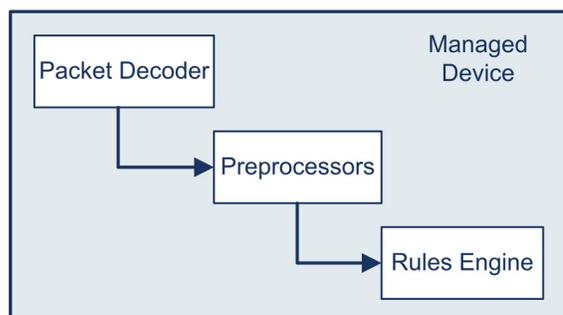
Decoded Packets

TCP/IP LAYER	DECODED PACKETS
Data Link	<ul style="list-style-type: none"> Ethernet Virtual local area network (VLAN) Multiprotocol Label Switching (MPLS)
Network	<ul style="list-style-type: none"> Encapsulated Remote Switched Port Analyzer (ERSPAN) Type II, Type III Internet Protocol version 4 (IPv4) Internet Protocol version 6 (IPv6) Internet Control Message Protocol version 4 (ICMPv4) Internet Control Message Protocol version 6 (ICMPv6) Point-to-Point Protocol (PPP) Point-to-Point Protocol over Ethernet (PPPoE) Generic Routing Encapsulation (GRE) Encapsulating Security Protocol (ESP) Teredo tunneling GPRS Tunneling Protocol (GTP)
Transport	<ul style="list-style-type: none"> Transmission Control Protocol (TCP) User Datagram Protocol (UDP)

Processing Packets

LICENSE: Protection

After the packets are decoded through the first three TCP/IP layers, they are sent to preprocessors, which normalize traffic at the application layer and detect protocol anomalies. After the packets have passed through the preprocessors, they are sent to the rules engine. The rules engine inspects the packet headers and payloads to determine whether they trigger any shared object rules or standard text rules.



You can enable and disable preprocessors and preprocessor options to suit your environment. For example, one of the preprocessors 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.

The rules engine takes three tracks as it inspects packets from the preprocessors:

- the rule optimizer
- the multi-rule search engine
- the event selector

For more information on preprocessors, see [Using Advanced Settings in an Intrusion Policy](#) on page 773.

The 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. As packets arrive at the rules engine, it selects the appropriate rule subsets to apply to each packet.

After the rule subsets are selected, the multi-rule search engine performs three different types of searches:

- 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.

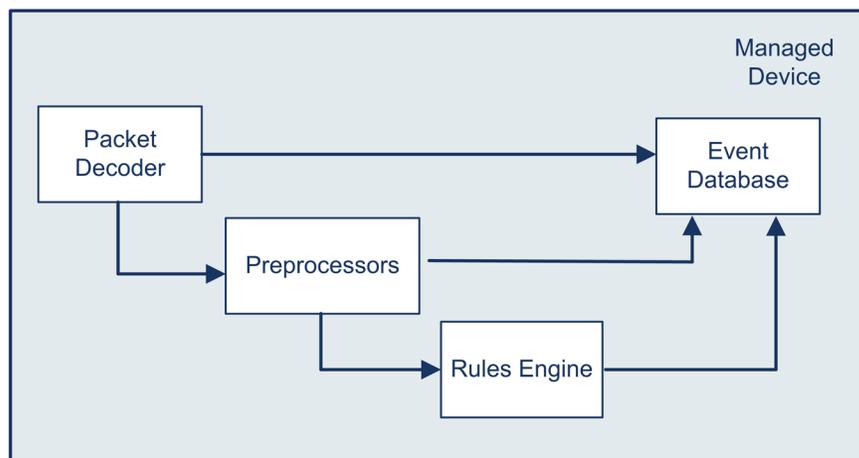
After the multi-rule search engine examines the packets, it generates an event for every rule triggered and adds it to an event queue. The event selector prioritizes the events in the queue and logs an event to the event database. These are the intrusion events that appear in the intrusion event statistics and intrusion event reports.

Generating Events

LICENSE: Protection

Packets are evaluated by the packet decoder, the preprocessors, and the rules engine. At each step of the process, a packet could cause the system to generate an event, which is an indication that the packet or its contents may be a risk to the

security of your network, or, in the case of an attack that originates from within your network, to the security of either your network or an external network.



For example, if the packet decoder 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 and, when the accompanying decoder rule is enabled, generates an event. Similarly, at the preprocessing step, 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, generates an event. The same kind of response occurs within the rules engine, with most rules written so that they also generate events when triggered by packets.

Each event in the database includes two sources of information about the potential attack. The first is called 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. The second is the packet log and includes a copy of the decoded packet header and packet payload.

Analyzing Intrusion Event Data

LICENSE: Protection

As the system accumulates intrusion events, you can begin your analysis of potential attacks. The Sourcefire 3D 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. These tools include:

- an Intrusion Event Statistics page that gives you an overview of the current activity on your managed device

For more information, see [Viewing Intrusion Event Statistics](#) on page 595.

- text-based and graphical reports that you can generate for any time period you choose; you can also design your own event reports and then configure them to run at scheduled intervals
For more information, see [Working with Reports](#) on page 1724.
- an incident-handling tool that you can use to gather event and packet data related to an attack; you can also add notes to help you track your investigation and response
For more information, see [Handling Incidents](#) on page 656.
- predefined and custom workflows that you can use to drill down through the intrusion events and to identify the events that you want to investigate further
For more information, see [Understanding and Using Workflows](#) on page 1791 and [Working with Intrusion Events](#) on page 593.

Using Intrusion Event Responses

LICENSE: Protection

In addition to generating intrusion events based on attacks, you can use an extensive list of alerting mechanisms to make sure that specific attacks are brought to your attention immediately. Conversely, you can suppress events that are less likely to affect critical systems or set a threshold that the number of events must reach before you are alerted.

For more information about automated alerting, see [Configuring External Responses to Intrusion Events](#) on page 1033.

You can use the following tools to set up automatic responses to intrusion events:

- automated alerting that you can configure for SNMP, email, and syslog
For more information, see [Configuring External Responses to Intrusion Events](#) on page 1033.
- on a Defense Center, automated correlation policies that you can use to respond to and remediate specific intrusion events
For more information, see [Configuring Remediations](#) on page 1606.

Understanding Intrusion Prevention Deployments

LICENSE: Protection

You can configure a passive deployment for a managed device so the device senses traffic out of band from the packet stream. Similarly, you can configure an inline, switched, or routed deployment where using an intrusion policy set to drop packets allows you to drop or replace packets that you know to be harmful.

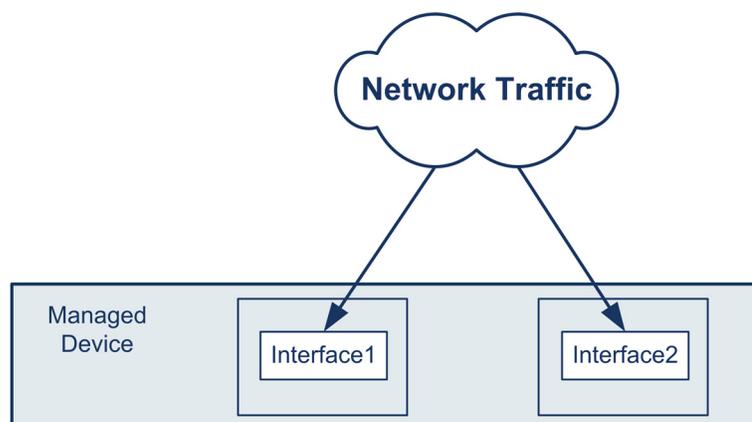
You can tailor intrusion policies for each managed device so they generate events only for the attacks that are likely to affect the security of the hosts on specific

portions of your network. You can specify which rules do not alert, which rules generate events and, except for a passive deployment, which rules generate events and also drop the malicious traffic.

For either type of system, you connect sensing interfaces to the appropriate segments on your network and add those interfaces to an interface set. These interfaces are configured in stealth mode so, to other devices on the network, the device itself does not appear to be connected to the network at all. Additionally, the interfaces are configured in promiscuous mode so that they detect all of the traffic on the network segment regardless of where the traffic is going. In this configuration, the device can see all of the traffic on the network segment, but is itself invisible.

The key deployment difference between an out-of-band deployment and an inline, switched, or routed deployment lies in the interface sets used by each system. An out-of-band deployment uses a passive interface set; an inline, switched, or routed deployment uses an inline set. The interfaces for a passive interface set passively analyze the traffic on the segments they monitor, while traffic flows between pairs of interfaces in an inline set.

The following illustration shows an example of a managed device deployed passively and with two passive interface sets. Each interface is monitoring a different network segment.

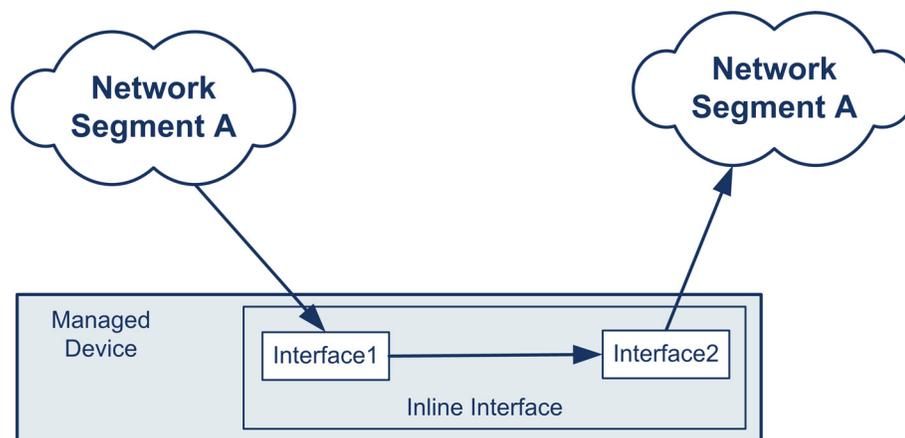


Sensing traffic out-of-band allows you to devote almost all of the device's sensing bandwidth and computational power to monitoring traffic, reconstructing datagrams and streams, normalizing packets, detecting anomalies, and alerting you to possible intrusions. Moreover, because the interface is deployed out of band and operates in stealth mode, attackers are unlikely to know of its existence, which renders it less of a target for attacks.

In an inline deployment, by comparison, you configure a managed device that uses an inline interface set. To do this, you connect the device to your network so that traffic flows between the device's network interfaces. When the interface set is configured as Inline, the interfaces are again configured in promiscuous mode so that they detect all of the traffic on the network segment regardless of where the traffic is going. The device can see all of the traffic on the network segment

but is itself invisible. However, when the interface set is deployed inline, you can configure rules to drop suspicious packets or, for custom standard text rules, to replace malicious portions of a packet payload with more benign content.

For example, the following illustration shows a device deployed inline. The device uses an interface set containing two of the network interfaces monitoring a single network segment.



Similar to a device using a passive interface set, the device using an inline interface set can see all of the traffic that passes through the interfaces in its interface set, regardless of the traffic's destination. However, because the traffic flows between the interfaces, you can modify or block suspicious packets. For example, if the device detects a packet whose payload contains a known exploit to which your network may be vulnerable, you can configure the system to drop the packet. In this case, the malicious packet never reaches its intended target.

In an inline deployment, you can also replace a portion of the payload with content of your own choosing. Consider a simple example where the device detects a packet that contains `bin/sh`, which often indicates a shellcode attack. You can write a custom intrusion rule that replaces all or part of this string with exactly the same number of characters. For example, replacing `bin/sh` with `foo/sh` and then passing the packet on to its destination causes the shellcode attack to fail without tipping off the attacker that the packet was altered.

Compare this result with the result when the same traffic is inspected passively. In that scenario, the same rule detects the exploit, but instead of having an option to drop the packet, you can only alert on its presence.

As you consider the benefits of deploying intrusion protection and prevention, you should weigh some of the trade-offs. First, you must choose a managed device model that matches or exceeds the traffic bandwidth of the network segment. Also, depending on the criticality of the hosts on the network segment, you should consider deploying the managed device with the optional *bypass* network card. The bypass card ensures that traffic continues to pass through the interfaces even if the appliance itself fails or loses power (although you may lose a few packets when you reboot the appliance). For more information on inline sets,

see [Configuring Inline Sets](#) on page 271. You can learn more about deployment options in your managed device's installation guide.

The Benefits of Custom Intrusion Policies

LICENSE: Protection

The system provides default intrusion policies suitable for both passive and inline deployments. However, you may find that the rules, preprocessor options, and other advanced settings configured in those policies do not address the security needs of your network. You can *tune* a policy by enabling, disabling, and setting specific configuration options for advanced settings and rules. Tuning advanced settings and rule sets allows you to configure, at a very granular level, how the system processes and inspects the traffic on your network.

For example, intrusion policies provide the following ways to tune preprocessors:

- Disable preprocessors that do not apply to the traffic on the subnet you are monitoring.
- Specify ports, where appropriate, to focus the activity of the preprocessor.
- Configure preprocessors to generate events when they encounter certain features in packets, for example, state problems or certain combinations of TCP flags.
- Configure adaptive profiles in combination with network discovery to use information about host operating systems from the network discovery map to switch to the most appropriate target-based profile for IP defragmentation and TCP stream preprocessing.

Note that the tuning options available vary by preprocessor or other advanced setting. For details on the available advanced settings, their options, and how to tune them, see [Using Advanced Settings in an Intrusion Policy](#) on page 773.

Additionally, within each intrusion policy, you can tune rules in the following ways:

- Improve performance by using fewer rules; disable rules that are not applicable to your environment.
- Verify that all rules applicable to your environment are enabled.
- For inline deployments, specify which rules should drop malicious packets from the packet stream.

TIP! You can use network discovery to identify the operating systems on your network. This allows you to more easily identify which rules are applicable to your environment.

- Modify the values of variables used in rules so the rules are more specific to your network.

Within the intrusion policy, you can also set suppression levels and thresholds to control how frequently you are notified of intrusion events. You can choose to

suppress event notifications and set thresholds for individual rules or entire intrusion policies. For more information, see [Setting Threshold Options within the Packet View](#) on page 631, [Setting Suppression Options within the Packet View](#) on page 633, and [Filtering Intrusion Event Notification Per Policy](#) on page 746.

Specifying the protocol analysis, data normalization, and traffic inspection performed by the system and saving this configuration as a whole allows you to control the kind of information the system provides you to best meet your enterprise security needs. It also provides a simple mechanism for changing as much or little of your policy as needed to continue to detect new attacks and exploits.

You can also tune rules in the following ways:

- Modify existing rules, if necessary, using the rule editor to correspond the rules to your network infrastructure.
- Write new standard text rules as needed using the Snort language and the rule editor to catch new exploits or to enforce your security policies.

For details on rule keywords, their arguments and syntax, and how to tune your rule set, see [Understanding and Writing Intrusion Rules](#) on page 1046.

CHAPTER 17

WORKING WITH INTRUSION EVENTS

The Sourcefire 3D System can help you monitor your network for traffic that could affect the availability, integrity, and confidentiality of a host and its data. By placing managed devices on key network segments, you can examine the packets that traverse your network for malicious activity. The system has several mechanisms it uses to look for the broad range of exploits that attackers have developed.

When the system identifies a possible intrusion, it generates an *intrusion event*, which is a record of the date, time, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded. Managed devices transmit their events to the Defense Center where you can view the aggregated data and gain a greater understanding of the attacks against your network assets.

You can also deploy a managed device as an inline, switched, or routed intrusion system, which allows you to configure the device to drop or replace packets that you know to be harmful.

The Sourcefire 3D System also 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. These tools include:

- an event summary page that gives you an overview of the current activity on your managed devices
- text-based and graphical reports that you can generate for any time period you choose; you can also design your own reports and configure them to run at scheduled intervals
- an incident-handling tool that you can use to gather event data related to an attack; you can also add notes to help you track your investigation and response

- automated alerting that you can configure for SNMP, email, and syslog
- automated correlation policies that you can use to respond to and remediate specific intrusion events
- predefined and custom workflows that you can use to drill down through the data to identify the events that you want to investigate further

See the following sections for more information:

- [Viewing Intrusion Event Statistics](#) on page 595 describes the Intrusion Event Statistics page, which provides you with an overview of the health of the appliance and a summary of the top threats to your network.
- [Viewing Intrusion Event Performance](#) on page 598 explains how to generate graphs of intrusion event performance statistics.
- [Viewing Intrusion Event Graphs](#) on page 601 explains how to generate charts that show event trends over time.
- [Viewing Intrusion Events](#) on page 602 describes how to use the web interface to view and investigate your intrusion events.
- [Understanding Workflow Pages for Intrusion Events](#) on page 612 describes the various pages that are available in intrusion event workflows and explains how you can use them to analyze your intrusion events.
- [Using Drill-Down and Table View Pages](#) on page 616 describes the features of two of the types of pages in an intrusion event workflow.
- [Using the Packet View](#) on page 622 explains how to use the packet view of intrusion events.
- [Using Impact Flags to Evaluate Events](#) on page 641 describes how you can use impact flags to evaluate intrusion events.
- [Searching for Intrusion Events](#) on page 644 explains how you can use the search feature to constrain a list of intrusion events to specific criteria.
- [Using the Clipboard](#) on page 652 describes how to add intrusion events to a holding area called the clipboard so that you can later add the events to incidents. This section also explains how to generate event reports based on the contents of the clipboard.

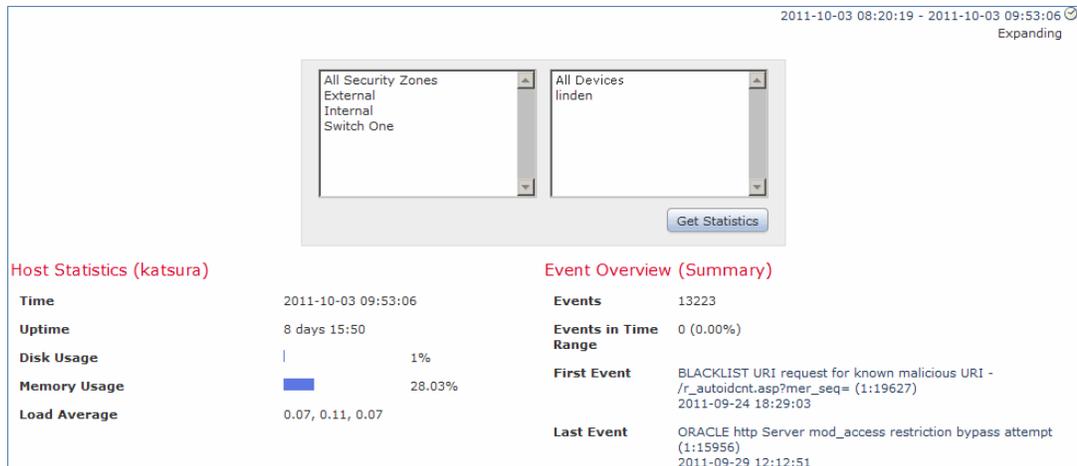
Also, see:

- [Handling Incidents](#) on page 656 for more information about incident handling and how you can use incidents to track the progress of an event analysis.
- [Configuring External Responses to Intrusion Events](#) on page 1033 for more information about automated alerting.
- [Working with Reports](#) on page 1724 for more information about intrusion event reports.
- [Using Geolocation](#) on page 1816 for more information about geolocation information in intrusion events.

Viewing Intrusion Event Statistics

LICENSE: Protection

The Intrusion Event Statistics page provides you with a quick summary of the current state of your appliance and any intrusion events generated for your network.



The Intrusion Event Statistics page has three main areas:

- [Host Statistics](#) on page 596 describes the Host Statistics section, which provides information about the appliance and, for Defense Centers, their managed devices.
- [Event Overview](#) on page 597 describes the Event Overview, which provides an overview of the information in the event database.
- [Event Statistics](#) on page 598 describes the Event Statistics, which provides more specific details about the information in the event database, such as the top 10 event types.

Each of the IP addresses, ports, protocols, event messages, and so on on the page is a link. Click any link to view the associated event information. For example, if one of the top 10 destination ports is **80 (http)/tcp**, clicking that link displays the first page in the default intrusion events workflow, and lists the events targeting that port. Note that only the events (and the managed devices that generate events) in the current time range appear. Also, intrusion events that you have marked reviewed continue to appear in the statistics. For example, if the current time range is the past hour but the first event was generated five hours ago, when you click the **First Event** link, the resulting event pages will not show the event until you change the time range.

To view intrusion event statistics:

ACCESS: Admin/Intrusion Admin

1. Select **Overview > Summary > Intrusion Event Statistics**.
The Intrusion Event Statistics page appears.
2. From the two selection boxes at the top of the page, select the zones and devices whose statistics you want to view, or select **All Security Zones** and **All Devices** to view statistics for all the devices that are collecting intrusion events.
3. Click **Get Statistics**.
The Intrusion Event Statistics page refreshes with data from the devices you selected.

TIP! To view data from a custom time range, click the link in the upper right page area and follow the directions in [Setting Event Time Constraints](#) on page 1820.

4. See the following sections for more information about the statistics that appear on the Intrusion Event Statistics page:
 - [Host Statistics](#) on page 596
 - [Event Overview](#) on page 597
 - [Event Statistics](#) on page 598

Host Statistics

LICENSE: Protection

The Host Statistics section of the Intrusion Event Statistics page provides information about the appliance itself. On the Defense Center, this section also provides information about any managed devices.

Host Statistics (katsura)	
Time	2011-09-27 10:26:50
Uptime	2 days 16:22
Disk Usage	 1%
Memory Usage	 18.71%
Load Average	0.32, 0.14, 0.07

This information includes the following:

- **Time** shows the current time on the appliance.
- **Uptime** shows the number of days, hours, and minutes since the appliance itself was restarted. On the Defense Center, the uptime also shows the last time each managed device was rebooted, the number of users logged in, and the load average.

- **Disk Usage** shows the percentage of the disk that is being used.
- **Memory Usage** shows the percentage of system memory that is being used.
- **Load Average** shows the average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.

Event Overview

LICENSE: Protection

The Event Overview section of the Intrusion Event Statistics page provides an overview of the information in the intrusion event database.

Event Overview (Summary)	
Events	6773
Events in Time Range	220 (3.25%)
First Event	BLACKLIST URI request for known malicious URI - /r_autoidcnt.asp?mer_seq= (1:19627) 2011-09-24 18:29:03
Last Event	BLACKLIST URI request for known malicious URI - vic.aspx?ver= (1:19623) 2011-09-27 10:25:37

These statistics include the following:

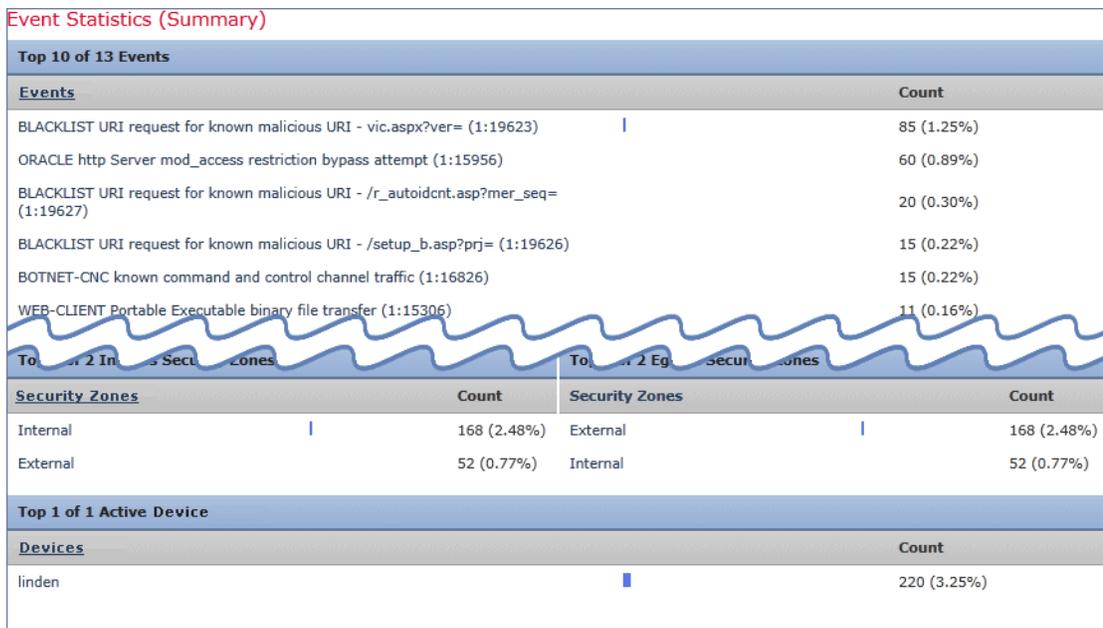
- **Events** shows the number of events in the intrusion event database.
- **Events in Time Range** shows the currently selected time range as well as the number and percentage of events from the database that fall within the time range.
- **First Event** shows the event message for the first event in the event database.
- **Last Event** shows the event message for the last event in the event database.

IMPORTANT! On the Defense Center, note that if you selected a managed device, the Event Overview section for that device appears instead.

Event Statistics

LICENSE: Protection

The Event Statistics section of the Intrusion Event Statistics page provides more specific information about the information in the intrusion event database.



This information includes details on:

- the top 10 event types
- the top 10 source IP addressees
- the top 10 destination IP addresses
- the top 10 destination ports
- the protocols, ingress and egress security zones, and devices with the greatest number of events

Viewing Intrusion Event Performance

LICENSE: Protection

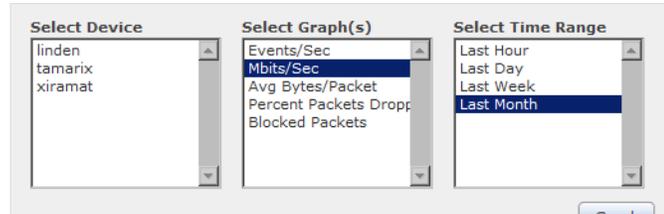
The intrusion event performance page allows you to generate graphs that depict performance statistics for intrusion events over a specific period of time. Graphs can be generated to reflect number of intrusion events per second, number of megabits per second, average number of bytes per packet, the percent of packets uninspected by Snort, and the number of packets blocked as the result of TCP normalization. These graphs can show statistics for the last hour, last day, last week, or last month of operation.

To view the intrusion event performance statistics:

ACCESS: Admin/Maint

- ▶ Select **Overview > Summary > Intrusion Event Performance**.

The Intrusion Event Performance page appears. The Defense Center version of the page is shown below.



See [Generating Intrusion Event Performance Statistics Graphs](#) on page 599 for more information.

Generating Intrusion Event Performance Statistics Graphs

LICENSE: Protection

You can generate graphs that depict performance statistics for a Defense Center or a managed device based on the number of events per second, megabits per second, average bytes per packet, percent of packets uninspected by Snort, and the number of packets blocked as the result of TCP normalization.

IMPORTANT! New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs.

The following table lists the available graph types.

Intrusion Event Performance Graph Types

GRAPH TYPE	OUTPUT
Events/Sec	Displays a graph that represents the number of events that are generated on the device per second
Mbits/Sec	Displays a graph that represents the number of megabits of traffic that pass through the device per second
Avg Bytes/ Packet	Displays a graph that represents the average number of bytes included in each packet

Intrusion Event Performance Graph Types (Continued)

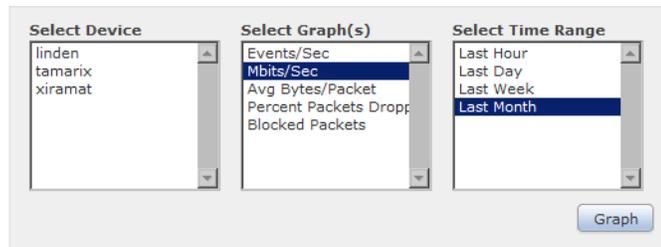
GRAPH TYPE	OUTPUT
Percent Packets Dropped	This graph depicts the average percentage of uninspected packets across all selected devices. For example, if you select two devices, then an average of 50% may indicate that one device has a 90% drop rate and the other has a 10% drop rate. It may also indicate that both devices have a drop rate of 50%. The graph only represents the total % drop when you select a single device.
Blocked Packets	Displays a graph that represents the number of packets blocked as the result of TCP normalization when the inline normalization Normalize TCP option is enabled. See TCP Normalization on page 921 for more information.

To generate intrusion event performance graphs:

ACCESS: Admin/Maint

1. Select **Overview > Summary > Intrusion Event Performance**.

The Intrusion Event Performance page appears. The Defense Center version of the page is shown below.

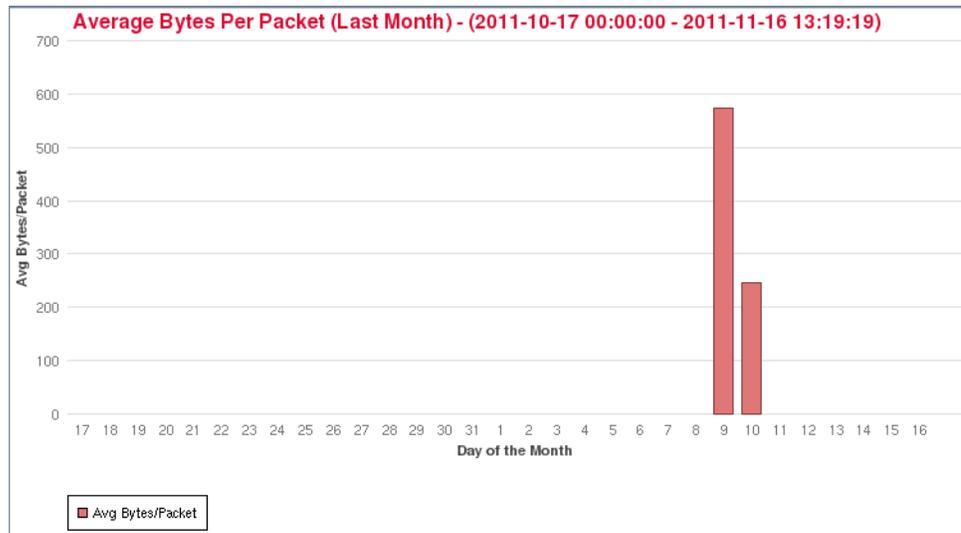


2. From the **Select Device** list, select the devices whose data you want to view.
3. From the **Select Graph(s)** list, select the type of graph you want to create.
4. From the **Select Time Range** list, select the time range you would like to use for the graph.

You can choose from last hour, last day, last week, or last month.

5. Click **Graph**.

The graph appears, displaying the information you specified.



6. To save the graph, right-click it and follow the instructions for your browser to save the image.

Viewing Intrusion Event Graphs

LICENSE: Protection

The Sourcefire 3D System provides graphs that show you intrusion event trends over time. You can generate intrusion event graphs over time ranging from the last hour to the last month, for the following:

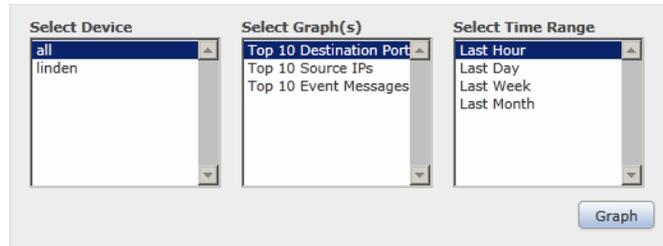
- one or all managed devices
- top 10 destination ports
- top 10 source IP addresses
- top 10 event messages

To generate an event graph:

ACCESS: Admin/Intrusion Admin

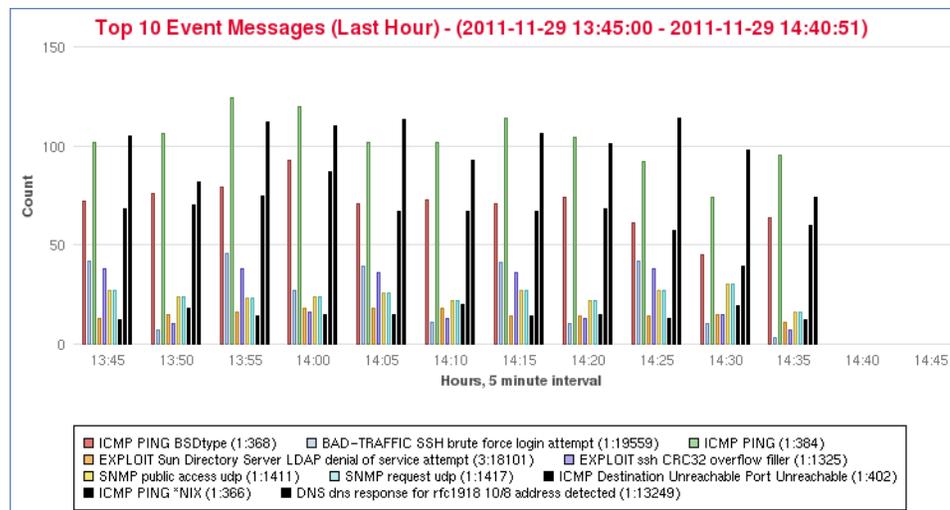
1. Select **Overview > Summary > Intrusion Event Graphs**.

The Intrusion Event Graphs page appears. Three selection boxes at the top of the page control which graph is generated.



2. Under **Select Device**, select **all** to include all devices, or select the specific device you want to include in the graph.
3. Under **Select Graph(s)**, select the type of graph you want to generate.
4. Under **Select Time Range**, select the time range for the graph.
5. Click **Graph**.

The graph is generated.



Viewing Intrusion Events

LICENSE: Protection

When the system recognizes a packet that is potentially malicious, it generates an intrusion event and adds the event to the database.

The initial intrusion events view differs depending on the workflow you use to access the page. You can use one of the predefined workflows, which includes one or more drill-down pages, a table view of intrusion events, and a terminating packet view, or you can create your own workflow. You can also view workflows based on custom tables, which may include intrusion events. Note that an event view may be slow to display if it contains a large number of IP addresses and you have enabled the **Resolve IP Addresses** event view setting. See [Configuring Event View Settings](#) on page 2218 for more information.

You view an intrusion event to determine whether there is a threat to your network security. If you are confident that an intrusion event is not malicious, you can mark the event reviewed. Your name appears as the reviewer, and the reviewed event is no longer listed in the default intrusion events view. You can return a reviewed event to the default intrusion events view by marking the event unreviewed.

You can view intrusion events that you have marked reviewed. Reviewed events are stored in the event database and are included in the event summary statistics, but no longer appear in the default event pages. See [Reviewing Intrusion Events](#) on page 611 for more information.

If you perform a backup and then delete reviewed intrusion events, restoring your backup restores the deleted intrusion events but does not restore their reviewed status. You view those restored intrusion events under Intrusion Events, not under Reviewed Events.

To quickly view connection events associated with one or more intrusion events, select the intrusion events using the check boxes in the event viewer, then click the **Connections** tab above the column headings. This is most useful when navigating between table views of events. You can also view the intrusions associated with particular connections in a similar way. Note that when a pre-Version 5.1.1 managed device generates the event, you cannot view associated events in this way on a Version 5.1.1 or greater Defense Center. For more information, see [Working With Connection Data](#) on page 539 and [Navigating Between Workflows](#) on page 1836.

For more information, see the following sections:

- [Understanding Intrusion Events](#) on page 604
- [Creating Custom Workflows](#) on page 1841
- [Using Drill-Down and Table View Pages](#) on page 616
- [Using the Packet View](#) on page 622
- [Viewing Connection Data Associated with Intrusion Events](#) on page 610
- [Reviewing Intrusion Events](#) on page 611
- [Viewing a Workflow Based on a Custom Table](#) on page 1787

To view intrusion events:

ACCESS: Admin/Intrusion Admin

► Select **Analysis > Intrusion > Events**.

The first page of the default intrusion events workflow appears. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

TIP! If you are using a custom workflow that does not include the table view of intrusion events, select any of the predefined workflows that ship with the appliance by clicking **(switch workflow)** next to the workflow title.

See [Understanding Intrusion Events](#) on page 604 to learn more about the events that appear in intrusion event views. See [Understanding Workflow Pages for Intrusion Events](#) on page 612 to learn more about how to narrow your view to the intrusion events that are important to your analysis.

Understanding Intrusion Events

LICENSE: Protection

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, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded.

The following list describes the information that an intrusion event contains. Note that some fields in the table view of intrusion events are disabled by default. To enable a field for the duration of your session, click the expand arrow (▶) to expand the search constraints, then click the column name under **Disabled Columns**.

Time

The date and time of the event.

Priority

The event priority as determined by the Sourcefire VRT.

Impact Flag

The impact flag in this field indicates the correlation between intrusion data, network discovery data, and vulnerability information. For more information, see [Using Impact Flags to Evaluate Events](#) on page 641.

Note that because there is no operating system information available for hosts added to the network map based on NetFlow data, the Defense Center cannot assign Vulnerable (impact level 1: red) impact flags for intrusion events involving those hosts, unless you use the host input feature to manually set the host operating system identity.

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 the system would have dropped the packet in an inline deployment if you enabled the **Drop when Inline** intrusion policy option
- 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. For more information, see [Setting Rule States](#) on page 743, [Setting Drop Behavior in an Inline Deployment](#) on page 689, [Configuring Passive Interfaces](#) on page 267, and [Tap Mode](#) on page 275.

Source IP

The IP address used by the sending host.

Source Country

The country of the sending host.

Destination IP

The IP address used by the receiving host.

Destination Country

The country of the receiving host.

Original Client IP

The original client IP address that was extracted from an X-Forwarded-For (XFF) or True-Client-IP HTTP header. To display a value for this field, you must enable the HTTP Inspect preprocessor **Extract Original Client IP Address** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

This field is disabled by default.

Source Port / ICMP Type

The port number on the sending host. For ICMP traffic, where there is no port number, the system displays the ICMP type.

Destination Port / ICMP Code

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, the system displays the ICMP code.

VLAN ID

The innermost VLAN ID associated with the packet that triggered the intrusion event.

MPLS Label

The Multiprotocol Label Switching label associated with the packet that triggered this intrusion event.

This field is disabled by default.

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.

Classification

The classification where the rule that generated the event belongs. See the [Rule Classifications table](#) on page 1058 for list of rule classification names and numbers.

Generator

The component that generated the event. See the [Generator IDs table](#) on page 785 for a list of intrusion event generator IDs.

Source User

The User ID for any known user logged in to the source host.

Destination User

The User ID for any known user logged in to the destination host.

Application Protocol

The application protocol, if available, which represents communications between hosts, detected in the traffic that triggered the intrusion event. For information on how the Sourcefire 3D System identifies detected application protocols in the Defense Center web interface, see the [Sourcefire 3D System Identification of Application Protocols table](#) on page 1267.

Client

The client application, if available, which represents software running on the monitored host detected in the traffic that triggered the intrusion event.

Web Application

The web application, which represents the content or requested URL for HTTP traffic detected in the traffic that triggered the intrusion event.

Note that if the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation here.

Category, Tag (Application Protocol, Client, Web Application)

Criteria that characterize an application to help you understand the application's function. For more information, see the [Application Characteristics table](#) on page 1264.

Application Risk

The risk associated with detected applications in the traffic that triggered the intrusion event. Each type of application detected in a connection has an associated risk; this field displays the highest risk of those. For more information, see the [Application Characteristics table](#) on page 1264.

Business Relevance

The business relevance associated with detected applications in the traffic that triggered the intrusion event. Each type of application detected in a connection has an associated business relevance; this field displays the lowest (least relevant) of those. For more information, see the [Application Characteristics table](#) on page 1264.

Ingress Security Zone

The ingress security zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment. See [Working with Security Zones](#) on page 187.

Egress Security Zone

For an inline deployment, the egress security zone of the packet that triggered the event. This security zone field is not populated in a passive deployment. See [Working with Security Zones](#) on page 187.

Device

The managed device where the access control policy was applied. See [Managing Devices](#) on page 190.

Ingress Interface

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface. See [Configuring Interfaces](#) on page 258.

Egress Interface

For an inline set, the egress interface of the packet that triggered the event. This interface column is not populated for a passive interface. See [Configuring Interfaces](#) on page 258.

Intrusion Policy

The intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event was enabled. You can select an intrusion policy as the default action for an access control policy, or you can associate an intrusion policy with an access control rule. See [Setting the Default Action](#) on page 428 and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

Access Control Policy

The access control policy that includes the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled. See [Using Access Control Policies](#) on page 424.

Access Control Rule

The access control rule associated with an intrusion rule that generated the event; see [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511. **Default Action** indicates that the intrusion policy where the rule is enabled is not associated with an access control rule but, instead, is configured as the default action of the access control policy; see [Setting the Default Action](#) on page 428.

HTTP Hostname

The host name, if present, that was extracted from the HTTP request Host header. Note that request packets do not always include the host name.

To display host names, you must enable the HTTP Inspect preprocessor **Log Hostname** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

This column displays the first fifty characters of the extracted host name. You can hover your pointer over the displayed portion of an abbreviated host name to display the complete name, up to 256 bytes. You can also display the complete host name, up to 256 bytes, in the packet view. See [Viewing Event Information](#) on page 625 for more information.

This field is disabled by default.

HTTP URI

The raw URI, if present, associated with the HTTP request packet that triggered the intrusion event. Note that request packets do not always include a URI.

To display the extracted URI, you must enable the HTTP Inspect preprocessor **Log URI** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly. See [Selecting Stream Reassembly Options](#) on page 950.

This column displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view. See [Viewing Event Information](#) on page 625 for more information.

This field is disabled by default.

Email Sender

The address of the email sender that was extracted from the SMTP MAIL FROM command. To display a value for this field, you must enable the SMTP preprocessor **Log From Address** option. Multiple sender addresses are supported. See [Understanding SMTP Decoding](#) on page 890 for more information.

This field is disabled by default.

Email Recipient

The address of the email recipient that was extracted from the SMTP RCPT TO command. To display a value for this field, you must enable the SMTP preprocessor **Log To Addresses** option. Multiple recipient addresses are supported. See [Understanding SMTP Decoding](#) on page 890 for more information.

This field is disabled by default.

Email Attachments

The MIME attachment file name that was extracted from the MIME Content-Disposition header. To display attachment file names, you must enable the SMTP preprocessor **Log MIME Attachment Names** option. Multiple attachment file names are supported. See [Understanding SMTP Decoding](#) on page 890 for more information.

This field is disabled by default.

Reviewed By

The name of the user who reviewed the event. See [Reviewing Intrusion Events](#) on page 611.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Viewing Connection Data Associated with Intrusion Events

LICENSE: Protection

When you associate an intrusion policy with an access control rule or the default action of an access control policy, the system can log the connections where intrusion events are detected, as follows:

- Although this logging is automatic for access control rules, you must manually enable connection logging to see associated connection data for the default action; see [Logging Connections for the Default Action](#) on page 431.
- Managed devices **must** be running at least Version 5.1.1 of the Sourcefire 3D System to associate intrusion events with connections. Connections detected by pre-5.1.1 devices are not associated with intrusion events.

To view connection data associated with one or more intrusion events:

ACCESS: Admin

1. Select **Analysis > Intrusion > Events**.

The first page of the default intrusion events workflow appears.

Viewing associated data is most useful when navigating between table views of events. See [Understanding Workflow Pages for Intrusion Events](#) on page 612 to learn more about how to narrow your view to the intrusion events that are important to your analysis.

2. Select the intrusion events using the check boxes in the event viewer, then click the **Connections** tab above the column headings.

You can view the intrusion events associated with particular connections in a similar way. For more information, see [Navigating Between Workflows](#) on page 1836.

When you view associated events, the Defense Center uses your default connection data workflow. For more information on connection data, see [Working With Connection Data](#) on page 539.

TIP! If you are using a custom workflow that does not include the table view of intrusion events, select any of the predefined workflows that ship with the appliance by clicking **(switch workflow)** next to the workflow title.

Reviewing Intrusion Events

LICENSE: Protection

If you have examined an intrusion event and are confident that the event does not represent a threat to your network security (perhaps because you know that none of the hosts on your network are vulnerable to the detected exploit), you can mark the event reviewed. Your name appears as the reviewer, and the reviewed event is no longer listed in the default intrusion events view. Events that you mark reviewed remain in the event database, but no longer appear in intrusion event views.

To mark an intrusion event reviewed:

ACCESS: Admin/Intrusion Admin

- ▶ On a page that displays intrusion events, you have two options:
 - To mark one or more intrusion events from the list of events, select the check boxes next to the events and click **Review**.
 - To mark all intrusion events from the list of events, click **Review All**.

A success message appears and the list of reviewed events is updated.

See [Understanding Intrusion Events](#) on page 604 to learn more about the events that appear in intrusion event views. See [Understanding Workflow Pages for Intrusion Events](#) on page 612 to learn more about how to narrow your view to the intrusion events that are important to your analysis.

IMPORTANT! Although they do not appear on intrusion event-related workflow pages, reviewed events are included in the event summary statistics.

To view events previously marked reviewed:

ACCESS: Admin/Intrusion Admin

► Select **Analysis > Intrusion > Reviewed Events**.

The first page of the default reviewed intrusion events workflow appears. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

TIP! If you are using a custom workflow that does not include the table view of intrusion events, select any of the predefined workflows that ship with the appliance by clicking (**switch workflow**) next to the workflow title.

See [Understanding Intrusion Events](#) on page 604 to learn more about the events that appear in reviewed intrusion event views. See [Understanding Workflow Pages for Intrusion Events](#) on page 612 to learn more about how to narrow your view to the intrusion events that are important to your analysis.

To mark reviewed events unreviewed:

ACCESS: Admin/Intrusion Admin

► On a page that displays reviewed events, you have two options:

- To remove individual intrusion events from the list of reviewed events, select the check boxes next to the events and click **Unreview**.
- To remove all intrusion events from the list of reviewed events, click **Unreview All**.

A success message appears and the list of reviewed events is updated.

Understanding Workflow Pages for Intrusion Events

LICENSE: Protection

The preprocessor, decoder, and intrusion rules that are enabled in the current intrusion policy generate intrusion events whenever the traffic that you monitor violates the policy.

The Sourcefire 3D System provides a set of predefined workflows, populated with event data, that you can use to view and analyze intrusion events. Each of these workflows steps you through a series of pages to help you pinpoint the intrusion events that you want to evaluate.

The predefined intrusion event workflows contain three different types of pages, or event views:

- one or more drill-down pages
- the table view of intrusion events
- a packet view

Drill-down pages generally include two or more columns in a table (and, for some drill-down views, more than one table) that allow you to view one specific type of information. For example, the following graphic shows a drill-down page with the number of events generated for each destination port.

<input type="checkbox"/>	Destination Port	Count
<input type="checkbox"/>	2049 (nfsd) / tcp	68
<input type="checkbox"/>	53 (domain) / udp	46
<input type="checkbox"/>	161 (snmp) / udp	39
<input type="checkbox"/>	31337 / udp	19
<input type="checkbox"/>	3389 / tcp	19
<input type="checkbox"/>	80 (http) / tcp	17
<input type="checkbox"/>	845 / tcp	3
<input type="checkbox"/>	58156 / tcp	3

Displaying rows 1–25 of 198 rows << Page 1 of 8 >>

View Copy Delete Review Download Packets
 View All Copy All Delete All Review All Download All Packets

When you “drill down” to find more information for one or more destination ports, you automatically select those events and the next page in the workflow appears. In this workflow, if you select the check box for the row with 80 (http)/ tcp and click **View**, the next page in the workflow appears (in this case, a page showing the number of events with each event message).

<input type="checkbox"/>	Message	Count
<input type="checkbox"/>	SERVER-IIS cmd.exe access (1:1002)	13
<input type="checkbox"/>	PUA-ADWARE Wajam Monitizer url outbound connection - post install (1:23246)	2
<input type="checkbox"/>	MALWARE-CNC TDS Sutra - request hi.cgi (1:21850)	2

Displaying rows 1–3 of 3 rows << Page 1 of 1 >>

View Copy Delete Review Download Packets
 View All Copy All Delete All Review All Download All Packets

In this way, drill-down tables help you reduce the number of events you are analyzing at one time. On the first page, you eliminated all but 17 events from your analysis. On the second page, if you select the check box for the first row of events (SERVER-ISS cmd.exe access (1:1002)) and click **View**, then the next page in the workflow appears and you have narrowed your list to 13 events.

The initial *table view* of intrusion events lists each intrusion event in its own row. The columns in the table list information such as the time, the source IP address and port, the destination IP address and port, the event priority, the event

message, and more. Note that several of the columns were removed from the table view to simplify the following graphic.

<input type="checkbox"/>	Time ×	Priority ×	Source IP ×	Destination IP ×	Source ×	Destination ×	Message ×
↓	2013-02-19 17:39:45	high	10.10.10.3	10.10.10.2	56820 / tcp	80 (http) / tcp	SERVER-IIS cmd.exe access (1:1002)
↓	2013-02-19 17:39:45	high	10.10.10.3	10.10.10.2	32913 / tcp	80 (http) / tcp	SERVER-IIS cmd.exe access (1:1002)
↓	2013-02-19 16:00:18	high	10.10.10.3	10.10.10.2	4258 / tcp	80 (http) / tcp	SERVER-IIS cmd.exe access (1:1002)
↓	2013-02-19 16:00:17	high	10.10.10.3	10.10.10.2	1079 / tcp	80 (http) / tcp	SERVER-IIS cmd.exe access (1:1002)

When you select events on a table view, instead of selecting events and displaying the next page in the workflow, you add to what are called *constraints*. Constraints are limits that you impose on the types of events that you want to analyze. In the previous two drill-down pages, the events were constrained in two ways: by destination port (80 (http)/tcp) and by event message (SERVER-IIS cmd.exe access (1:1002)). These constraints are carried forward to the table view. You can see the constraints by clicking the black arrow at the top of the table.

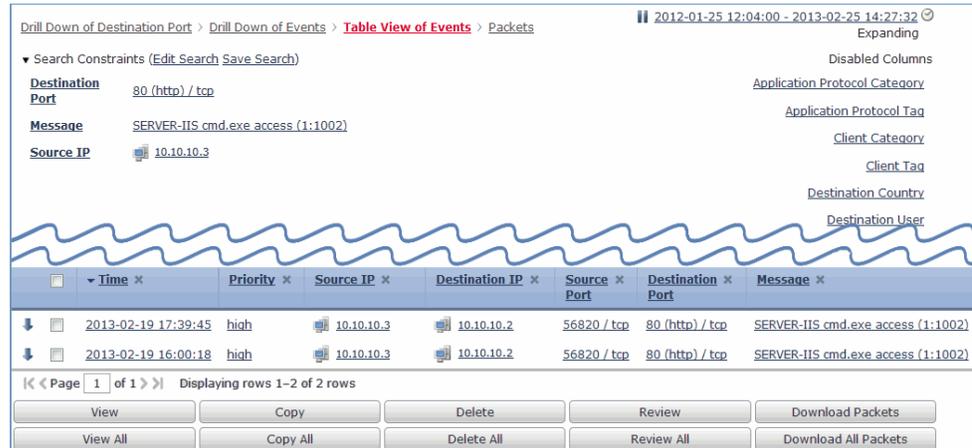
▼ Search Constraints (Edit Search Save Search)

Destination Port 80 (http) / tcp

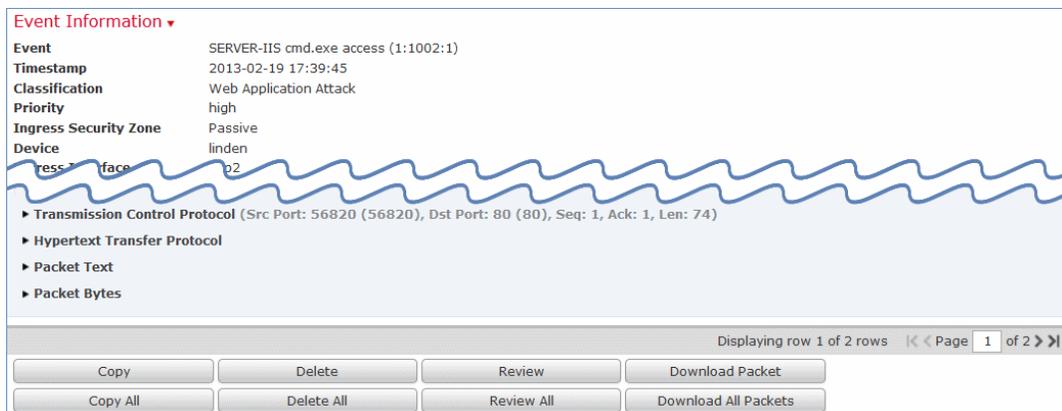
Message SERVER-IIS cmd.exe access (1:1002)

For example, if you click the close column icon (×) in any column and clear **Time** from the drop-down list, you can remove Time as one of the columns. To narrow the list of events in your analysis, you can click the link for a value in one of the rows in the table view. For example, to limit your analysis to the events generated from one of the source IP addresses (presumably, a potential attacker), click the IP address in the **Source IP Address** column. The following graphic shows the result, rows that provide a count of the number of times the attacker (10.10.10.3)

attempted to exploit the specific vulnerability (SERVER-ISS cmd.exe access (1:1002)) against, in this case, a single destination IP address (10.10.10.2).



If you select one or more rows in a table view and then click **View**, the packet view appears. A *packet view* provides information about the packet that triggered the rule or the preprocessor that generated the event. Each section of the packet view contains information about a specific layer in the packet. You can expand collapsed sections to see more information.



IMPORTANT! Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view. See [Detecting Portscans](#) on page 961 for more information.

If the predefined workflows do not meet your specific needs, you can create custom workflows that display only the information you are interested in. Custom intrusion event workflows can include drill-down pages, a table view of events, or both; the system automatically includes a packet view as the last page. You can

easily switch between the predefined workflows and your own custom workflows depending on how you want to investigate events.

TIP! [Understanding and Using Workflows](#) on page 1791 explains how to use workflows and the features common to all workflow pages. This chapter also explains how to create and use custom intrusion event workflows.

For more information, see:

- [Using Drill-Down and Table View Pages](#) on page 616, which explains how to use drill-down pages and the table view of events, which share many common features.
- [Using the Packet View](#) on page 622, which explains how to use the features in the packet view.
- [Searching for Intrusion Events](#) on page 644 explains how to search the event database for specific intrusion events.

Using Drill-Down and Table View Pages

LICENSE: Protection

The workflows that you can use to investigate intrusion events take advantage of three different types of pages:

- drill-down pages
- the table view of intrusion events
- the packet view

Each of these pages is described in [Understanding Workflow Pages for Intrusion Events](#) on page 612.

The drill-down views and table view of events share some common features that you can use to narrow a list of events and then concentrate your analysis on a

group of related events. The [Intrusion Event Common Features](#) table describes these features.

Intrusion Event Common Features

To...	YOU CAN...
learn more about the columns that appear	find more information in Understanding Intrusion Events on page 604.
view a host's profile	click the host profile icon () that appears next to the host IP address.
view geolocation details	click the flag icon that appears in the Source Country or Destination Country columns.
modify the time and date range for displayed events	find more information in Setting Event Time Constraints on page 1820. Note that events generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
sort and constrain events on the current workflow page	find more information in: <ul style="list-style-type: none"> • Sorting Drill-Down Workflow Pages on page 1835 • the Constraining Events on Drill-Down Pages table on page 620 • the Constraining Events on the Table View of Events table on page 621
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1836. TIP! To avoid displaying the same intrusion events on different workflow pages, the time range pauses when you click a link at the bottom of the page to display another page of events, and resumes when you click to take any other action on the subsequent page. For more information, see Setting Event Time Constraints on page 1820.
navigate between pages in the current workflow, keeping the current constraints	click the appropriate page link at the top left of the workflow page. For more information, see Using Workflow Pages on page 1812.

Intrusion Event Common Features (Continued)

To...	YOU CAN...
add events to the clipboard so you can transfer them to an incident at a later time	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To copy several intrusion events on a workflow page to the clipboard, select the check boxes next to events you want to copy, then click Copy. • To copy all the intrusion events in the current constrained view to the clipboard, click Copy All. <p>The clipboard stores up to 25,000 events per user. For more information, see Using the Clipboard on page 652.</p>
delete events from the event database	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To delete selected intrusion events, select the check boxes next to events you want to delete, then click Delete. • To delete all the intrusion events in the current constrained view, click Delete All, then confirm you want to delete all the events.
mark events reviewed to remove them from intrusion event pages, but not the event database	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To review selected intrusion events, select the check boxes next to events you want to review, then click Review. • To review all the intrusion events in the current constrained view, click Review All. <p>For more information, see Reviewing Intrusion Events on page 611.</p>
download a local copy of the packet (a packet capture file in libpcap format) that triggered each selected event	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To download the packets that triggered the selected intrusion events, select the check boxes next to events triggered by the packets you want to download, then click Download Packets. • To download all packets that triggered the intrusion events in the current constrained view, click Download All Packets. <p>Captured packets are saved in libpcap format. This format is used by several popular protocol analyzers.</p>
navigate to other event views to view associated events	find more information in Navigating Between Workflows on page 1836.
temporarily use a different workflow	click (switch workflow) . For more information, see Selecting Workflows on page 1809.
bookmark the current page so that you can quickly return to it	click Bookmark This Page . For more information, see Using Bookmarks on page 1838.

Intrusion Event Common Features (Continued)

To...	You CAN...
view the Intrusion Events section of the Summary Dashboard	click Dashboards . For more information, see Working with Dashboards on page 106.
navigate to the bookmark management page	click View Bookmarks . For more information, see Using Bookmarks on page 1838.
generate a report based on the data in the current view	click Report Designer . For more information, see Creating a Report Template from an Event View on page 1725.

The number of intrusion events that appear on the event views may be quite large, depending on:

- the time range you select
- the amount of traffic on your network
- the intrusion policy you apply

To make it easier to analyze intrusion events, you can constrain the event pages. The constraining processes are slightly different for drill-down views and the table view of intrusion events.

TIP! The time range pauses when you click one of the links at the bottom of the intrusion event workflow page to navigate to another page, and resumes when you click to take any other action on the subsequent page, including exiting the workflow; this reduces the likelihood of displaying the same events as you navigate to other pages in the workflow to see more events. For more information, see [Setting Event Time Constraints](#) on page 1820 and [Navigating to Other Pages in the Workflow](#) on page 1836.

The [Constraining Events on Drill-Down Pages](#) table describes how to use the drill-down pages.

Constraining Events on Drill-Down Pages

To...	YOU CAN...
drill down to the next workflow page constraining on a specific value	<p>click the value.</p> <p>For example, on the Destination Port workflow, to constrain the events to those with a destination of port 80, click 80/tcp in the DST Port/ICMP Code column. The next page of the workflow, Events, appears and contains only port 80/tcp events.</p>
drill down to the next workflow page constraining on selected events	<p>select the check boxes next to the events you want to view on the next workflow page, then click View.</p> <p>For example, on the Destination Port workflow, to constrain the events to those with destination ports 20/tcp and 21/tcp, select the check boxes next to the rows for those ports and click View. The next page of the workflow, Events, appears and contains only port 20/tcp and 21/tcp events.</p> <p>IMPORTANT! If you constrain on multiple rows and the table has more than one column (not including a Count column), you build what is called a compound constraint. Compound constraints ensure that you do not include more events in your constraint than you mean to. For example, if you use the Event and Destination workflow, each row that you select on the first drill-down page creates a compound constraint. If you pick event 1:100 with a destination IP address of 10.10.10.100 and you also pick event 1:200 with a destination IP address of 192.168.10.100, the compound constraint ensures that you do not also select events with 1:100 as the event type and 192.168.10.100 as the destination IP address or events with 1:200 as the event type and 10.10.10.100 as the destination IP address.</p>
drill down to the next workflow page keeping the current constraints	click View All .

The [Constraining Events on the Table View of Events](#) table describes how to use the table view.

Constraining Events on the Table View of Events

To...	YOU CAN...
constrain the view to events with a single attribute	click the attribute. For example, to constrain the view to events with a destination of port 80, click 80/tcp in the DST Port/ICMP Code column.
remove a column from the table	click the close icon (✕) in the column heading that you want to hide. In the pop-up window that appears, click Apply . TIP! To hide or show other columns, select or clear the appropriate check boxes before you click Apply . To add a disabled column back to the view, click the expand arrow (▶) to expand the search constraints, then click the column name under Disabled Columns .
view the packets associated with one or more events	either: <ul style="list-style-type: none">• click the down arrow icon (↓) next to the event whose packets you want to view.• select one or more events whose packets you want to view, and, at the bottom of the page, click View.• at the bottom of the page, click View All to view the packets for all events that match the current constraints.

TIP! At any point in the process, you can save the constraints as a set of search criteria. For example, if you find that over the course of a few days your network is being probed by an attacker from a single IP address, you can save your constraints during your investigation and then use them again later. You cannot, however, save compound constraints as a set of search criteria. For more information, see [Performing and Saving Searches](#) on page 1770.

TIP! If no intrusion events appear on the event views, adjusting the selected time range might return results. If you selected an older time range, events in that time range might have been deleted. Adjusting the rule thresholding configuration might generate events.

Using the Packet View

LICENSE: Protection

A packet view provides information about the packet that triggered the rule that generated an intrusion event.

TIP! The packet view on a Defense Center does not contain packet information when the **Transfer Packet** option is disabled for the device detecting the event.

Event Information ▾

Event SERVER-IIS cmd.exe access (1:1002:1)
Timestamp 2013-02-19 17:39:45
Classification Web Application Attack
Priority high
Ingress Security Zone Passive
Device linden

▸ Transmission Control Protocol (Src Port: 56820 (56820), Dst Port: 80 (80), Seq: 1, Ack: 1, Len: 74)
▸ Hypertext Transfer Protocol
▸ Packet Text
▸ Packet Bytes

Displaying row 1 of 2 rows << Page 1 of 2 >>

Copy Delete Review Download Packet
Copy All Delete All Review All Download All Packets

The packet view indicates why a specific packet was captured by providing information about the intrusion event that the packet triggered, including the event's time stamp, message, classification, priority, and, if the event was generated by a standard text rule, the rule that generated the event. The packet view also provides general information about the packet, such as its size.

In addition, the packet view has a section that describes each layer in the packet: data link, network, and transport, as well as a section that describes the bytes that comprise the packet. You can expand collapsed sections to display detailed information.

IMPORTANT! Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view. See [Detecting Portscans](#) on page 961 for more information.

The [Packet View Actions](#) table describes the actions you can take on the packet view.

Packet View Actions

To...	You CAN...
modify the date and time range in the packet views	find more information in Setting Event Time Constraints on page 1820.
learn more about the information displayed in the packet view	find more information in: <ul style="list-style-type: none">• Viewing Event Information on page 625• Viewing Frame Information on page 634• Viewing Data Link Layer Information on page 635• Viewing Network Layer Information on page 636• Viewing Transport Layer Information on page 638• Viewing Packet Byte Information on page 641
add an event to the clipboard so you can transfer it to the incidents at a later time	either: <ul style="list-style-type: none">• click Copy to copy the event whose packet you are viewing• click Copy All to copy all the events whose packets you previously selected <p>The clipboard stores up to 25,000 events per user. For more information on the clipboard, see Using the Clipboard on page 652.</p>
delete an event from the event database	either: <ul style="list-style-type: none">• click Delete to delete the event whose packet you are viewing• click Delete All to delete all the events whose packets you previously selected

Packet View Actions (Continued)

To...	You CAN...
mark an event reviewed to remove it from event views, but not the event database.	<p>either:</p> <ul style="list-style-type: none">• click Review to review the event whose packet you are viewing• click Review All to review all the events whose packets you previously selected <p>For more information, see Reviewing Intrusion Events on page 611. Note that reviewed events continue to be included in the event statistics on the Intrusion Event Statistics page.</p>
download a local copy of the packet (a packet capture file in libpcap format) that triggered the event	<p>either:</p> <ul style="list-style-type: none">• click Download Packet to save a copy of the captured packet for the event you are viewing• click Download All Packets to save copies of the captured packets for all the events whose packets you previously selected <p>The captured packet is saved in libpcap format. This format is used by several popular protocol analyzers.</p> <p>Note that you cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan view provides all usable packet information. See Understanding Portscan Events for more information.</p> <p>Note that you must have at least 15% available disk space in order to download.</p>
expand or collapse a page section	click the arrow next to the section.

To display the packet view:

ACCESS: Admin/Intrusion Admin

- ▶ On the table view of intrusion events, select packets to view. See the [Constraining Events on the Table View of Events table](#) on page 621 for more information.

The packet view appears. If you selected more than one event, you can page through the packets by using the page numbers at the bottom of the page.

Viewing Event Information

LICENSE: Protection

On the packet view, you can view information about the packet in the Event Information section.

Event Information ▾	
Event	INDICATOR-SHELLCODE x86 OS agnostic fnstenv geteip dword xor decoder (1:17322:1)
Timestamp	2013-03-01 11:05:25
Classification	Executable Code was Detected
Priority	high
Ingress Security Zone	External
Egress Security Zone	Internal
Device	linden
Ingress Interface	s1p2
Egress Interface	s1p1
Source IP	10.10.10.5
Source Port	55411 / udp
Destination IP	10.10.10.6
Destination Port	2049 (nfsd) / udp
Intrusion Policy	Balanced Security and Connectivity
Access Control Policy	IPS: balanced
Access Control Rule	Inspection
Rule	alert ip \$EXTERNAL_NET any -> \$HOME_NET any (msg:"INDICATOR-SHELLCODE x86 OS agnostic fnstenv geteip dword xor decoder"; content:"[D9 EE D9 74 24 F4]"; content:"[81]"; distance:1; content:"[13]"; distance:1; content:"[83]"; distance:1; content:"[FC E2 F4]"; distance:1; metadata:policy balanced-ips drop, policy connectivity-ips drop, policy security-ips drop; classtype:shellcode-detect; sid:17322; rev:2;)
Summary	This event is generated when shellcode is detected in network traffic.
Actions ▾	
Rule Actions	Edit View Documentation Rule Comment Set this rule to generate events in all locally created policies Disable this rule in all locally created policies Set this rule to drop the triggering packet and generate an event in all locally created inline intrusion policies
Set Thresholding Options	▶ in all locally created policies
Set Suppression Options	▶ in all locally created policies

Event

The event message. For rule-based events, this corresponds to the rule message. For other events, this is determined by the decoder or preprocessor.

The ID for the event is appended to the message in the format (*GID:SID:Rev*). *GID* is the generator ID of the rules engine, the decoder, or the preprocessor that generated the event. *SID* is the identifier for the rule, decoder message, or preprocessor message. *Rev* is the revision number of the rule. For more information, refer to [Reading Preprocessor Generator IDs](#) on page 784.

Timestamp

The time that the packet was captured.

Classification

The event classification. For rule-based events, this corresponds to the rule classification. For other events, this is determined by the decoder or preprocessor.

Priority

The event priority. For rule-based events, this corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other events, this is determined by the decoder or preprocessor.

Ingress Security Zone

The ingress security zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment. See [Working with Security Zones](#) on page 187.

Egress Security Zone

For an inline deployment, the egress security zone of the packet that triggered the event. See [Working with Security Zones](#) on page 187.

Device

The managed device where the access control policy was applied. See [Managing Devices](#) on page 190.

Ingress Interface

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface. See [Configuring Interfaces](#) on page 258.

Egress Interface

For an inline set, the egress interface of the packet that triggered the event. See [Configuring Interfaces](#) on page 258.

Source/Destination IP

The host IP address or domain name where the packet that triggered the event (source) originated, or the target (destination) host of the traffic that triggered the event.

Note that to display the domain name, you must enable IP address resolution; for more information, see [Configuring Event View Settings](#) on page 2218.

Click the address or domain name to view the context menu, then select **Whois** to do a whois search on the host, **View Host Profile** to view host information, or **Blacklist Now** or **Whitelist Now** to add the address to a global blacklist or whitelist. See [Using Host Profiles](#) on page 1334 and [Working with the Global Whitelist and Blacklist](#) on page 166.

Source Port/ICMP Type

Source port of the packet that triggered the event. For ICMP traffic, where there is no port number, the system displays the ICMP type.

Destination Port/ICMP Code

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, the system displays the ICMP code.

Email Headers

The data that was extracted from the email header. Note that email headers do not appear in the table view of intrusion events, but you can use email header data as a search criterion.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option. See [Understanding SMTP Decoding](#) on page 890 for more information. For rule-based events, this row appears when email data is extracted.

HTTP Hostname

The host name, if present, extracted from the HTTP request Host header. This row displays the complete host name, up to 256 bytes. Click the expand arrow (▶) to display the complete host name when longer than a single row.

To display host names, you must enable the HTTP Inspect preprocessor **Log Hostname** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

Note that HTTP request packets do not always include a host name. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

HTTP URI

The raw URI, if present, associated with the HTTP request packet that triggered the intrusion event. This row displays the complete URI, up to 2048 bytes. Click the expand arrow (▶) to display the complete URI when it is longer than a single row.

To display the URI, you must enable the HTTP Inspect preprocessor **Log URI** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

Note that HTTP request packets do not always include a URI. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly. See [Selecting Stream Reassembly Options](#) on page 950.

Intrusion Policy

The intrusion policy, if present, where the intrusion, preprocessor, or decoder rule that generated the intrusion event was enabled. You can select an intrusion policy as the default action for an access control policy or associate an intrusion policy with an access control rule. See [Setting the Default Action](#) on page 428 and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

Access Control Policy

The access control policy that includes the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled. See [Using Access Control Policies](#) on page 424.

Access Control Rule

The access control rule associated with an intrusion rule that generated the event; see [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511. **Default Action** indicates that the intrusion policy where the rule is enabled is not associated with an access control rule but, instead, is configured as the default action of the access control policy; see [Setting the Default Action](#) on page 428.

Rule

For standard text rule events, the rule that generated the event.

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Because rule data may contain sensitive information about your network, administrators may toggle users' ability to view rule information in the packet view with the View Local Rules permission in the user role editor. For more information, see [Modifying User Privileges and Options](#) on page 1912.

Actions

For standard text rule events, expand **Actions** to take any of the following actions on the rule that triggered the event:

- edit, the rule
- view documentation for the revision of the rule
- add a comment to the rule
- change the state of the rule
- set a threshold for the rule
- suppress the rule

See [Using Packet View Actions](#) on page 629, [Setting Threshold Options within the Packet View](#) on page 631, and [Setting Suppression Options within the Packet View](#) on page 633 for more information.

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Using Packet View Actions

LICENSE: Protection

On the packet view, you can take several actions in the Event Information section on the rule that triggered the event. Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available. You must expand **Actions** to display rule actions.



Edit

For standard text rule events, click **Edit** to modify the rule that generated the event.

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

IMPORTANT! If you edit a rule provided by Sourcefire (as opposed to a custom standard text rule), you actually create a new local rule. Make sure you set the local rule to generate events and also disable the original rule in the current intrusion policy. Note, however, that you **cannot** enable local rules in the default policies. For more information, see [Modifying Existing Rules](#) on page 1184.

View Documentation

For standard text rule events, click **View Documentation** to learn more about the rule revision that generated the event.

Rule Comment

For standard text rule events, click **Rule Comment** to add a text comment to the rule that generated the event.

This allows you to provide additional context and information about the rule and the exploit or policy violation it identifies. You can also add and view rule comments in the rule editor. For more information, see [Adding Comments to Rules](#) on page 1187.

Disable this rule

If this event is generated by a standard text rule, you can disable the rule, if necessary. You can set the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

For more information, see [Setting Rule States](#) on page 743.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Sourcefire.

IMPORTANT! You **cannot** disable shared object rules from the packet view, nor can you disable rules in the default policies.

Set this rule to generate events

If this event is generated by a standard text rule, you can set the rule to generate events in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

For more information, see [Setting Rule States](#) on page 743.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Sourcefire.

IMPORTANT! You **cannot** set shared object rules to generate events from e packet view, nor can you disable rules in the default policies.

Set this rule to drop

If your managed device is deployed inline on your network, you can set the rule that triggered the event to drop packets that trigger the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Sourcefire. Note also that this option appears only when **Drop when Inline** is enabled in the current policy. See [Setting Drop Behavior in an Inline Deployment](#) on page 689 for more information.

Set Thresholding Options

You can use this option to create a threshold for the rule that triggered this event in all policies that you can edit locally. Alternately, you create a threshold only for the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

The thresholding options are described in [Setting Threshold Options within the Packet View](#) on page 631.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default intrusion policy provided by Sourcefire.

Set Suppression Options

You can use this object to suppress the rule that triggered this event in all policies that you can edit locally. Alternately, you can suppress the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

The suppression options are described in [Setting Suppression Options within the Packet View](#) on page 633.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Sourcefire.

Setting Threshold Options within the Packet View

LICENSE: Protection

You can control the number of events that are generated per rule over time by setting the threshold options in the packet view of an intrusion event. You can set threshold options in all policies that you can edit locally or, when it can be edited locally, only in the in the current policy (that is, the policy that caused the event to be generated).

To set the threshold options within the packet view:

ACCESS: Admin/Intrusion Admin

1. Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section; expand **Set Thresholding Options** and select one of the two possible options:

- **in the current policy**
- **in all locally created policies**

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Sourcefire.

The thresholding options appear.

The screenshot shows a form for configuring thresholding options. It includes radio buttons for 'Type' (limit, threshold, both) and 'Track By' (Source, Destination). There are input fields for 'Count' and 'Seconds'. A checkbox labeled 'Override any existing settings for this rule' is present, along with a 'Save Thresholding' button.

2. 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.
3. Select the appropriate radio button to indicate whether you want the event instances tracked by **Source** or **Destination** IP address.
4. In the **Count** field, type the number of event instances you want to use as your threshold.
5. In the **Seconds** field, type a number between 1 and 86400 that specifies the time period for which event instances are tracked.
6. If you want to override any current thresholds for this rule in existing intrusion policies, select **Override any existing settings for this rule**.
7. Click **Save Thresholding**.

The system adds your threshold and displays a message indicating success. If you chose not to override existing settings, a message appears informing you of any conflicts.

Setting Suppression Options within the Packet View

LICENSE: Protection

You can use the suppression options to suppress intrusion events altogether, or based on the source or destination IP address. You can set suppression options in all policies that you can edit locally. Alternately, you can set suppression options only in the current policy (that is, the policy that generated the event) when the current policy can be edited locally.

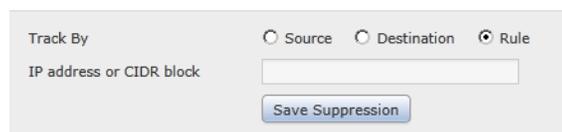
To suppress intrusion events within the packet view:

ACCESS: Admin/Intrusion Admin

1. Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section; expand **Set Suppression Options** and click one of the two possible options:
 - **in the current policy**
 - **in all locally created policies**

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Sourcefire.

The suppression options appear.



The screenshot shows a dialog box for setting suppression options. At the top, it says "Track By" followed by three radio buttons: "Source", "Destination", and "Rule". The "Rule" radio button is selected. Below the radio buttons is a text input field labeled "IP address or CIDR block". At the bottom of the dialog box is a "Save Suppression" button.

2. Select one of the following **Track By** options:
 - To completely suppress events for the rule that triggered this event, select **Rule**.
 - To suppress events generated by packets originating from a specified source IP address, select **Source**.
 - To suppress events generated by packets going to a specified destination IP address, select **Destination**.
3. In the **IP address or CIDR block** field, enter the IP address or CIDR block/prefix length you want to specify as the source or destination IP address.
For information on using CIDR notation and prefix lengths in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
4. Click **Save Suppression**.

The suppression options within your intrusion policies are modified according to your specifications. If you chose not to override existing settings, a message appears informing you of any conflicts.

Viewing Frame Information

LICENSE: Protection

On the packet view, click the arrow next to **Frame** to view information about the captured frame. The packet view may display a single frame or multiple frames. Each frame provides information about an individual network packet. You would see multiple frames, for example, in the case of tagged packets or packets in reassembled TCP streams. For information on tagged packets, see [Evaluating Post-Attack Traffic](#) on page 1166. For information on reassembled TCP streams, see [Reassembling TCP Streams](#) on page 949.

Frame n

The captured frame, where n is 1 for single-frame packets and the incremental frame number for multi-frame packets. The number of captured bytes in the frame is appended to the frame number.

Arrival Time

The date and time the frame was captured.

Time delta from previous captured frame

For multi-frame packets, the elapsed time since the previous frame was captured.

Time delta from previous displayed frame

For multi-frame packets, the elapsed time since the previous frame was displayed.

Time since reference or first frame

For multi-frame packets, the elapsed time since the first frame was captured.

Frame Number

The incremental frame number.

Frame Length

The length of the frame in bytes.

Capture Length

The length of the captured frame in bytes.

Frame is marked

Whether the frame is marked (true or false).

Protocols in frame

The protocols included in the frame.

Viewing Data Link Layer Information

LICENSE: Protection

On the packet view, click the arrow next to the data link layer protocol (for example, **Ethernet II**) to view the data link layer information about the packet, which contains the 48-bit media access control (MAC) addresses for the source and destination hosts. It may also display other information about the packet, depending on the hardware protocol.

IMPORTANT! Note that this example discusses Ethernet link layer information; other protocols may also appear.

The packet view reflects the protocol used at the data link layer. The following listing describes the information you might see for an Ethernet II or IEEE 802.3 Ethernet packet in the packet view.

Destination

The MAC address for the destination host.

IMPORTANT! Ethernet can also use multicast and broadcast addresses as the destination address.

Source

The MAC address for the source host.

Type

For Ethernet II packets, the type of packet that is encapsulated in the Ethernet frame; for example, IPv6 or ARP datagrams. Note that this item only appears for Ethernet II packets.

Length

For IEEE 802.3 Ethernet packets, the total length of the packet, in bytes, not including the checksum. Note that this item only appears for IEEE 802.3 Ethernet packets.

Viewing Network Layer Information

LICENSE: Protection

On the packet view, click the arrow next to the network layer protocol (for example, **Internet Protocol**) to view more detailed information about network layer information related to the packet.

IMPORTANT! Note that this example discusses IP packets; other protocols may also appear.

See the following sections for more information:

- [Viewing IPv4 Network Layer Information](#) on page 636
- [Viewing IPv6 Network Layer Information](#) on page 637

Viewing IPv4 Network Layer Information

LICENSE: Protection

The following listing describes protocol-specific information that might appear in an IPv4 packet.

Version

The Internet Protocol version number.

Header Length

The number of bytes in the header, including any IP options. An IP header with no options is 20 bytes long.

Differentiated Services Field

The values for differentiated services that indicate how the sending host supports Explicit Congestion Notification (ECN):

- 0x0 — does not support ECN-Capable Transport (ECT)
- 0x1 and 0x2 — supports ECT
- 0x3 — Congestion Experienced (CE)

Total Length

The length of the IP packet, in bytes, minus the IP header.

Identification

The value that uniquely identifies an IP datagram sent by the source host. This value is used to trace fragments of the same datagram.

Flags

The values that control IP fragmentation, where:

values for the Last Fragment flag indicate whether there are more fragments associated with the datagram:

- 0 — there are no more fragments associated with the datagram
- 1 — there are more fragments associated with the datagram

values for the Don't Fragment flag control whether the datagram can be fragmented:

- 0 — the datagram can be fragmented
- 1 — the datagram must **not** be fragmented

Fragment Offset

The value for the fragment offset from the beginning of the datagram.

Time to Live (ttl)

The remaining number of hops that the datagram can make between routers before the datagram expires.

Protocol

The transport protocol that is encapsulated in the IP datagram; for example, ICMP, IGMP, TCP, or UDP.

Header Checksum

The indicator for whether the IP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an intrusion evasion attempt.

Source/Destination

The IP address or domain name for the source (or destination) host.

Note that to display the domain name, you must enable IP address resolution; for more information, see [Configuring Event View Settings](#) on page 2218.

Click the address or domain name to view the context menu, then select **Whois** to do a whois search on the host, **View Host Profile** to view host information, or **Blacklist Now** or **Whitelist Now** to add the address to a global blacklist or whitelist. See [Using Host Profiles](#) on page 1334 and [Working with the Global Whitelist and Blacklist](#) on page 166.

Viewing IPv6 Network Layer Information

LICENSE: Protection

The following listing describes protocol-specific information that might appear in an IPv6 packet.

Traffic Class

An experimental 8-bit field in the IPv6 header for identifying IPv6 packet classes or priorities similar to the differentiated services functionality provided for IPv4. When unused, this field is set to zero.

Flow Label

A optional 20-bit IPv6 hexadecimal value 1 to FFFFF that identifies a special flow such as non-default quality of service or real-time service. When unused, this field is set to zero.

Payload Length

A 16-bit field identifying the number of octets in the IPv6 payload, which is comprised of all of the packet following the IPv6 header, including any extension headers.

Next Header

An 8-bit field identifying the type of header immediately following the IPv6 header, using the same values as the IPv4 Protocol field.

Hop Limit

An 8-bit decimal integer that each node that forwards the packet decrements by one. The packet is discarded if the decremented value reaches zero.

Source

The 128-bit IPv6 address for the source host.

Destination

The 128-bit IPv6 address for the destination host.

Viewing Transport Layer Information

LICENSE: Protection

On the packet view, click the arrow next to the transport layer protocol (for example, **TCP**, **UDP**, or **ICMP**) to view more information about the packet.

TIP! Click **Data** when present to view the first twenty-four bytes of the payload for the protocol immediately above it in the Packet Information section of the packet view.

The contents of the transport layer for each of the following protocols is described below:

- [TCP Packet View](#) on page 639
- [UDP Packet View](#) on page 640
- [ICMP Packet View](#) on page 640

IMPORTANT! Note that these examples discuss TCP, UDP, and ICMP packets; other protocols may also appear.

TCP Packet View

LICENSE: Protection

This section describes the protocol-specific information for a TCP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Sequence number

The value for the first byte in the current TCP segment, keyed to initial sequence number in the TCP stream.

Next sequence number

In a response packet, the sequence number of the next packet to send.

Acknowledgement number

The TCP acknowledgement, which is keyed to the sequence number of the previously accepted data.

Header Length

The number of bytes in the header.

Flags

The six bits that indicate the TCP segment's transmission state:

- **U** — the urgent pointer is valid
- **A** — the acknowledgement number is valid
- **P** — the receiver should push data
- **R** — reset the connection

- S — synchronize sequence numbers to start a new connection
- F — the sender has finished sending data

Window size

The amount of unacknowledged data, in bytes, that the receiving host will accept.

Checksum

The indicator for whether the TCP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an evasion attempt.

Urgent Pointer

The position, if present, in the TCP segment where the urgent data ends. Used in conjunction with the u flag.

Options

The values, if present, for TCP options.

UDP Packet View

LICENSE: Protection

This section describes the protocol-specific information for a UDP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Length

The combined length of the UDP header and data.

Checksum

The indicator for whether the UDP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

ICMP Packet View

LICENSE: Protection

This section describes the protocol-specific information for an ICMP packet.

Type

The type of ICMP message:

- 0 — echo reply
- 3 — destination unreachable
- 4 — source quench
- 5 — redirect
- 8 — echo request
- 9 — router advertisement
- 10 — router solicitation
- 11 — time exceeded
- 12 — parameter problem
- 13 — timestamp request
- 14 — timestamp reply
- 15 — information request (obsolete)
- 16 — information reply (obsolete)
- 17 — address mask request
- 18 — address mask reply

Code

The accompanying code for the ICMP message type. ICMP message types 3, 5, 11, and 12 have corresponding codes as described in RFC 792.

Checksum

The indicator for whether the ICMP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

Viewing Packet Byte Information

LICENSE: Protection

On the packet view, click the arrow next to **Packet Bytes** to view hexadecimal and ASCII versions of the bytes that comprise the packet.

Using Impact Flags to Evaluate Events

LICENSE: Protection

To help you evaluate the impact an event has on your network, the Defense Center displays an impact flag in the table view of intrusion events. For each event, the Defense Center adds an impact flag whose color indicates the

correlation between intrusion data, network discovery data, and vulnerability information.

IMPORTANT! Because there is no operating system information available for hosts added to the network map based on NetFlow data, the Defense Center cannot assign impact Vulnerable (impact level 1: red) impact flags for intrusion events involving those hosts, unless you use the host input feature to manually set the hosts' operating system identity.

The [Impact Flags](#) table describes the possible values for the impact flags.

Impact Flags

IMPACT LEVEL	VULNERABILITY	COLOR	DESCRIPTION
0	Unknown	gray	Neither the source nor the destination host is on a network that is monitored by network discovery.
1	Vulnerable	red	Either: <ul style="list-style-type: none">the source or the destination host is in the network map, and a vulnerability is mapped to the hostthe source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software; see Setting a Red (Level 1) Impact Flag on page 1100 for more information
2	Potentially Vulnerable	orange	Either the source or the destination host is in the network map and one of the following is true: <ul style="list-style-type: none">for port-oriented traffic, the port is running a server application protocolfor non-port-oriented traffic, the host uses the protocol

Impact Flags (Continued)

IMPACT LEVEL	VULNERABILITY	COLOR	DESCRIPTION
3	Currently Not Vulnerable	yellow	Either the source or the destination host is in the network map and one of the following is true: <ul style="list-style-type: none"> for port-oriented traffic (for example, TCP or UDP), the port is not open for non-port-oriented traffic (for example, ICMP), the host does not use the protocol
4	Unknown Target	blue	Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.

To use the impact flag on the table view to evaluate events:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Intrusion > Events**.

The first page of the default intrusion events workflow appears. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

2. Constrain the event view to view only those events that you want to evaluate.

For more information, see [Using Drill-Down and Table View Pages](#) on page 616.

3. At the top of the page, click **Table View of Events**.

The table view of events appears. **Impact Flag** can have any of the values described in the [Impact Flags table](#) on page 642. Several of the columns in the following graphic were removed from the table view to simplify the graphic.

	Time	Priority	Impact Flag	Inline Result	Source IP	Destination IP	Source Port / ICMP Type	Destination Port / I
↓	2011-09-29 12:12:51	high	0	↓	10.10.10.4	10.10.10.5	36157/tcp	80 (http)/tcp
↓	2011-09-29 12:12:51	high	1	↓	10.10.10.4	10.10.10.5	42398/tcp	80 (http)/tcp
↓	2011-09-29 12:10:50	high	2	↓	10.10.10.4	10.10.10.5	35666/tcp	80 (http)/tcp
↓	2011-09-29 12:10:50	high	3	↓	10.10.10.4	10.10.10.5	41907/tcp	80 (http)/tcp
↓	2011-09-29 12:09:21	high	4	↓	10.10.10.4	10.10.10.5	1035/tcp	80 (http)/tcp

4. To sort the table by impact level, click **Impact Flag**.
The events are sorted by impact level.

TIP! To reverse the sort order, click **Impact Flag** again.

Searching for Intrusion Events

LICENSE: Protection

You can search for specific intrusion events by using a predefined search delivered with the Sourcefire 3D System or by creating your own search criteria.

The predefined searches serve as examples and can provide quick access to important information about your network. You may want to modify specific fields within the default searches to customize them for your network environment, then save them to reuse later. The search criteria you can use are described in the following list.

TIP! For information about the syntax for specifying IP addresses and ports in an intrusion event search, see [Specifying IP Addresses in Searches](#) on page 1776 and [Specifying Ports in Searches](#) on page 1777.

For more information on searching, including how to load and delete saved searches, see the [Searching for Events table](#) on page 1770.

Priority

Specify the priority of the events you want to view. The priority corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other intrusion events, the priority is determined by the decoder or preprocessor. Valid values are **high**, **medium**, and **low**.

Impact

Specify the impact flag assigned to the intrusion event based on the correlation between intrusion data and network discovery data. Valid case-insensitive values are **Impact Level 0**, **Impact Level 1**, **Impact Level 2**, **Impact Level 3**, and **Impact Level 4**.

Do not use impact flag colors or partial strings (for example, do not use **blue**, **level 1**, or **0**).

For more information, see [Using Impact Flags to Evaluate Events](#) on page 641.

Inline Result

Type either:

- **dropped**, to specify whether the packet is dropped in an inline deployment
- **would have dropped**, to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline deployment

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. For more information, see [Setting Rule States](#) on page 743, [Setting Drop Behavior in an Inline Deployment](#) on page 689, [Configuring Passive Interfaces](#) on page 267, and [Tap Mode](#) on page 275.

Source IP

Specify the IP address used by the source host involved in the intrusion events.

Destination IP

Specify the IP address used by the destination host involved in the intrusion events.

Source/Destination IP

Specify the source or destination IP address used by the host whose intrusion events you want to view.

Source Country

Specify the country of the source host involved in the intrusion events.

Destination Country

Specify the country of the destination host involved in the intrusion events.

Source/Destination Country

Specify the country of the source or destination host involved in the intrusion events you want to view.

Source Continent

Specify the continent of the source host involved in the intrusion events.

Destination Continent

Specify the continent of the destination host involved in the intrusion events.

Source/Destination Continent

Specify the continent of the source or destination host involved in the intrusion events you want to view.

Original Client IP

Specify the original client IP address extracted from the X-Forwarded-For (XFF) or True-Client-IP HTTP headers. To extract a value for this field in an intrusion event, you must enable the HTTP Inspect preprocessor **Extract Original Client IP Address** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

Protocol

Type the name or number of the transport protocol used in the connection as listed in <http://www.iana.org/assignments/protocol-numbers>.

Note that there is no Protocol column in the intrusion event table view. This is the protocol associated with the source and destination port/ICMP column.

Source Port / ICMP Type

Specify the source port associated with the intrusion event.

TIP! For ICMP traffic, which does not target ports, you can use this field to search for events with specific ICMP types.

Destination Port / ICMP Code

Specify the destination port associated with the intrusion event.

TIP! For ICMP traffic, which does not target ports, you can use this field to search for events with specific ICMP codes.

VLAN ID

Specify the innermost VLAN ID associated with the packet that triggered the intrusion event.

MPLS Label

Specify the Multiprotocol Label Switching label of the packet associated with the packet that triggered the intrusion event.

Message

Specify all or part of the event message for the events you want to view.

Classification

Enter the classification number, or all or part of the classification name or description for the rule that generated the events you want to view. You can also enter a comma-separated list of numbers, names, or descriptions. Finally, if you add a custom classification, you can also search using all or part of its name or description. See the [Rule Classifications table](#) on page 1058 for a list of classification numbers, names, and descriptions.

Generator

Specify the component that generated the events you want to view, as listed in the [Generator IDs table](#) on page 785.

Snort ID

Specify the Snort ID (SID) of the rule that generated the event or, optionally, specify the combination generator ID (GID) and SID of the rule, where the GID and SID are separated with a colon (:) in the format GID:SID. You can specify any of the values in the [Snort ID Search Values](#) table:

Snort ID Search Values

VALUE	EXAMPLE
a single SID	10000
a SID range	10000-11000
greater than a SID	>10000
greater than or equal to a SID	>=10000
less than a SID	<10000
less than or equal to a SID	<=10000
a comma-separated list of SIDs	10000,11000,12000
a single GID:SID combination	1:10000
a comma-separated list of GID:SID combinations	1:10000,1:11000,1:12000
a comma-separated list of SIDs and GID:SID combinations	10000,1:11000,12000

For more information, see [Reading Preprocessor Generator IDs](#) on page 784.

Note that the Snort ID column does not appear in search results; the SID of the events you are viewing is listed in the Message column.

Source User

Specify the User ID for a user logged in to the source host.

Destination User

Specify the User ID for a user logged in to the destination host.

Source/Destination User

Specify the User ID for a user logged in to the source or destination host.

Application Protocol

Type the name of the application protocol, which represents communications between hosts, detected in the traffic that triggered the intrusion event.

Client

Type the name of the client application, which represents software running on the monitored host detected in the traffic that triggered the intrusion event.

Web Application

Type the name of the web application, which represents the content or requested URL for HTTP traffic detected in the traffic that triggered the intrusion event.

Category, Tag (Application Protocol, Client, Web Application)

Type a category or tag associated with the application detected in the session. Use a commas to separate multiple categories or tags. These fields are case-insensitive.

Application Risk

Type the highest risk associated with the application detected in the session. Valid criteria are: **Very High**, **High**, **Medium**, **Low**, and **Very Low**. These fields are case-insensitive.

Business Relevance

Type the lowest business relevance associated with an application detected in the session. Valid criteria are: **Very High**, **High**, **Medium**, **Low**, and **Very Low**. These fields are case-insensitive.

Security Zone (Ingress, Egress, Ingress/Egress)

Type the name of a security zone associated with the packet that triggered the event. These fields are case-insensitive. See [Working with Security Zones](#) on page 187.

Device

Specify the device where the access control policy was applied. You can specify a device name, device group, or IP address. See [Managing Devices](#) on page 190, [Editing Assigned Device Names](#) on page 244, and [Managing Device Groups](#) on page 217.

Interface (Ingress, Egress)

Type the name of an interface associated with the packet that triggered the event; see [Configuring Interfaces](#) on page 258.

Intrusion Policy

Type the name of the intrusion policy associated with the event; see [Configuring Intrusion Policies](#) on page 667.

Access Control Policy

Type the name of the access control policy associated with the event; see [Using Access Control Policies](#) on page 424.

Access Control Rule

Type the name of the access control policy associated with the event; see [Understanding and Writing Access Control Rules](#) on page 472.

HTTP Hostname

Specify a single host name that was extracted from the HTTP request Host header.

To associate host names with intrusion events for HTTP client traffic, you must enable the HTTP Inspect preprocessor **Log Hostname** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

HTTP URI

Specify a single URI associated with the HTTP request packet that triggered the intrusion event.

To associate URIs with intrusion events for HTTP traffic, you must enable the HTTP Inspect preprocessor **Log URI** option. See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for more information.

Email Sender

Specify the address of the email sender that was extracted from the SMTP MAIL FROM command. You can also enter a comma-separated list to search for events associated with all specified addresses. See [Understanding Intrusion Events](#) on page 604 for more information.

Email Recipient

Specify the address of the email recipient that was extracted from the SMTP RCPT TO command. You can also enter a comma-separated list to search for events associated with all specified addresses. See [Understanding Intrusion Events](#) on page 604 for more information.

Email Attachments

Specify the MIME attachment file name that was extracted from the MIME Content-Disposition header. Enter a comma-separated list to search for events associated with all attachment file names in the list. See [Understanding Intrusion Events](#) on page 604 for more information.

Email Headers

Specify data that was extracted from the email header. Note that email headers do not appear in the table view of intrusion events, but you can use email header data as a search criterion.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option. See [Understanding SMTP Decoding](#) on page 890 for more information.

Reviewed By

Specify the name of the user who reviewed the event. See [Reviewing Intrusion Events](#) on page 611.

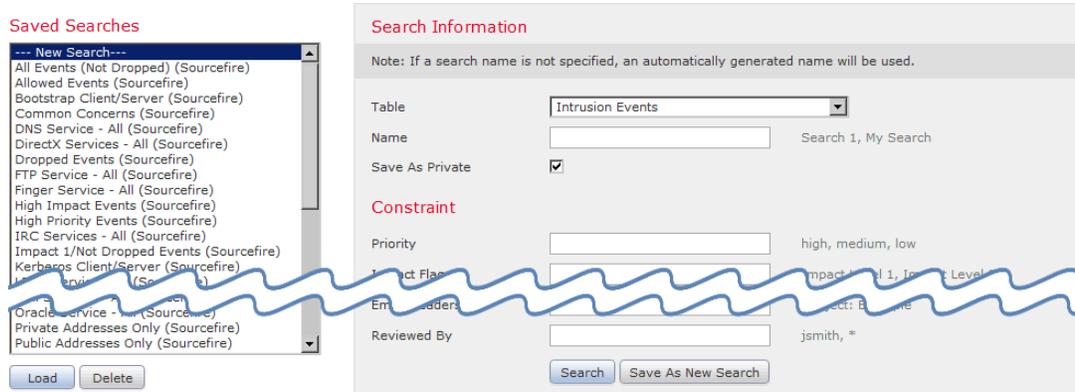
TIP! You can enter **unreviewed** to search for events that have not been reviewed.

To search for intrusion events:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Search**.

The Intrusion Events search page appears.



You can also click **Search** while viewing lists of intrusion events (**Analysis > Intrusion > Events**).

2. Optionally, if you want to save the search, enter a name for the search in the **Name** field.

If you do not enter a name, one is automatically created when you save the search.

3. Enter your search criteria in the appropriate fields, as described in the list above the procedure.

For more information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

4. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private. Note that users with the Administrator role can still view searches that you save as private.

If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

5. You have the following options:

- Click **Search** to start the search.

Your search results appear in the default intrusion events workflow, constrained by the current time range. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

- Click **Save** if you are modifying an existing search and want to save your changes.
- Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**) so that you can run it at a later time.

Using the Clipboard

LICENSE: Protection

The clipboard is a holding area where you can copy intrusion events from any of the intrusion event views. For information on how to add events to the clipboard, see [Using Drill-Down and Table View Pages](#) on page 616 and [Using the Packet View](#) on page 622.

The contents of the clipboard are sorted by the date and time that the events were generated. After you add intrusion events to the clipboard, you can delete them from the clipboard as well as generate reports on the contents of the clipboard.

You can also add intrusion events from the clipboard to incidents, which are compilations of events that you suspect are involved in a possible violation of your security policies. For more information about adding events from the clipboard to an incident, see [Creating an Incident](#) on page 661.

See the following sections for more information:

- [Generating Clipboard Reports](#) on page 652
- [Deleting Events from the Clipboard](#) on page 654

Generating Clipboard Reports

LICENSE: Protection

You can generate a report for the events on the clipboard just as you would from any of the event views.

To generate a report on intrusion events from the clipboard:

ACCESS: Admin/Intrusion Admin

1. Add one or more events to the clipboard:
 - For information on how to add events to the clipboard from a drill-down page or table view of events, see [Using Drill-Down and Table View Pages](#) on page 616.
 - For information on how to add events to the clipboard from the packet view, see [Using the Packet View](#) on page 622.

2. Select **Analysis > Intrusion > Clipboard**.

The clipboard appears.

<input type="checkbox"/>	Time	Priority	Device	Source IP	Destination IP	Source Port / ICMP Type	Destination Port / ICMP Code	Message
<input type="checkbox"/>	2011-09-29 12:12:51	high	linden	10.10.10.4	10.10.10.5	36157/tcp	80 (http/tcp)	ORACLE http Server mod_access restriction bypass attempt (1:15956)
<input type="checkbox"/>	2011-09-29 12:12:51	high	linden	10.10.10.4	10.10.10.5	42398/tcp	80 (http/tcp)	ORACLE http Server mod_access restriction bypass attempt (1:15956)
<input type="checkbox"/>	2011-09-29 11:54:12	high	linden	10.10.10.4	10.10.10.5	1061/tcp	88 (kerberos-sec)/tcp	NET-C known command and control channel traffic (1:16826)
<input type="checkbox"/>	2011-09-29 11:53:41	low	linden	10.10.10.4	10.10.10.5	81 (hosts2-ns)/tcp	1042/tcp	WEB-CLIENT Portable Executable binary file transfer (1:15306)

Displaying rows 1–25 of 1675 rows << Page 1 of 67 >>

3. You have the following options:

- To include specific events from a page on the clipboard, navigate to that page, select the check box next to the events, and click **Generate Report**.
- To include all the events from the clipboard, click **Generate Report All**.

In either case, the Report Templates page appears.

Save Generate Advanced

Report Title Report from Clipboard (1)

Report Sections

Clipboard

Table Clipboard

Preset None

Format

Search Clipboard Query

Fields Time, Priority, Device, Source IP, Destination

Section Description <<Time Window>><<Constraints>>

Time Window Inherit Time Window Last hour

Maximum Results Top 10000

Preview

4. Specify how you want your report to look, then click **Generate**.

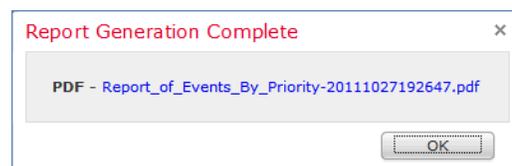
The Generate Report pop-up dialog appears.

5. Select one or more output formats (HTML, PDF, CSV) and, optionally, modify any of the other settings.

TIP! For more information about using the Report Designer, see [Working with Reports](#) on page 1724.

6. Click **Generate**, then click **Yes**.

The Report Generation Complete pop-up window appears with a link to view your report.



7. Click either:
 - a report link, which opens a new window to display the report you selected.
 - **OK** to return to the Report Templates page where you can modify your report design.

Deleting Events from the Clipboard

LICENSE: Protection

If you have intrusion events on the clipboard that you do not want to add to an incident, you can delete the events.

IMPORTANT! Deleting an event from the clipboard does **not** delete the event from the event database. However, deleting an event from the event database does delete the event from the clipboard.

To delete events from the clipboard:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Intrusion > Clipboard**.
The clipboard appears.

2. You have the following options:
 - To delete specific intrusion events from a page on the clipboard, navigate to the page, select the check box next to the events, and click **Delete**.
The events are deleted.
 - To delete all the intrusion events from the clipboard, click **Delete All**.
All the events are deleted from the clipboard. Note that if you select the **Confirm 'All' Actions** option in the Event Preferences, you are first prompted to confirm that you want to delete all the events.

CHAPTER 18

HANDLING INCIDENTS

Incident handling refers to the response an organization takes when a violation of its security policies is suspected. The Sourcefire 3D System includes features to support you as you collect and process information that is relevant to your investigation of an incident. You can use these features to gather intrusion events and packet data that may be related to the incident. You can also use the incident as a repository for notes about any activity that you take outside of the Sourcefire 3D System to mitigate the effects of the attack. For example, if your security policies require that you quarantine compromised hosts from your network, you can note that in the incident.

The Sourcefire 3D System also supports an incident life cycle, allowing you to change an incident's status as you progress through your response to an attack. When you close an incident, you can note any changes you have made to your security policies as a result of any lessons learned.

See the following sections for more information about handling incidents in the Sourcefire 3D System:

- [Incident Handling Basics](#) on page 657
- [Creating an Incident](#) on page 661
- [Editing an Incident](#) on page 663
- [Generating Incident Reports](#) on page 664
- [Creating Custom Incident Types](#) on page 665

Incident Handling Basics

LICENSE: Protection

Each organization is likely to have its own process for discovering, defining, and responding to violations of its security policies. The sections that follow describe some of the basics of incident handling and how you can incorporate the Sourcefire 3D System in your incident response plan:

- [Definition of an Incident](#) on page 657
- [Common Incident Handling Processes](#) on page 657
- [Incident Types in the Sourcefire 3D System](#) on page 660

Definition of an Incident

LICENSE: Protection

Generally, an *incident* is defined as one or more intrusion events that you suspect are involved in a possible violation of your security policies. Sourcefire also uses the term to describe the feature you use in the Sourcefire 3D System to track your response to an incident.

As explained in [Working with Intrusion Events](#) on page 593, some intrusion events are more important than others to the availability, confidentiality, and integrity of your network assets. For example, the port scan detection features provided by the Sourcefire 3D System can keep you informed of port scanning activity on your network. Your security policy, however, may not specifically prohibit port scanning or see it as a high priority threat, so rather than take any direct action, you may instead want to keep logs of any port scanning for later forensic study.

On the other hand, if the system generates events that indicate hosts within your network have been compromised and are participating in distributed denial-of-service (DDoS) attacks, then this activity is likely a clear violation of your security policy, and you should create an incident in the Sourcefire 3D System to help you track your investigation of these events.

Common Incident Handling Processes

LICENSE: Protection

Each organization is likely to define its own process for handling security incidents. Most methodologies include some or all of the following phases:

- [Preparation](#) on page 658
- [Detection and Notification](#) on page 658
- [Investigation and Qualification](#) on page 658
- [Communication](#) on page 659
- [Containment and Recovery](#) on page 660
- [Lessons Learned](#) on page 660

Each of these phases is described in the sections that follow. The descriptions also explain how the Sourcefire 3D System fits into each phase.

Preparation

You can prepare for incidents in two ways:

- by having clear and comprehensive security policies in place, as well as the hardware and software resources to enforce them
- by having a clearly defined plan to respond to incidents and a properly trained team that can implement the plan

A key part of incident handling is understanding which parts of your network are at the greatest risk. By deploying Sourcefire 3D System components on those network segments, you can increase your awareness of when and how incidents occur. Also, by taking the time to carefully tune the intrusion policy for each managed device, you can ensure that the events that are generated are of the highest quality.

Detection and Notification

You cannot respond to an incident unless you can detect it. Your incident handling process should note the kinds of security-related events that you can detect and the mechanisms, both software and hardware, that you use to detect them. You should also note where you can detect violations of your security policies. If your network includes segments that are not actively or passively monitored, then you need to note that as well.

The managed devices that you deploy on your network are responsible for analyzing the traffic on the segments where they are installed, for detecting intrusions, and for generating events that describe them. Keep in mind that the access control policy you apply to each of the managed devices governs what kinds of activity they detect and how it is prioritized. You can also set notification options for certain types of intrusion events so that the incident team does not need to sift through hundreds of events. You can specify that you are notified automatically when certain high priority, high severity events are detected.

Investigation and Qualification

Your incident handling process should specify how, after a security incident is detected, an investigation is conducted. In some organizations, junior members of the team triage all the incidents and handle the less severe or lower priority cases themselves. High severity and high priority incidents are handled by more senior members of the team. You should carefully outline the escalation process so that each team member understands the criteria for raising an incident's importance.

Part of the escalation process is tied to understanding how a detected event can affect the security of your network assets. For example, an attack against hosts running Microsoft SQL Server is not a high priority for organizations that use a different database server. Similarly, the attack is less important to you if you use

SQL Server on your network, but you are confident that all the servers are patched and are not vulnerable to the attack. However, if someone has recently installed a copy of the vulnerable version of the software (perhaps for testing purposes), you may have a greater problem than a cursory investigation would suggest.

The Sourcefire 3D System is particularly well suited to supporting the investigation and qualification process. You can create your own event classifications, and then apply them in a way that best describes the vulnerabilities on your network. When traffic on your network triggers an event, that event is automatically prioritized and qualified for you with special indicators showing which attacks are directed against hosts that are known to be vulnerable.

The incident tracking feature in the Sourcefire 3D System also includes a status indicator that you can change to show which incidents have been escalated.

Communication

All incident handling processes should specify how an incident is communicated between the incident handling team and both internal and external audiences. For example, you should consider what kinds of incidents require management intervention and at what level. Also, your process should outline how and when you communicate with outside organizations. Will some incidents require that you notify law enforcement agencies? If your hosts are participating in a distributed denial of service (DDoS) against a remote site, will you inform them? Do you want to share information with organizations such as the CERT Coordination Center (CERT/CC) or FIRST?

Sourcefire 3D System has features that you can use to gather intrusion data in standard formats such as HTML, PDF, and CSV (comma-separated values) so that you can easily share intrusion data with others.

For example, CERT/CC collects standard information about security incidents on its web site. CERT/CC looks for the kinds of information that you can easily extract from the Sourcefire 3D System, such as:

- information about the affected machines, including:
 - the host name and IP
 - the time zone
 - the purpose or function of the host
- information about the sources of the attack, including:
 - the host name and IP
 - the time zone

- whether you had any contact with an attacker
- the estimated cost of handling the incident
- a description of the incident, including:
 - dates
 - methods of intrusion
 - the intruder tools involved
 - the software versions and patch levels
 - any intruder tool output
 - the details of vulnerabilities exploited
 - the source of the attack
 - any other relevant information

You can also use the comment section of an incident to record when you communicate issues and with whom.

Containment and Recovery

Your incident handling process should clearly indicate what steps are taken when a host or other network component is compromised. The range of containment and recovery options stretches from applying patches to vulnerable hosts to shutting down the target and removing it from the network. You should also consider the importance, depending upon the nature and severity of the attack, of preserving evidence in case you pursue criminal charges.

You can use the incident feature of Sourcefire 3D System to maintain a record of the actions you take during the containment and recovery phase of the incident.

Lessons Learned

Each security incident, whether or not it is a successful attack, is an opportunity to review your security policies. Do you need to update your firewall rules? Do you need a more structured approach to patch management? Are unauthorized wireless access points a new security issue? Each lesson learned should feed back into your security policies and help you prepare better for the next incident.

Incident Types in the Sourcefire 3D System

LICENSE: Protection

You can assign an incident type to each incident you create. The following types are supported by default in the Sourcefire 3D System:

- Intrusion
- Denial of Service
- Unauthorized Admin Access
- Web Site Defacement

- Compromise of System Integrity
- Hoax
- Theft
- Damage
- Unknown

You can also create your own incident types, as explained in [Creating Custom Incident Types](#) on page 665.

Creating an Incident

LICENSE: Protection

This section explains how you create an incident.

To create an incident:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Intrusion > Incidents**.

The Incidents page appears.

2. Click **Create Incident**.

The Create Incident page appears.

The screenshot shows a web form titled "Create New Incident". It contains the following elements:

- Status:** A dropdown menu with "New" selected.
- Type:** A dropdown menu with "Intrusion" selected, followed by the word "Types".
- Time Spent:** A text input field containing "1m". To its right is the text "e.g. 1d 1h 1m 1s".
- Total Time Spent:** An empty text input field.
- Summary:** A large text area with a vertical scrollbar.
- Add Comment:** A larger text area with a vertical scrollbar.
- Buttons:** "Save" and "Reset" buttons at the bottom.

If you previously copied intrusion events to the clipboard, they are displayed at the bottom of the page, as in the following graphic. See [Using the Clipboard](#) on page 652 for information about using the clipboard.

Events in Your Clipboard

<input type="checkbox"/>	Time	Priority	Protocol	Source IP	Destination IP	Source Port / ICMP Type	Destination Port / ICMP Code	Message
<input type="checkbox"/>	2011-09-27 13:33:30	high	tcp	10.10.10.3	10.10.10.4	4675/tcp	445 (microsoft-ds)/tcp	SHELLCOD x86 OS agnostic xi dword decoder (1:17344)
<input type="checkbox"/>	2011-09-25 18:06:28	high	tcp	10.10.10.3	10.10.10.4	139 (netbios-ssn)/tcp	35361/tcp	SHELLCOD x86 OS agnostic alpha numeric upper case decoder (1:17340)

Displaying rows 1-2 of 2 rows << Page 1 of 1 >>

Add to Incident Add All to Incident

- From the **Type** drop-down menu, select the option that best describes the incident.
- In the **Time Spent** field, enter the amount of time you spent on the incident in the #d #h #m #s format, where # represents the number of days, hours, minutes, or seconds.
- In the **Summary** text box, type a short description (up to 255 alphanumeric characters spaces, and symbols) of the incident.
- In the **Add Comment** text box, type a more complete description (up to 8191 alphanumeric characters, spaces and symbols) for the incident.

7. Do you want to add events to the incident?
 - If **yes**, select the events on the clipboard and click **Add to Incident**.
You can also add all the events from the clipboard by clicking **Add All to Incident**.
 - If **no**, click **Save**.

In either case, the incident is saved with the information you entered.

IMPORTANT! If you want to add individual events from more than one page on the clipboard, you must add the events from one page, then add the events from the other pages separately.

Editing an Incident

LICENSE: Protection

You can update an incident as you collect more information. You can also add or delete events from the incident as your investigation progresses.

To edit an incident:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Intrusion > Incidents**.
The Incidents page appears.
2. Click the edit icon (✎) next to the incident you want to edit.
3. You can edit any of the following aspects of the incident:
 - change the status
 - change the type
 - add events from the clipboard
 - delete events
4. In the **Time Spent** field, enter the amount of additional time you spent on the incident.
5. In the **Add Comment** text box, indicate your changes to the incident (up to 8191 alphanumeric characters, spaces and symbols) for the incident.
6. Optionally, you can add or delete events from the incident:
 - To add events from the clipboard, select the events on the clipboard and click **Add to Incident**.
 - To add all the events from the clipboard, click **Add All to Incident**.
 - To delete specific events from the incident, select the events and click **Delete**.

- To delete all events from the incident, click **Delete All**.
 - To update the incident without adding or deleting events, click **Save**.
- Your changes to the incident are saved.

Generating Incident Reports

LICENSE: Protection

You can use the Sourcefire 3D System to generate incident reports that can include the incident summary, incident status, and any comments along with information from the events you add to the incident. You can also specify whether you want to include event summary information in the report.

To generate an incident report:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Intrusion > Incidents**.
The Incidents page appears.
2. Click the edit icon (✎) next to the incident you want to include in your report.
3. You have two options:
 - To include all the events from the incident in the report, click **Generate Report All**.
 - To include specific events from the incident in the report, select the check boxes next to the events you want and click **Generate Report**.

In either case, the Generate Report page appears, including the options for incident reports.
4. Type a name for the report. You can use alphanumeric characters, periods, and spaces.
5. In **Incident Report Sections**, select the check boxes for the portions of the incident that you want to include in the report: **status**, **summary**, and **comments**.
6. If you want to include event information in the report, select the workflow you want to use and then, in **Report Sections**, specify whether you want to include event summary information.
7. Select the check boxes next to the workflow pages you want to include in the report.
8. Select the check boxes next to the output formats you want to use for the report: **PDF**, **HTML**, and **CSV**.

IMPORTANT! CSV-based incident reports include only event information. They do **not** include the status, summary, or comments from the incident.

9. Click **Generate Report** and confirm that you want to update the report profile.
The report is generated.

Creating Custom Incident Types

LICENSE: Protection

The Sourcefire 3D System is delivered with the following incident types that you can use to classify your incidents:

- Compromise of System Integrity
- Damage
- Denial of Service
- Hoax
- Intrusion
- Theft
- Unauthorized Admin Access
- Unknown
- Web Site Defacement

If these incident types do not meet your needs, you can add your own. Note that you cannot delete any custom incident types.

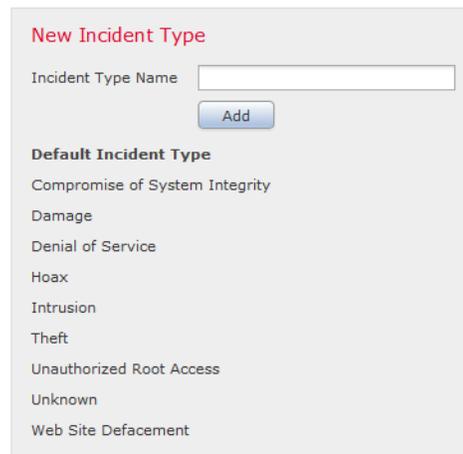
To create a new incident type:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Intrusion > Incidents**.
The Incident page appears.
2. Click **Create Incident**.
The Create Incident page appears.

3. In the **Type** area, click **Types**.

The incident management Types page appears. The default incident types are listed at the bottom of the page.



New Incident Type

Incident Type Name

Add

Default Incident Type

- Compromise of System Integrity
- Damage
- Denial of Service
- Hoax
- Intrusion
- Theft
- Unauthorized Root Access
- Unknown
- Web Site Defacement

4. In the **Incident Type Name** field, type a name for the new incident type. Use alphanumeric characters and spaces.
5. Click **Add**.
The new incident type is added.
6. Click **Done** to close the pop-up window and return to the Incidents page.
You can use the new incident type the next time you create or edit an incident.

CHAPTER 19

CONFIGURING INTRUSION POLICIES

An *intrusion policy* is a defined set of intrusion detection and prevention configurations. You can create an intrusion policy using the settings in the default intrusion policies that Sourcefire provides, or you can tailor your own policies to inspect the traffic that traverses your network. You can modify your intrusion policy to improve performance in your environment and to provide a focused view of the traffic on your network. At a minimum, you should consciously choose whether to configure the following settings in your intrusion policy:

- Specify whether you want to drop packets that trigger rules set to Drop and Generate events in an inline deployment. See [Setting Drop Behavior in an Inline Deployment](#) on page 689 for more information.
- Set your variables to accurately reflect your home and external networks and, as appropriate, the servers on your network. See [Managing Variables](#) on page 698 for more information.

You should also consider whether to take advantage of the following procedures, which can improve performance and better focus your network:

- Disable rules that do not apply to your environment, verify that all rules that **do** apply to your environment are enabled, and set rule attributes such as suppression, thresholding, and alerting. See [Setting Rule States](#) on page 743 for more information.
- Associate hosts and applications on your network with rules written to protect those hosts and applications and recommend rule state changes. See [Managing FireSIGHT Rule State Recommendations](#) on page 764 for more information.

See the following sections for more information:

- [Planning and Implementing an Intrusion Policy](#) on page 668 describes, at a high level, the process you use to create an intrusion policy.
- [Managing Intrusion Policies](#) on page 670 explains how to view a listing of your intrusion policies, and create and edit policies.
- [Setting Drop Behavior in an Inline Deployment](#) on page 689 explains how to set whether your policy drops offending packets for rules set to Drop and Generate Events in an inline deployment.
- [Understanding the Base Policy](#) on page 691 explains how to replace your base policy with a different default intrusion policy provided by Sourcefire or a custom base policy that you create.
- [Managing Variables](#) on page 698 explains how you can use policy-specific and system default variables to tailor your policy to your network.
- [Managing Rules in an Intrusion Policy](#) on page 717 explains how you can enable and disable rules and configure other rule attributes such as thresholds, suppression, and so on.
- [Managing FireSIGHT Rule State Recommendations](#) on page 764 explains how you can generate rule state recommendations for intrusion rules based on the hosts and applications on your network.
- [Using Advanced Settings in an Intrusion Policy](#) on page 773 explains how you can enable, disable, and configure preprocessors and other advanced detection and performance features.
- [Defining IP Addresses and Ports for Your Network](#) on page 711 provides the syntax used to specify IP addresses and port numbers in the variables and rules in your policy.
- [Using Layers in an Intrusion Policy](#) on page 792 explains how you can use intrusion policy layers to more efficiently manage multiple intrusion policies in a complex network environment.

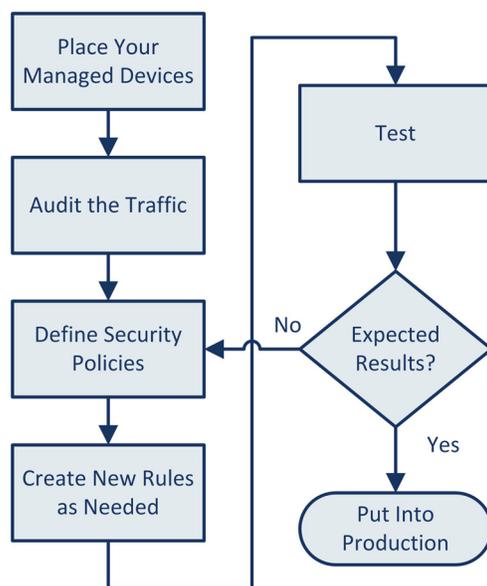
Planning and Implementing an Intrusion Policy

LICENSE: Protection

Building custom intrusion 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.

Traffic profiles and characteristics may change either by design or from the result of malicious action. Sourcefire recommends building a customized intrusion policy to ensure successful monitoring under a wide range of traffic conditions.

The following illustrates the process you use to define your intrusion policy and tune your system.



When planning your intrusion policy:

1. Decide where to place your managed devices.
There are a variety of deployment options in tuning your device. For details on deciding where to place your managed devices to best monitor the traffic that matters to you, see the *Installation Guide* for your device.
2. Understand the traffic that traverses the network segment.
Before tuning your intrusion policy, it pays to understand the traffic it will monitor. For example, if you are monitoring traffic in the DMZ, you may want to pay special attention to web servers and verify that all applicable web server rules are active. If you are monitoring an internal subnet with no external facing servers, you may want to tune your system differently.
3. Define your security policies.
Security policies include your internal security guidelines, as well as your variable, preprocessor, and rules configurations. You should:
 - Define the security guidelines that govern the hosts on that subnet.
Your internal security policies guide how you tune the decoder engine, preprocessor engine, and rules engine. For example, if your security policies prohibit instant messaging, you may want to identify instant message traffic traversing your network.

- Optionally, configure your preprocessors, enabling and disabling options as appropriate.

For more information on the preprocessors provided in Sourcefire 3D System, as well as details on how to configure them, see [Using Advanced Settings in an Intrusion Policy](#) on page 773.

- Set your variables to accurately reflect your home and external networks.

Defining variables makes rule inspection more effective and efficient by directing rules to inspect the traffic to and from specific IP addresses and ports. Setting these at the policy or system level allows you to tune your policy or system without editing every rule. Variables can also be used when suppressing rules and configuring the advanced adaptive profiles feature. For details on managing variables, see [Managing Variables](#) on page 698.

- Disable shared object rules and standard text rules that do not apply to your environment and verify that all rules that **do** apply to your environment are enabled. For inline deployments, carefully choose the intrusion rules that you want to drop packets rather than simply generate events. For more information on setting rule states, see [Setting Rule States](#) on page 743.

4. If none of the existing intrusion rules meet your needs, write new rules that inspect for intrusion attempts.

For information on the rule keywords you can use to construct custom standard text rules, and their syntax, see [Understanding and Writing Intrusion Rules](#) on page 1046.

5. Test your configuration.

Managing Intrusion Policies

LICENSE: Protection

On the Intrusion Policy page (**Policies > Intrusion > Intrusion Policy**) you can view all your current intrusion policies by name with optional description along with the following information:

- the time and date the policy was last modified and the user who modified it.
- whether dropping packets in an inline deployment is enabled in the policy
- when a policy has unsaved changes, in *italicized* black text

Options on this page also allow you to create a new policy, compare two policies or two revisions of the same policy, view a report that lists all of the most recently saved settings in each policy, and edit, delete, or export a policy.

TIP! You can import intrusion policies from other Defense Centers in your deployment. See [Importing and Exporting Objects](#) on page 2225 for more information.

Note that the Intrusion Policy page displays the time a policy was last modified in local time, but intrusion policy reports list the time the policy was last modified in Coordinated Universal Time (UTC).

The [Intrusion Policy Management Actions](#) table describes the actions you can take to manage your policies on the Intrusion Policy page:

Intrusion Policy Management Actions

To...	You CAN...
compare the settings of two intrusion policies or two revisions of the same policy	click Compare Policies . Optionally, then click Comparison Report to open or save a PDF version of the report, or click New Comparison to begin a new comparison. See Comparing Two Intrusion Policies on page 685 for more information.
create a new intrusion policy	click Create Policy . See Creating an Intrusion Policy on page 672 for more information.
reapply an intrusion policy to your managed devices	click the apply icon (✓). See Reapplying an Intrusion Policy on page 680 for more information.
view a PDF report that lists the current configuration settings in an intrusion policy	click the report icon (📄). See Viewing an Intrusion Policy Report on page 682 for more information.
export an intrusion policy for use on another appliance of the same type	click the export icon (📎). See Exporting Objects on page 2226 for more information.

Intrusion Policy Management Actions (Continued)

To...	You CAN...
modify the settings in an existing intrusion policy	click the edit icon (✎). See Editing an Intrusion Policy on page 675 for information.
delete an intrusion policy	click the delete icon (🗑), then click OK , or click Cancel if you decide not to delete the policy. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy. You cannot delete an intrusion policy if an access control policy references it.

See the following sections for more information:

- [Creating an Intrusion Policy](#) on page 672
- [Editing an Intrusion Policy](#) on page 675
- [Using the Navigation Panel](#) on page 678
- [Committing Intrusion Policy Changes](#) on page 679
- [Reapplying an Intrusion Policy](#) on page 680
- [Viewing an Intrusion Policy Report](#) on page 682
- [Comparing Two Intrusion Policies](#) on page 685

To manage your intrusion policies:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

Compare Policies Create Policy

Intrusion Policy	Drop when Inline	Status	Last Modified	
Initial Inline Policy - katsura Default policy	Yes	Policy not applied on any devices	2012-04-06 12:43:17 Modified by "admin"	✔️ 📄 ✎ 🗑
Initial Passive Policy - katsura Default policy	No	Policy not applied on any devices	2012-04-06 12:43:15 Modified by "admin"	✔️ 📄 ✎ 🗑

2. Take any of the actions described in the [Intrusion Policy Management Actions table](#) on page 671.

Creating an Intrusion Policy

LICENSE: Protection

You can create one or more intrusion policies. For example, you can create policies that monitor traffic on your network. You can also create policies that you use for testing in a safe network environment, or for familiarizing yourself with

features such as FireSIGHT Recommended Rules or the different default policies provided by Sourcefire.

When you create a policy, a pop-up window provides immediate access to the features you are most likely to configure. You can create your intrusion policy using only the options in the pop-up window, or you can save your changes and continue to the advanced intrusion policy editor, where you can configure any intrusion policy features.

TIP! You can import intrusion policies from other Defense Centers in your deployment. See [Importing and Exporting Objects](#) on page 2225 for more information.

To create an intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

2. Click **Create Policy**.

If you have unsaved changes in another policy, click **Cancel** when prompted to return to the Intrusion Policy page. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Create Intrusion Policy pop-up window appears.

The screenshot shows a 'Create Intrusion Policy' dialog box. It contains the following fields and controls:

- Policy Information:**
 - Name * (text input, highlighted with a red box)
 - Description (text input)
 - Drop when Inline (checkbox, checked)
 - Base Policy (dropdown menu, showing 'Connectivity Over Security')
- Variables:**
 - Use the system default value (checkbox, checked)
 - Networks to protect (text input, containing 'any')
- Buttons:** 'Create Policy', 'Create and Edit Policy', and 'Cancel'.
- Legend:** '* Required'

3. Type a unique name of 50 characters or less that identifies your policy and, optionally, a description that differentiates it from other policies.

4. Specify whether you want the system to drop the packet and generate an event when a packet triggers a rule set to Drop and Generate Events in an inline deployment:
 - To drop the packet and generate an event, select the **Drop when Inline** check box.
 - To generate an event but not drop the packet, clear the **Drop when Inline** check box.

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. For more information, see [Setting Rule States](#) on page 743, [Setting Drop Behavior in an Inline Deployment](#) on page 689, [Configuring Passive Interfaces](#) on page 267, and [Tap Mode](#) on page 275.

5. Optionally, select a different Sourcefire default or custom policy that you want to use as the base policy for your intrusion policy from the **Base Policy** drop-down list. See [Understanding the Base Policy](#) on page 691 for more information.

6. Select or clear the **Use the system default value** check box to specify which hosts on your network you want the policy to monitor:

- Select the check box to monitor the hosts specified by default for your system.

Note that the system monitors all traffic on your network unless you reconfigure the system default value for the \$HOME_NET variable. See [Understanding Existing Variables](#) on page 699 and [Modifying Variables](#) on page 701 for more information.

- Clear the check box to use the **Networks to protect** field to identify specific hosts to monitor.

Note that this option sets a policy-specific value for the \$HOME_NET variable, which overrides the system default value of \$HOME_NET.

You can specify **any** (that is, any IPv4 or IPv6 address), or you can specify one or more IP addresses, address blocks, or variables, including negation (!). For more information, see [Managing Variables](#) on page 698, [Defining IP Addresses and Ports for Your Network](#) on page 711, and [IP Address Conventions](#) on page 57.

7. You have the following options:
 - To exit the pop-up window without creating a policy, click **Cancel**. The Intrusion Policy page appears.

- To save your changes, click **Create Policy**.
The Intrusion Policy page appears.
You must apply the appropriate access control policy to put your changes into effect. See [Applying an Access Control Policy](#) for more information.
- To open the advanced intrusion policy editor, click **Create and Edit Policy**.
See [Editing an Intrusion Policy](#) on page 675 for more information.

Editing an Intrusion Policy

LICENSE: Protection

You can use the advanced intrusion policy editor to configure any intrusion policy feature. You can configure most commonly used settings on or directly from the Policy Information page. For information on more advanced intrusion policy features, see [Using Advanced Settings in an Intrusion Policy](#) on page 773 and [Using Layers in an Intrusion Policy](#) on page 792.

The following table explains the most common actions taken when editing an intrusion policy:

Common Intrusion Policy Editing Actions

To...	You CAN...
specify a different drop behavior in an inline deployment	select or clear the Drop when Inline check box. See Setting Drop Behavior in an Inline Deployment on page 689 for more information.
select a different base policy	click Select Base Policy. See Understanding the Base Policy on page 691 for more information.
view the advanced settings that are enabled by default in your base policy	click Manage Base Policy . See Using Advanced Settings in an Intrusion Policy on page 773 for more information.

Common Intrusion Policy Editing Actions (Continued)

To...	You CAN...
tailor variables for your specific network environment	click Manage Variables . See Managing Variables on page 698 for more information. Make sure you accurately specify the values for \$HOME_NET, which represents the IP addresses on the network segment protected by the system, and \$EXTERNAL_NET, which represents all the IP addresses that are not on the network segment. For example, if you set the rule state to Drop and Generate Events for a rule that uses these variables and the variables themselves are not set correctly, your policy could drop the wrong packets or pass packets that you expect to drop.
display or modify configured rule attributes for the rules in your intrusion policy	click Manage Rules . See Managing Rules in an Intrusion Policy on page 717 for more information.
display a filtered view of the intrusion policy Rules page showing rules enabled in your policy by current rule state and, optionally, set rule attributes for specified rules	click View next to the number of rules under Manage Rules that are set to Generate Events or to Drop and Generate Events. See Managing Rules in an Intrusion Policy on page 717 for more information.
display the FireSIGHT Recommended Rules configuration page	click FireSIGHT Recommendations in the navigation panel. Alternately, click Click here to set up FireSIGHT recommendations on the Policy Information page if you have not generated recommendations, or Click to change recommendations if you have generated recommendations. See Managing FireSIGHT Rule State Recommendations on page 764 for more information.
display a filtered view of the Rules page showing rules with recommended rule states and, optionally, can set rule attributes for specified rules	click View next to the number of recommendations to generate events, drop and generate events, or disable rules, or click View Recommended Changes to view all recommendations. These options appear only when you have generated recommendations. See Managing FireSIGHT Rule State Recommendations on page 764 for more information.

Common Intrusion Policy Editing Actions (Continued)

To...	You CAN...
edit advanced settings	click Advanced Settings in the navigation panel. See Using Advanced Settings in an Intrusion Policy on page 773 for more information.
revert advanced settings configuration to default configuration settings in the base policy layer	click Revert to Defaults on an advanced settings configuration page, then click OK at the prompt. See Understanding the Base Policy on page 691 for more information.
manage policy layers	click policy layers in the navigation panel. See Using Layers in an Intrusion Policy on page 792 for more information.
save changes to your policy	click Commit Changes . You must apply the appropriate access control policy to put your changes into effect. See Editing an Intrusion Policy on page 675, Committing Intrusion Policy Changes on page 679, and Applying an Access Control Policy for more information.
discard all unsaved changes	click Discard Changes , then click OK to discard your changes and go to the Intrusion Policy page, or click Cancel to keep your changes and return to the Policy Information page.
exit the policy, leaving changes to the policy in the system cache	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. See Committing Intrusion Policy Changes on page 679 for information on how the system caches one policy per user.

To edit an intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
 The Intrusion Policy page appears.

2. Click the edit icon (✎) next to the policy you want to edit, or the intrusion policy name.

The Policy Information page appears.



3. You can take any of the actions described in [Editing an Intrusion Policy](#) on page 675.
4. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Using the Navigation Panel

LICENSE: Protection

A navigation panel appears on the left side of the web interface when you are editing an intrusion policy.



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.

The two major links above the dividing line separate intrusion policy settings into Policy Information (the most commonly used settings) and Advanced Settings (settings that typically require little or no modification, and require specific expertise to configure).

Click **Policy Information** to display the Policy Information page, which includes configuration options for commonly used settings and links to configuration

pages for other commonly used settings. Sublinks beneath **Policy Information** provide direct access to the same configuration pages.

Click **Advanced Settings** to display the Advanced Settings page, where you can enable or disable advanced settings and access configuration pages for advanced settings in your intrusion policy. Note that you cannot access advanced intrusion policy settings from the Policy Information page.

Expanding the **Advanced Settings** link displays sublinks to individual configuration pages for all advanced settings that are enabled in your intrusion policy. Clicking any of these sublinks takes you to the same advanced settings configuration pages that you can access from the Advanced Settings page. See [Using Advanced Settings in an Intrusion Policy](#) on page 773 for more information.

You can click **Policy Layers** to display a summary of the intrusion policy layers that comprise your intrusion policy. Expanding the Policy Layers link displays sublinks to summary pages for the layers in your intrusion policy. Expanding each layer sublink displays further sublinks to the configuration pages for all advanced settings that are enabled in the layer, and to a layer-filtered view of intrusion rule settings. See [Using Layers in an Intrusion Policy](#) on page 792 and [Managing Rules in an Intrusion Policy](#) on page 717 for more information.

Dark shading of an item in the navigation panel highlights your current location in the intrusion policy. For example, in the illustration above the Policy Information page would be displayed to the right of the navigation panel.

A policy change icon (⚠) appears next to **Policy Information** when your intrusion policy contains unsaved changes. This icon disappears when you save your changes from the Policy Information page. You can click the policy change icon or **Policy Information** to display the Policy Information page.

Committing Intrusion Policy Changes

LICENSE: Protection

You must save (that is, commit) changes to your intrusion policy before the system recognizes the changes. When you associate an intrusion policy with an access control policy, the system associates the most recently saved configuration. See [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511 for more information.

The system caches changes to your policy on the system disk when you exit the policy without saving your changes. The system cache stores unsaved changes for one policy per user and you must commit or discard your changes before editing another policy when you are logged in as the same user.

Your changes are cached even when you log out of the system or experience a system crash. The system discards the cached changes when you edit another policy as the same user without saving your changes, or when you import a rule update. See [Importing Rule Updates and Local Rule Files](#) on page 2072 for more information.

The following might also occur when you commit your changes:

- If the **Write changes in Intrusion Policy to audit log** Intrusion Policy Preferences option in the system policy is enabled, the system logs a description of the changes in the audit log. See [Editing a System Policy](#) on page 1963 and [Viewing Audit Records](#) on page 2188 for more information.
- Depending on the configuration of the **Comments on policy change** Intrusion Policy Preferences option in the system policy, the Description of Changes pop-up window might appear when you save your changes, and you might be required to provide a description of your changes. Optionally or if required, provide a description of your changes, then click **OK** to save your changes, or click **Cancel** to return to the advanced editor without saving your changes. See [Configuring Intrusion Policy Preferences](#) on page 1985 for more information.
- If your configuration includes a standard text rule or a shared object rule that requires a disabled preprocessor or other advanced feature, click **OK** when prompted to automatically enable the feature in your policy and commit the policy. Click **Cancel** to return to the Policy Information page. See [Automatically Enabling Advanced Settings](#) on page 787 for more information.
- If your intrusion policy contains a system default variable with the invalid entry `!any`, a pop-up window appears and notifies you. For more information about correctly defining system default variables, see [Managing Variables](#) on page 698.
- If you are editing a policy at the same time another user is editing the same policy, and the other user saves their changes to the policy, you are warned when you commit the policy that you will overwrite the other user's changes. Click **OK** to continue and overwrite the changes, or click **Cancel** to return to the Policy Information page without saving your changes.
- If you are editing the same policy via multiple web interface instances as the same user, and you save your changes for one instance, you are prompted for any other instance if you try to commit the policy where you cannot save your changes. Click **OK** to discard your changes and go to the Intrusion Policy page.

Reapplying an Intrusion Policy

LICENSE: Protection

After you apply an intrusion policy to a managed device using access control (see [Applying an Access Control Policy](#) on page 466), 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](#) on page 1938 for more information.
- An intrusion policy reapply fails on invalid target devices. For example, if you apply an access control policy that removes a previously applied intrusion policy from a device and then attempt to reapply the intrusion policy before the access control policy apply task resolves, the intrusion policy reapply fails.
- You cannot apply intrusion policies to stacked devices running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). You can reapply an intrusion policy to a device stack, but not to individual devices within the stack. See [Managing Stacked Devices](#) on page 237 for more information.
- 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](#) on page 2072 for more information.

To reapply an intrusion policy:

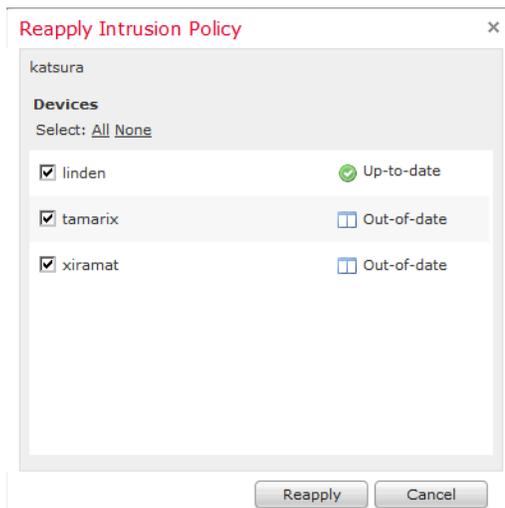
ACCESS: Admin/Security Approver

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

Intrusion Policy	Drop when Inline	Status	Last Modified	
Example Policy 1 Example Intrusion Policy	Yes	Policy not applied on any devices	2012-04-10 09:36:37 Modified by "admin"	   
Example Policy 2 Example Intrusion Policy	Yes	Policy out-of-date on 3 of 3 devices	2012-04-11 12:05:30 Modified by "admin"	   
Example Policy 3 Example Intrusion Policy	Yes	Policy not applied on any devices	2012-04-11 12:01:35 Modified by "admin"	   
Initial Inline Policy - katsura Default policy	Yes	Policy up-to-date on all 3 devices	2012-04-06 12:43:17 Modified by "admin"	   
Initial Passive Policy - katsura Default policy	No	Policy not applied on any devices	2012-04-06 12:43:15 Modified by "admin"	   

- Click the apply icon (✓) next to the policy you want to reapply.
The Reapply Intrusion Policy window appears, listing the devices where the policy is currently applied.



- Specify the devices where you want to reapply the policy.

TIP! Optionally, if a device is listed as **Out-of-date**, click the comparison icon (□) to view a report that compares the currently applied intrusion policy and the updated intrusion policy. See the [Comparing Two Intrusion Policies table](#) on page 685 for more information.

- Click **Reapply**.

The policy is reapplied. You can monitor the status of the apply using the task queue (**System > Monitoring > Task Status**). See [Viewing the Task Queue](#) on page 2238 for more information.

Viewing an Intrusion Policy Report

LICENSE: Protection

An intrusion policy report is a record of all enabled intrusion policy features and settings 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 use the report for auditing purposes or to inspect the current configuration of an intrusion policy.

Remember to commit any potential changes before you generate an intrusion policy report; only committed changes appear in the report.

TIP! You can also generate an intrusion policy comparison report that compares two intrusion policies, or two revisions of the same intrusion policy. For more information, see [Comparing Two Intrusion Policies](#) on page 685.

Depending on your configuration, an intrusion policy report can contain one or more sections as described in the [Intrusion Policy Report Sections](#) table.

Intrusion Policy Report Sections

SECTION	DESCRIPTION
Title Page	Identifies the name of the intrusion policy report, the date and time the intrusion policy was last modified, and the name of the user who made that modification. Note that the Intrusion Policy Report lists the Last Modified time in UTC, but the Intrusion Policy page lists the modified time in local time.
Table of Contents	Describes the contents of the report. Only enabled intrusion policy features appear on the report. For example, if the DNS Configuration feature is not enabled in your intrusion policy, it does not appear in the table of contents or in the report.
Policy Information	Provides the name and description of the intrusion policy, whether dropping packets in an inline deployment is enabled or disabled, current rule update version, whether the base policy is locked to the current rule update, the date and time the intrusion policy was last modified, and the name of the user who made that modification. See Editing an Intrusion Policy on page 675.
Variables	Lists any policy-specific variables. See Managing Variables on page 698.
FireSIGHT Recommendations	Provides information on any recommended rule states based on the hosts and applications in your network. Optionally, you can set your intrusion policy to Include all differences between recommendations and rule states in policy reports . See Managing FireSIGHT Rule State Recommendations on page 764.

Intrusion Policy Report Sections (Continued)

SECTION	DESCRIPTION
Advanced Settings	Lists all advanced feature settings (such as Checksum Verification, DCE/RPC Configuration, and so on) and their configurations (such as enabled, default, stateful, and so on). See Using Advanced Settings in an Intrusion Policy on page 773.
Rules	Provides a list of all enabled rules (such as Backdoor — Dagger, DDOS TFN Probe, and so on) and their actions (such as Generate events, Drop and generate events, and so on). See Managing Rules in an Intrusion Policy on page 717.

The following sample graphic displays the Advanced Settings section of an example intrusion policy report, and lists the configuration for each advanced setting. Other sections listed in the table of contents use the same format and provide the same level of detail for their respective sections.

Advanced Settings

Back Orifice Detection

There are no user-configurable options for Back Orifice

Checksum Verification

ICMP Checksums	Enable
IP Checksums	Enable
TCP Checksums	Enable
UDP Checksums	Enable

DCE/RPC Configuration

Global Settings	
Maximum Fragment Size	
Reassembly Threshold	0
Enable Defragmentation	Enabled
Memory cap reached	Disabled
SMB traffic	Enabled
Connection-oriented DCE/RPC traffic	Enabled
Connectionless DCE/RPC traffic	Enabled
Servers	
default	
Networks	default
Policy	WinXP
SMB Invalid Shares	
SMB Maximum AndX Chain	3
RPC over HTTP proxy Ports	
RPC proxy traffic only	Disabled

To view an intrusion policy report:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy.**

The Intrusion Policy page appears.

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. 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 Two Intrusion 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 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 or intrusion policy revisions:

- 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 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 web 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](#) on page 686
- [Using the Intrusion Policy Comparison Report](#) on page 687

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.

Example Policy 1 (2011-10-10 10:12:16 by admin)	Example Policy 2 (2011-10-10 10:13:01 by admin)
Policy Information	Policy Information
Name Example Policy 1	Name Example Policy 2
Modified 2011-10-10 10:12:16 by adi	Modified 2011-10-10 10:13:01 by adi
Base Policy Connectivity Over Security	Base Policy Security Over Connectivity
Advanced Settings	Advanced Settings
Back Orifice Detection Enabled	Back Orifice Detection Disabled
Servers default	Servers default
Client Flow Depth 300	Client Flow Depth 0
Server Flow Depth 300	Server Flow Depth 1460
Latency-Based Packet Handling Enabled	Latency-Based Packet Handling Disabled
Threshold 128	Threshold 2000000
Latency-Based Rule Handling Enabled	Latency-Based Rule Handling Disabled
Threshold 256	
Consecutive Threshold Violations Before 5	
Suspension Time 30	
Sun RPC Configuration Disabled	Sun RPC Configuration Enabled
Logged Events 1	Logged Events 3
Rules	Rules

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 in the [Intrusion Policy Comparison View Actions](#) table.

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.
determine which layer contains the configuration for a specific advanced setting	hover over the advanced configuration icon (⚙) next to the configuration you want to view. The window displays the name of the layer that contains the advanced configuration. See Using Layers in an Intrusion Policy on page 792 for more information.
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 intrusion policy comparison report creates a PDF that lists only the differences between the two intrusion policies or intrusion policy versions.

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.

The following sample graphic displays the Advanced Settings section of an intrusion policy comparison report, and lists the configuration for each advanced setting for both intrusion policy configurations. Each section uses the same format and provides the same level of detail. Note that the Value A and Value B columns represent the policies or policy revisions you configured in the comparison view.

IP Defragmentation

Field	Value A	Value B
IP Defragmentation > Servers > default > Overlap Limit		10
IP Defragmentation > Servers > default > Minimum Fragment Size		100

Regular Expression Limits

Field	Value A	Value B
Regular Expression Limits > Match Limit	1500	3500

You use a similar procedure to compare other types of policies on the Sourcefire 3D System. For more information, see:

- [Comparing System Policies](#) on page 1965
- [Comparing Health Policies](#) on page 2150

To compare two intrusion policies or two revisions of the same policy:

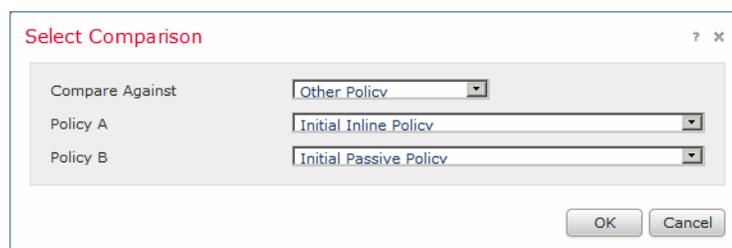
ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

2. Click **Compare Policies**.

The **Select Comparison** window appears.



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.

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.
5. Click **OK** to display the intrusion policy comparison view.
The comparison view appears.
6. Click **Comparison Report** to generate the intrusion policy comparison report.
7. The intrusion policy 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.

Setting Drop Behavior in an Inline Deployment

LICENSE: Protection

A *drop rule* is an intrusion rule or preprocessor rule whose rule state is set to Drop and Generate Events. You can use the **Drop when Inline** option in your intrusion policy to determine how the system handles drop rules in an inline deployment; see [Setting Rule States](#) on page 743 for information on setting rule states in your intrusion policy.

In an inline deployment, you would typically set your intrusion policy to drop packets that trigger drop rules. However, you might also set your policy to not drop packets so you can assess how your configuration functions on your network. In this case, the system would generate events but would not drop packets that trigger your drop rules. When you are satisfied with the results, you can set your policy to drop packets; then you can reapply the access control policy that includes your policy.

When you set your intrusion policy to drop packets in an inline deployment, the system drops packets that trigger enabled drop rules and generates events for the triggered rules.

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 file transfers while using 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.

Note that in a passive deployment, including when an inline interface is in tap mode, the system treats rules set to Drop and Generate Events the same as rules set to Generate Events; that is, the system generates events but does not drop packets that trigger the rules regardless of the drop behavior of your policy. See

[Tap Mode](#) on page 275 for more information.

Note also that the table view of intrusion events indicates when packets are dropped if **Drop when Inline** is enabled in an inline deployment, and when packets would have dropped if **Drop when Inline** is disabled. In a passive deployment, including when an inline interface is in tap mode, the table view of intrusion events always shows that drop rules would have dropped packets in a inline deployment, regardless of the setting for **Drop when Inline**. See [Understanding Intrusion Events](#) on page 604 for more information.

The [Drop Rule Behavior](#) table summarizes drop rule behavior in passive and inline deployments.

Drop Rule Behavior

WHEN THE DEPLOYMENT IS...	AND DROP WHEN INLINE IS...	OFFENDING PACKETS ARE...	AND THE EVENT TYPE IS...
inline	enabled	dropped	Dropped
inline	disabled	not dropped	Would have dropped
inline (tap mode)	enabled	not dropped	Would have dropped
inline (tap mode)	disabled	not dropped	Would have dropped
passive	enabled	not dropped	Would have dropped
passive	disabled	not dropped	Would have dropped

Note that setting what is called a *pass rule* to Generate Events has a different effect. For information, see [Specifying Rule Actions](#) on page 1050.

Note also that your inline intrusion policies can include rules that use the `replace` keyword. For information, see [Replacing Content in Inline Deployments](#) on page 1078.

To set the drop behavior of your intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Specify whether you want the system to drop the packet and generate an event when the packet triggers a rule set to Drop and Generate Events in an inline deployment:
 - To drop the packet and generate an event, select the **Drop when Inline** check box.
 - To generate an event but not drop the packet, clear the **Drop when Inline** check box.

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. For more information, see [Setting Rule States](#) on page 743, [Setting Drop Behavior in an Inline Deployment](#) on page 689, [Configuring Passive Interfaces](#) on page 267, and [Tap Mode](#) on page 275.

TIP! On 3D9900 and Series 3 devices, an inline set can use *tap mode*, which allows you to passively monitor traffic.

4. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions](#) table on page 675 for more information.

Understanding the Base Policy

LICENSE: Protection

The base policy in an intrusion policy defines the default settings for all rules and advanced settings in the policy. You can use a default policy provided by the Sourcefire Vulnerability Research Team (VRT) as your base policy, or you can use a custom policy that you create as your base policy.

Note the following important information regarding base policies:

- The base policy includes configurations for rules and advanced settings. It does not include variables or FireSIGHT Recommended Rules.
- Modifying a rule or advanced setting in your policy overrides the corresponding default setting in the base policy.
- The base policy is the lowest layer in an intrusion policy. For information on using policy layers to more effectively manage multiple intrusion policies, see [Using Layers in an Intrusion Policy](#) on page 792.
- Depending on your configuration, importing rule updates may modify settings in your base policy. However, changes that a rule update makes to your base policy do not override changes that you make to rules or advanced settings in your policy. See [Importing Rule Updates and Local Rule Files](#) on page 2072 for more information.

See the following sections for more information:

- [Using Default Intrusion Policies](#) on page 692
- [Using a Custom Base Policy](#) on page 693
- [Allowing Rule Updates to Modify the Base Policy](#) on page 694
- [Selecting the Base Policy](#) on page 695
- [Accepting Rule Setting Changes from a Custom Base Policy](#) on page 696

Using Default Intrusion Policies

LICENSE: Protection

Five default intrusion policies are delivered with the Sourcefire 3D System. You can use four of these default policies. Sourcefire uses the fifth, Experimental Policy 1, for testing purposes and you should not use it unless instructed to do so by a Sourcefire representative.

The Sourcefire Vulnerability Research Team (VRT) sets the state of each intrusion and preprocessor rule in each default policy. The VRT also sets the default state, enabled or disabled, of each preprocessor and of other advanced features, and the default option settings for each. For example, a rule might be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. Intrusion protection features in an intrusion policy you create inherit the default settings in a default policy that you use to create your policy. By using the policies provided by Sourcefire as a basis for your intrusion policy, you can take advantage of the experience of the VRT.

The default intrusion policies that you can use are:

- **Balanced Security and Connectivity**
This policy is built for both speed and detection. It serves as a good starting point for most organizations. It is also a good starting point for any type of deployment.
- **Connectivity Over Security**
This policy is built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. This policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

- **No Rules Active**
All intrusion rules, preprocessors, and other configurable intrusion policy features in this policy are disabled by default. This policy provides a starting point if you want to create your own policy instead of basing it on the enabled rules and features in one of the other policies provided by Sourcefire. The system automatically enables any preprocessor required by rules you enable.
Note that all rules and most preprocessors and other advanced features are disabled in this policy.
- **Security Over Connectivity**
This policy is built for organizations where network infrastructure security takes precedence over user convenience. This policy enables numerous network anomaly rules that could alert on or drop legitimate traffic.

You can use copies of Sourcefire default policies or create your own policies with tuned rule sets and advanced settings configurations to inspect traffic in the way that matters most to you. By doing this, you can improve both the performance of your managed device and your ability to respond effectively to the events it generates.

Note that the following initial policies, which come with your system, are custom policies provided by Sourcefire; they are not default policies:

- Initial Inline Policy
- Initial Passive Policy

Each of these custom policies uses a default policy as its base policy.

Using a Custom Base Policy

LICENSE: Protection

Custom policies include policies you create and the following two initial policies that come with your system:

- Initial Inline Policy
- Initial Passive Policy

You can use a custom policy as your base policy. Changes that you make to rules and advanced settings in a custom policy are automatically included in your base policy when you commit, that is, save changes in, the custom policy. However, you can override a default setting by modifying it in the policy that uses the custom policy as its base policy.

In a custom base policy, you do not have the option of allowing rule updates to modify the base policy. However, in some cases importing a rule update may impact the custom base policy when the parent policy, that is, the original policy that you use as your custom base policy, allows rule updates to modify its base policy. See [Allowing Rule Updates to Modify the Base Policy](#) on page 694 for more information.

Note you must create a policy-specific variable in both policies to use it in both. See [Creating New Policy-Specific Variables](#) on page 705 for more information.

Allowing Rule Updates to Modify the Base Policy

LICENSE: Protection

Rule updates that you import provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default intrusion policy settings. Rule updates can also delete rules and provide new rule categories and system variables. See [Importing Rule Updates and Local Rule Files](#) on page 2072 for more information.

Rule updates always modify the default policies provided by Sourcefire with any changes that a rule update makes to rules and advanced settings. Changes to variables and rule categories are handled at the system level. See [Using Default Intrusion Policies](#) on page 692, [Understanding Existing Variables](#) on page 699, and [Understanding Rule Categories](#) on page 739 for more information.

When you use a default policy provided by Sourcefire as your base policy, you can choose whether to allow rule updates to modify your base 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 rules and advanced settings in the default policy that you use as your base policy. If you have not modified the corresponding setting, the setting in your base policy determines the setting in your policy. However, a new rule update will not override any changes you have made 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.

Note that rule updates always delete rules that VRT deletes, regardless of the rule state in your policy or whether you allow rule updates to update your base policy. Until you reapply an access control policy that includes your policy after a rule update deletes a rule, rules in your currently applied intrusion policies will behave as follows:

- Disabled rules will remain disabled.
- Rules set to Generate Events will continue to generate events when triggered.
- Rules set to Drop and Generate Events will continue to generate events and drop offending packets when triggered.

Note also that, in a custom base policy, you do not have the option of allowing rule updates to modify the base policy, because in this case the base policy is not

a default policy provided by Sourcefire. However, a rule update can modify the custom base policy when both of the following conditions are met:

- You allow rule updates to modify the 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 rule, and you have not modified the rule's state in the parent policy, the modified rule state will be passed to custom base policy when you save the parent policy. See [Using a Custom Base Policy](#) on page 693 for more information.

Selecting the Base Policy

LICENSE: Protection

You can select the base policy for your intrusion policy and, when your base policy is a default policy provided by Sourcefire, choose whether to allow rule updates to update your base policy on the Base Policy summary page. You can also view but not change the default state, enabled or disabled, of preprocessors and other advanced features. From this page, you can access the configuration pages for advanced features where you can view but not change their default option settings. You can also access a read-only display of the Rules page, where you can view the default states of all rules in your base policy, filter the display to view subset of rules, and view details of individual rules.

To select the base policy in your intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Manage Base Policy** on the Policy Information page.
The Base Policy summary page appears.
4. Select the Sourcefire default or custom policy that you want to use as the base policy for your intrusion policy from the **Base Policy** drop-down list. See [Understanding the Base Policy](#) on page 691 for more information.

5. Optionally, select or clear the **Update when a new Rule Update is installed** check box to specify whether you want new rule updates to update your base policy.

When you save your changes 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.

See [Allowing Rule Updates to Modify the Base Policy](#) on page 694 for more information.

6. Optionally, take any of the following actions on the page:
 - To display all rules in your base policy on the Rules page in read-only mode, click **View Rule**.
In the read-only display in this page, you can filter the view to display subsets of rules in your base policy. You can also display details of individual rules. See [Managing Rules in an Intrusion Policy](#) on page 717 for more information.
 - To view which preprocessors and other advanced features are enabled or disabled in your base policy, scroll down the page. See [Using Advanced Settings in an Intrusion Policy](#) on page 773 for more information.
 - To display the configuration page and default settings for an advanced feature in read-only mode, click **View** next to the feature whose default settings you want to see. For an overview of advanced features that you can enable or disable and whose default settings you can modify, see [Using Advanced Settings in an Intrusion Policy](#) on page 773.
7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Accepting Rule Setting Changes from a Custom Base Policy

LICENSE: Protection

When you set event filters, dynamic states, and alerting for selected rules in a custom policy that you use as your base policy, then remove those settings in the policy that uses the custom policy as its base policy, your intrusion policy ignores subsequent setting changes that you make to the affected rules in the custom policy you use as your base policy.

The following procedure explains how to set a policy where you have not added layers to accept changes to rule settings that you make in the custom policy that you use as your base policy. See [Removing Multi-Layer Rule Settings](#) on page 797 to accept settings for these rules in a policy where you have added layers.

To accept rule setting changes in a policy where you have not added layers:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
2. Click the edit icon () next to the intrusion policy where you want to unblock settings.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Expand **Policy Layers** in the navigation panel.
4. Expand the link beneath **Policy Layers**, which is named **My Changes** if you have not renamed it.
5. Click **Rules** beneath **My Changes**.
The Rules page for My Changes appears.
6. Locate the rule or rules whose settings you want to accept. 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 [Understanding Rule Filtering in an Intrusion Policy](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.The page refreshes to display all matching rules.
7. 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.
8. Select **Inherit** from the **Rule State** drop-down list.
9. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Managing Variables

LICENSE: Protection

Variables represent values that are commonly used within intrusion rules, and when configuring some features such as rule suppressions and the advanced adaptive profiles feature. You can view, modify, and create variables. You can also delete variables if they are not used in any active or inactive rules.

Variables are important for several reasons. The following list describes why it is important to review, modify, and even add new variables:

- Most of the shared object rules and standard text rules that the Sourcefire 3D System provides use predefined variables to define networks and port numbers. The majority of the rules use these variables to specify the protected network (`$HOME_NET`) and the unprotected (or outside) network (`$EXTERNAL_NET`). 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. By ensuring that a variable such as `$HOME_NET` correctly defines your entire network and `$HTTP_SERVERS` includes all web servers on your network, you can be sure that processing is optimized and all relevant systems are monitored for suspicious activity.
- If you create custom standard text rules, you can use variables 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.
- Variables make it easier to edit preexisting shared object rules, because values are not hard-coded in many of these. You only have to change the variable value, rather than creating your own rules. This is also effective when you have multiple policies in place, as you can maintain multiple variable definitions that are unique to each policy.

See the following sections for more information:

- [Understanding Existing Variables](#) on page 699 describes the predefined system default variables provided with Sourcefire 3D System and includes information about how and why to modify them.
- [Modifying Variables](#) on page 701 explains the syntax used for IP- and port-based variables, and describes how to modify existing variables and variables that you create.
- [Creating New Variables](#) on page 703 describes how to create new variables for use in rules and system configuration.

- [Deleting Unused Variables](#) on page 709 describes how to delete unused variables.
- [Understanding Custom Variables](#) on page 710 describes custom variables you can use to configure intrusion protection capabilities not otherwise configurable via the web interface.

Understanding Existing Variables

LICENSE: Protection

The Sourcefire 3D System provides predefined *system default variables* for use within rules. You should set appropriate values for these variables to optimize system performance across your network. Variables can use values that include:

- **any**, to specify any value
- specific IP addresses or ports (see [Defining IP Addresses and Ports for Your Network](#) on page 711 for more information about acceptable syntax for IP addresses and ports in variables)
- other variables
- a negation of any of the above (excluding **any**), using **!**

IMPORTANT! The **SNORT_BPF** variable, which appears in the list of system default variables, is a custom variable used to enable an advanced intrusion detection feature. Because its function differs from the other default variables, it is not described in the table that follows. For more information on the **SNORT_BPF** variable, see [Understanding Custom Variables](#) on page 710.

The [Variable Descriptions](#) table describes the predefined variables.

Variable Descriptions

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.	Not required in current rule set
\$EXTERNAL_NET	Defines the network that the Sourcefire 3D System 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 use !\$HOME_NET as the value for \$EXTERNAL_NET

Variable Descriptions (Continued)

VARIABLE NAME	DESCRIPTION	MODIFY?
\$FILE_DATA_PORTS	Defines non-encrypted ports used in intrusion rules that detect files in a network stream.	Not required
\$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 active policy monitors, and is used in many rules to define the internal network.	Yes, to include the IP address range 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 (which you can view in the web interface)
\$HTTP_SERVERS	Defines the web servers on your network. Used in web server exploit rules.	Yes, if you run HTTP servers
\$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
\$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
\$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 use \$HOME_NET as the value for \$SIP_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
\$SQL_SERVERS	Defines database servers on your network, and is used in rules that address database-targeted exploits.	Yes, if you run SQL servers

Variable Descriptions (Continued)

VARIABLE NAME	DESCRIPTION	MODIFY?
\$SSH_PORTS	Defines the ports of SSH servers on your network, and is used for web server exploit rules.	Yes, if your SSH servers use ports other than the default port (which you can view in the web 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 use \$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

Modifying Variables

LICENSE: Protection

Modifying variables is an essential step in tuning your intrusion policy. By customizing predefined variables to closely match your network configuration, you can optimize system performance and ensure that the system evaluates relevant network traffic.

You can also modify variables that you create. For information on creating variables, see [Creating New Variables](#) on page 703.

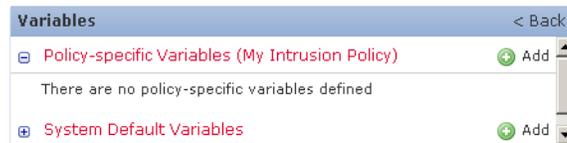
IMPORTANT! For more information about pre-configured variables and what they are used for, see [Understanding Existing Variables](#) on page 699.

To modify an existing variable:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.

3. Click **Manage Variables** on the Policy Information page.
The Variables page appears.



4. You have the following options:
 - Click the + or - sign to the left of the page section name for a variable type to expand or collapse a list of currently configured variables for that type.
 - Click the delete icon (🗑️) next to a user-defined variable to delete the variable.
 - Click the reset icon (↩️) next to a modified system default variable to reset the variable to its default setting.
 - Click **Add** next to a variable type to create a new variable of that type. See [Creating New Variables](#) on page 703 for more information.
5. Specify a new value for the variable in the field next to the variable you want to modify.

The value for the variable is modified.

Note that modifying the value for a system default variable from within an intrusion policy makes the new value available to all intrusion policies.

See the following for more information:

- If you are modifying an IP-based variable, use the syntax described in [Defining IP Addresses in Variables and Rules](#) on page 713.
- If you are modifying a port-based variable, use the syntax described in [Defining Ports in Variables and Rules](#) on page 715.
- If you are modifying a custom variable, use the syntax described in the [Custom Variables table](#) on page 711.

TIP! For a description of predefined system default variables, see the [Variable Descriptions table](#) on page 699.

6. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Creating New Variables

LICENSE: Protection

If you create your own Sourcefire 3D System rules, you can also create the following types of variables to use in your rules:

- custom system default variables which, like predefined system default variables, are local to the Sourcefire managed device or Defense Center where you create your intrusion policy

Note that the system exports custom system default variables and their values when you export an intrusion policy. See [Importing and Exporting Objects](#) on page 2225 for more information.

- policy-specific variables, which are specific to the policy in which you create them

Note that only policy-specific variables that you create are added to the policy where you create them. However, the intrusion policy also lists and provides access to system default variables, which reside on the appliance where you create the policy. You can create and modify both types of variables.

In the case of policy-specific variables, you can also use an existing system variable policy-specific value for the same variable. As shown in the following figure, a policy-specific value takes precedence over the value for the system default variable.

Policy-Specific Variable

Name: \$EXAMPLE_VARIABLE
Type: IP
Value: 192.168.1.1

Takes precedence over

System Default Variable

Name: \$EXAMPLE_VARIABLE
Type: IP
Value: any

Policy-specific and system default variables use the same syntax and must follow the same guidelines. You can create nested variables so long as the nesting is not circular.

TIP! Note that non-negated lists defined in the following examples are surrounded with brackets, which was required in previous software versions but is no longer required. However, you must surround negated lists in brackets.

In the following example, `!$SMTP_SERVERS`, `!$HTTP_SERVERS`, and `$OTHER_SERVERS` are valid nested variables; that is, they are variables that are included in other variables:

1. Define `SMTP_SERVERS` as `10.1.1.1`
2. Define `HTTP_SERVERS` as `10.1.1.2`.
3. Create a new variable called `OTHER_SERVERS` and define it as `10.2.2.0/24`.
4. Define `HOME_NET` as
`[10.1.1.0/24, !$SMTP_SERVERS, !$HTTP_SERVERS, $OTHER_SERVERS]`.

In the following example, however, `HOME_NET` is an invalid nested variable because the nesting of `HOME_NET` is circular; that is, the definition of `OTHER_SERVERS` includes `HOME_NET`, so you would be nesting `HOME_NET` in itself:

1. Define `SMTP_SERVERS` as `10.1.1.1`
2. Define `HTTP_SERVERS` as `10.1.1.2`.
3. Create a new variable called `OTHER_SERVERS` and define it as
`[10.2.2.0/24, $HOME_NET]`.
4. Define `HOME_NET` as
`[10.1.1.0/24, !$SMTP_SERVERS, !$HTTP_SERVERS, $OTHER_SERVERS]`.

Note that nested, negated variables are not supported. For example, if you want to create a variable, `NONCORE_NET`, that represents the IP addresses that are outside of your protected networks, you could not create it using the following procedure:

1. Define `HOME_NET` as `[10.1.0.0/16, 10.2.0.0/16, 10.3.0.0/16]`.
2. Define `EXTERNAL_NET` as `!$HOME_NET`.
3. Create a new variable called `DMZ_NET` and define it as `10.4.0.0/16`.
4. Create a new variable called `NOTDMZ_NET` and define it as `!$DMZ_NET`.
5. Create the `NONCORE_NET` variable and define it as
`[$EXTERNAL_NET, $NOTDMZ_NET]`.

Instead, create `NONCORE_NET` as follows:

1. Define `HOME_NET` as `[10.1.0.0/16, 10.2.0.0/16, 10.3.0.0/16]`.
2. Create a new variable called `DMZ_NET` and define it as `10.4.0.0/16`.
3. Create a new variable called `NONCORE_NET` and define it as
`![$HOME_NET, $DMZ_NET]`.

See the following sections for more information:

- [Creating New Policy-Specific Variables](#) on page 705
- [Creating New System Default Variables](#) on page 707

Creating New Policy-Specific Variables

LICENSE: Protection

A policy-specific variable sets an explicit value for the variable in the intrusion policy where you create it. You can add a new policy-specific variable from the Policy-Specific Variables section on the Variables page. You can also set a policy-specific value using an existing system default variable.

When you add a new policy-specific variable to an intrusion policy, the system also creates a system default variable with the same name and a value of **any** and lists the new system default variable in all intrusion policies on the system.

IMPORTANT! Creating a new policy-specific variable also creates a system default variable with the same name and a value of **any**. The new system default variable is available to all other intrusion policies. However, to use the policy-specific variable in another policy, you must add the policy-specific variable to the other policy and give it an explicit value. For example, suppose you create a new policy-specific variable, set the value to **80,8080**, and use it in a rule enabled in one of your custom policies. If you then decide to use the variable in a rule in another policy, you have to add a policy-specific variable with the same name to the second policy and set the value to **80,8080**. In any intrusion policy where you do not explicitly set a policy-specific value for the variable, the system will use the system default value of **any**.

To create a new policy-specific variable:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

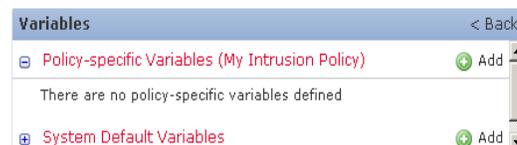
2. Click the edit icon () next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

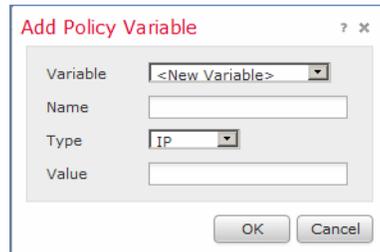
The Policy Information page appears.

3. Click **Manage Variables** on the Policy Information page.

The Variables page appears.



- In the **Policy-Specific Variables** section, click **Add**.
The Add Policy Variable pop-up window appears.



- You have two choices:
 - To create a new policy-specific variable, ensure that **<New Variable>** is selected from the Variable drop-down list and enter a **Name** for the new variable.
For example, if you are creating a variable that defines all Apache web servers on your network, you might decide to name the variable **APACHE_SERVERS**.

TIP! The **Name** field accepts only alphanumeric characters and underscores. Variable names **cannot** contain spaces or punctuation characters. In addition, the system automatically precedes each variable name with a dollar sign (\$), and you receive an error if you use \$ in a variable name.

- To set a policy-specific value using an existing system default variable, select the name of the existing variable from the Variable drop-down list.
The **Type** field is constrained. Go to step 7.
- If you created a new variable in step 5, select **IP**, **Port**, or **Custom** from the **Type** drop-down list:
 - See [Defining IP Addresses in Variables and Rules](#) on page 713 for more information if you are defining a IP address-based variable.
 - See [Defining Ports in Variables and Rules](#) on page 715 for more information if you are defining a port-based variable.
 - See [Understanding Custom Variables](#) on page 710 if you are defining a special-purpose custom variable with one of the reserved variable names described in the [Custom Variables table](#) on page 711.
When the new variable is a custom variable, the **Value** field expands.

7. Enter the new variable **Value**:
 - If the value describes IP addresses, see [Defining IP Addresses in Variables and Rules](#) on page 713 for information about the syntax you can use.
 - If the value describes port numbers, see [Defining Ports in Variables and Rules](#) on page 715 for information about the syntax you can use.
 - If the value is another variable or negation of a variable, enter it in the value field (for example, `$HOME_NET` or `!$HOME_NET`).
 - If the value is for a custom variable, use the syntax described in the [Custom Variables table](#) on page 711.

IMPORTANT! A variable value can include up to 8192 characters. Keep in mind, however, that this limit applies to the size of the expanded value of the variable. If you use one or more variables to define another variable, the total number of characters and spaces of all the variable values cannot exceed 8192 characters.

8. Click **OK**.

The new variable is created. If a corresponding system default variable did not exist, a system default variable with a value of **any** is added in the policy and appears in all future intrusion policies that you create.

Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Creating New System Default Variables

LICENSE: Protection

Each intrusion policy lists the existing, pre-configured system default variables described in the [Variable Descriptions table](#) on page 699. You can also add new system default variables from the System Default Variables section on the Variables page.

System default variables that you add from within an intrusion policy are listed in the policy and are available for use by all intrusion policies created on the Sourcefire Defense Center where you added the variables. When you enable an intrusion rule that uses a particular variable, the system uses the system default value for the variable unless it is overridden in the policy by a policy-specific value.

TIP! An exported policy includes system default variables that you create when you export an intrusion policy that uses them.

To create a new system default variable:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

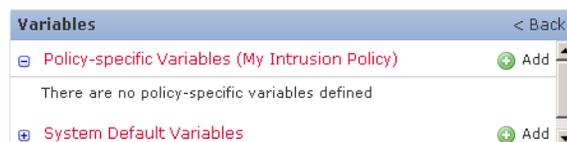
2. Click the edit icon () next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

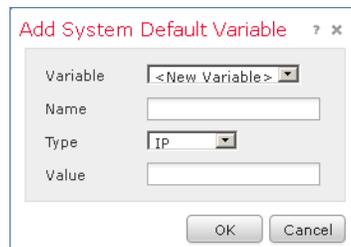
3. Click **Manage Variables** on the Policy Information page.

The Variables page appears.



4. In the **System Default Variables** section, click **Add**.

The Add System Default Variable pop-up window appears.



5. To add a new system default variable, ensure that **<New Variable>** is selected from the **Variable** drop-down list and enter a **Name** for the new variable.

For example, if you are creating a variable that defines all Apache web servers on your network, you might decide to name the variable `APACHE_SERVERS`.

TIP! The **Name** field accepts only alphanumeric characters and underscores. Variable names **cannot** contain spaces or punctuation characters. In addition, the system automatically precedes each variable name with a dollar sign (\$), and you receive an error if you use \$ in a variable name.

Note that you cannot use an existing policy-specific variable to set a value for a system default variable, because the system automatically creates a new system default variable with a value of `any` when you create a new policy-specific variable. See [Modifying Variables](#) on page 701 for information on modifying the value of a system default variable.

6. Select **IP**, **Port**, or **Custom** from the **Type** drop-down list:
 - See [Defining IP Addresses in Variables and Rules](#) on page 713 for more information if you are defining an IP address-based variable.
 - See [Defining Ports in Variables and Rules](#) on page 715 for more information if you are defining a port-based variable.
 - See [Understanding Custom Variables](#) on page 710 if you are defining a special-purpose custom variable with one of the reserved variable names described in the [Custom Variables table](#) on page 711.
When the new variable is a custom variable, the **Value** field expands.
7. Enter the new variable **Value**:
 - If the new variable describes IP addresses, see [Defining IP Addresses in Variables and Rules](#) on page 713 for information about the syntax you can use.
 - If the new variable describes port numbers, see [Defining IP Addresses in Variables and Rules](#) on page 713 for information about the syntax you can use.
 - If the value is another variable or negation of a variable, enter it in the value field (for example, `$HOME_NET` or `!$HOME_NET`).
 - If the new variable is a custom variable, use the syntax described in the [Custom Variables table](#) on page 711.

IMPORTANT! A variable value can include up to 8192 characters. Keep in mind, however, that this limit applies to the size of the expanded value of the variable. If you use one or more variables to define another variable, the total number of characters and spaces of all the variable values cannot exceed 8192 characters.

8. Click **OK**.
The new system default variable is available for use by, and listed in, all intrusion policies.
9. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Deleting Unused Variables

LICENSE: Protection

You can delete policy-specific variables from an intrusion policy. You cannot delete a variable that is used in any rule, regardless of whether the rule is set to Generate Events, Drop and Generate Events, or Disable. If you delete variables that are in use, an error message appears when you attempt to save the changes in your policy.

You can delete system default variables that you create, but you cannot delete system default variables with predefined default values.

To delete a variable:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Manage Variables** on the Policy Information page.
The Variables page appears.
4. Optionally, click the + sign to the left of the **Policy-Specific Variables** or **System Default Variables** section name to display the corresponding variables list.
5. Click **Delete** next to the variable that you want to delete.
The variable is deleted.
6. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Understanding Custom Variables

LICENSE: Protection

Custom variables allow you to configure features that you cannot otherwise configure via the web interface. You create a custom variable using a specific assigned name and a set of instructions appropriate to the function the variable provides. The system recognizes the name and implements the instructions accordingly. You cannot name another variable with a name reserved for a custom variable.

You can create a custom variable as a system default or policy-specific variable. See [Creating New Variables](#) on page 703 for more information.

IMPORTANT! In other variables, **any** might mean, for example, any port or any IP address, but **any** disables a custom variable.

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 in a custom variable 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. You can also modify a custom variable and remove or add instructions as needed. See the Value for Variable Binding column in [Custom Variables](#) table for the syntax to use when defining each type of custom variable.

The [Custom Variables](#) table describes the custom variables that are available.

Custom Variables

VARIABLE	DESCRIPTION	VALUE FOR VARIABLE BINDING
SNORT_BPF	<p>Allows you to apply a Berkeley Packet Filter (BPF) to filter traffic before it reaches the system.</p> <p>Note that with or without BPF, the system filters traffic internally so that only traffic covered by intrusion rules and preprocessors is processed.</p>	Use the same standard BPF instruction set as accepted by tcpdump.
USER_CONF	<p>Provides a general tool that allows you to configure one or more features not otherwise available via the web interface.</p> <p>WARNING! Do not use the reserved variable <code>USER_CONF</code> to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Sourcefire Support. Conflicting or duplicate configurations will halt the system.</p>	Use the feature-specific directives provided by Sourcefire Support or in the feature description.

Configuring Custom Variables

You can set an explicit value for the predefined system default SNORT_BPF custom variable by typing a new value in the field next to SNORT_BPF in the System Default Variables section of the intrusion policy Variables page. See [Modifying Variables](#) on page 701 for more information. You can also create a policy-specific variable for SNORT_BPF based on the existing system default variable. See [Modifying Variables](#) on page 701 for more information.

You can create a new USER_CONF variable as a policy-specific variable using the reserved name USER_CONF. See [Creating New Policy-Specific Variables](#) on page 705. You can also create a system default variable using the reserved name USER_CONF. See [Creating New Policy-Specific Variables](#) on page 705.

Defining IP Addresses and Ports for Your Network

LICENSE: Protection

When tuning your system, you can define IP addresses and ports in variables and rules. This lets you specify the level of granularity of inspection so that rules execute against the IP addresses and ports appropriate to your network needs.

Within the following existing, pre-configured variables and also within variables you create, you can define specific IP addresses:

- your home network
- the external network
- AOL Instant Messenger (AIM) servers
- domain name service (DNS) servers
- HTTP servers (or web servers)
- SIP servers
- SQL database servers
- SMTP servers
- SNMP servers
- SSH Servers
- Telnet servers

You can define specific port numbers within variables you create or for the following existing, pre-configured variables:

- ports susceptible to shell code exploits
- HTTP (or web server) ports
- Oracle database server ports
- ports to inspect for downloaded file type in HTTP, IMAP, and POP traffic pointed to by the `file_data` keyword
- non-encrypted ports that detect files in a network stream
- GTP ports
- SIP ports
- SSH server ports

For details on defining variables, see [Managing Variables](#) on page 698. For details on existing variables, see [Understanding Existing Variables](#) on page 699.

Within standard text rules, you can specify the source and destination IP addresses and ports the rule uses to determine whether it executes against a packet. For more information on defining IP addresses and port values in rule headers, see [Understanding Rule Headers](#) on page 1049.

You can set IP addresses and definitions in a variety of ways. See the following sections for more information:

- [Defining IP Addresses in Variables and Rules](#) on page 713 describes the syntax you can use to define single IP addresses, IP address ranges and lists, and any IP address in variables and rules.
- [Defining Ports in Variables and Rules](#) on page 715 describes the syntax you can use to define single and multiple ports in variables and rules.

Defining IP Addresses in Variables and Rules

LICENSE: Protection

When configuring variables and 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, 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.

See the following sections for more information on how to specify IP addresses:

- [Specifying Any IP Address](#) on page 713
- [Specifying Multiple IP Addresses](#) on page 713
- [Specifying Network Objects](#) on page 714
- [Excluding IP Addresses](#) on page 714

Specifying Any IP Address

LICENSE: Protection

You can specify the word **any** as a variable value or 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 IP** and **Destination IP** fields and evaluates packets with any IPv4 or IPv6 source or destination address:

```
a!ert 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.

IMPORTANT! You must surround negated lists with brackets. See [Excluding IP Addresses](#) on page 714 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](#) on page 57.
- 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](#) on page 161 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 in variables or rules 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](#) on page 714 for more information.

Excluding IP Addresses

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.

IMPORTANT! 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 Variables and Rules

LICENSE: Protection

The Sourcefire 3D System uses a specific type of syntax to define the port numbers used in variables and rule headers.

IMPORTANT! The system ignores port definitions in an intrusion rule header when the protocol is set to ip. For more information, see [Specifying Protocols](#) on page 1051.

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 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. See [Specifying Source and Destination Ports](#) on page 1053 for more information.

The following table describes the syntax you can use:

Port Syntax for Variables and Rules

SYNTAX	DESCRIPTION	EXAMPLE
any	Specifies any port.	any
port	Specifies a single port.	80
first_port-last_port	Indicates a port range, where <code>first_port</code> specifies the first port, and <code>last_port</code> specifies the last port.	20-21
-port	Indicates all ports less than or equal to port.	-21
port-	Indicates all ports greater than or equal to port.	80-
!	Indicates all ports except the specified port, port list, or range of ports. Note that you can logically use negation with all port designations except <i>any</i> , which if negated would indicate <i>no port</i> .	!80

Note that 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.

CHAPTER 20

MANAGING RULES IN AN INTRUSION POLICY

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](#) on page 689 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.

You can generate rule state recommendations based on vulnerabilities associated with the hosts and applications on your network and, optionally, update rules to reflect the recommended states.

See the following sections for more information:

- [Understanding Intrusion Prevention Rule Types](#) on page 718 describes the intrusion rules and preprocessor rules you can view and configure in an intrusion policy.
- [Viewing Rules in an Intrusion Policy](#) on page 719 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](#) on page 729 describes how you can use rule filters to find the rules for which you want to apply rule settings.

- [Setting Rule States](#) on page 743 describes how to enable and disable rules from the Rules page.
- [Filtering Intrusion Event Notification Per Policy](#) on page 746 explains how to set event filtering thresholds for specific rules and set suppression on specific rules.
- [Adding Dynamic Rule States](#) on page 756 explains how to set rule states that trigger dynamically when rate anomalies are detected in matching traffic.
- [Adding Alerts](#) on page 761 describes how to associate SNMP alerts with specific rules.
- [Automatically Enabling Advanced Settings](#) on page 787 explains how to enable preprocessors and other advanced features required by rules when those rules are set to Generate Events or Drop and Generate Events.
- [Adding Rule Comments](#) on page 762 describes how to add comments to rules in an intrusion policy.
- [Managing FireSIGHT Rule State Recommendations](#) on page 764 describes how to generate rule state recommendations based on vulnerabilities associated with the hosts and applications on your network.
- [Using Layers in an Intrusion Policy](#) on page 792 explains how you can more efficiently manage multiple intrusion policies in a complex network by adding intrusion policy layers comprised of individual configurations for rule attributes and advanced settings.

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 Sourcefire 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 Sourcefire's shared object rules, standard text rules, and preprocessor rules for each default intrusion policy included with the system.

The [Rule Types](#) table describes each type of rule included with the Sourcefire 3D System.

Rule Types

TYPE	DESCRIPTION
shared object rule	An intrusion rule created by the Sourcefire 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 on page 1184 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 on page 1184, Understanding and Writing Intrusion Rules on page 1046 and Importing Local Rule Files on page 2080 for more information. A standard text rule created by the VRT has a GID (generator ID) of 1. Custom standard text rule 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 Sourcefire 3D System. You must enable preprocessor rules if you want them to generate events. These rules have a decoder- or preprocessor-specific GID (generator ID). See the Generator IDs table on page 785 for more information.

Viewing Rules in an Intrusion Policy

LICENSE: Protection

You can adjust how rules are displayed in the intrusion policy. Rules can be sorted 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](#) on page 729
- the rule attribute menus — for more information, see [Setting Rule States](#) on page 743, [Filtering Intrusion Event Notification Per Policy](#) on page 746, [Adding Dynamic Rule States](#) on page 756, [Adding Alerts](#) on page 761, and [Adding Rule Comments](#) on page 762
- the rules listing — for more information, see the [Rules Page Columns table](#) on page 720
- the rule details — for more information, see [Viewing Rule Details](#) on page 723

You can also sort rules by different criteria; for more information, see [Sorting the Rule Display](#) on page 722.

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 [Rules Page Columns](#) table that follows describes the columns on the Rules page.

Rules Page Columns

HEADING	DESCRIPTION	FOR MORE INFORMATION, SEE...
GID	Integer which indicates the Generator ID (GID) for the rule.	Reading Preprocessor Generator IDs on page 784
SID	Integer which indicates the Snort ID (SID), which acts a unique identifier for the rule.	Reading Preprocessor Generator IDs on page 784
Message	Message included in events generated by this rule, which also acts as the name of the rule.	Defining the Event Message on page 1058

Rules Page Columns (Continued)

HEADING	DESCRIPTION	FOR MORE INFORMATION, SEE...
	<p>The rule state for the rule, which may be one of four states:</p> <ul style="list-style-type: none"> • drop and generate events () • generate events () • disable () • inherit (blank) <p>Note that you can access the Set rule state dialog box for a rule by clicking on its rule state icon.</p>	Setting Rule States on page 743
	FireSIGHT recommended rule state for the rule.	Managing FireSIGHT Rule State Recommendations on page 764
	Event filter, including event thresholds and event suppression, applied to the rule.	Filtering Intrusion Event Notification Per Policy on page 746
	Dynamic rule state for the rule, which goes into effect if specified rate anomalies occur.	Adding Dynamic Rule States on page 756
	Alerts configured for the rule, including SNMP alerts.	Adding Alerts on page 761
	Comments added to the rule.	Adding Rule Comments on page 762

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 either of these views is the same as making the changes in the other. See [Using Layers in an Intrusion Policy](#) on page 792 for more information. The drop-down list also lists the Rules page for the read-only base policy. See [Understanding the Base Policy](#) on page 691 for information on the base policy.

To view the rules in an intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

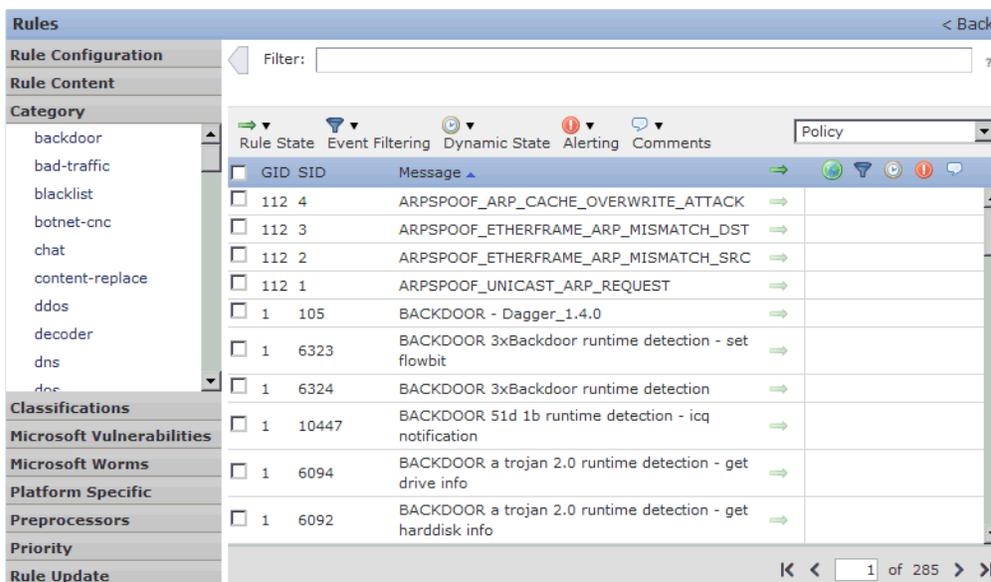
2. Click the edit icon (✎) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Manage 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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Manage Rules** in the Policy Information page.
The Rules page appears. By default, the page lists the rules alphabetically by message.
4. Click the title or icon in the top of the column by which you want to sort.
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, FireSIGHT recommendations, 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.

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 Event Information on page 625
Rule State	The current rule state for the rule. Also indicates the layer where the rule state is set.	Setting Rule States on page 743; Using Layers in an Intrusion Policy on page 792
FireSIGHT Recommendation	If FireSIGHT recommendations have been generated, the recommended rule state for the rule.	Managing FireSIGHT Rule State Recommendations on page 764
Rule Overhead	The rule's potential impact on system performance and the likelihood that the rule might generate false positives.	Understanding Rule Overhead on page 767
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 on page 725
Suppressions	Suppression settings currently set for this rule, as well as the facility to add suppressions for the rule.	Setting Suppression for a Rule on page 726
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 on page 727

Rule Details (Continued)

ITEM	DESCRIPTION	FOR MORE INFORMATION, SEE...
Alerts	Alerts currently set for this rule, as well as the facility to add an alert for the rule.	Setting an SNMP Alert for a Rule on page 728
Comments	Comments added to this rule, as well as the facility to add comments for the rule.	Adding a Rule Comment for a Rule on page 729
Documentation	The rule documentation for the current rule, supplied by the Sourcefire Vulnerability Research Team (VRT).	Using Packet View Actions on page 629

To view rule details:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Manage Rules** in the Policy Information page.
The Rules page appears. By default, the page lists the rules alphabetically by message.
4. Highlight the rule whose rule details you want to view.
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](#) on page 747.

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:

ACCESS: Admin/Intrusion Admin

1. Click **Add** next to Thresholds.
The Set Threshold dialog box appears.
2. 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.
3. Select the appropriate option for **Track By** to indicate whether you want the event instances tracked by source or destination IP address.
4. In the **Count** field, type the number of event instances you want to use as your threshold.
5. In the **Seconds** field, type a number between 1 and 86400 that specifies the time period for which event instances are tracked.
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](#) on page 753.

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:

ACCESS: Admin/Intrusion Admin

1. Click **Add** next to Suppressions.
The Add Suppression dialog box appears.

2. 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.
3. If you selected **Source** or **Destination** for the suppression type, in the **Network** field enter the IP address, an address block, a 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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](#) on page 757.

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:

ACCESS: Admin/Intrusion Admin

1. Click **Add** next to Dynamic State.
The Add Rate-Based Rule State dialog box appears.
2. Select the appropriate **Track By** 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.

3. Optionally, when 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
4. 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.
5. Select a **New State** radio button to 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 to generate an event in passive deployments.
 - Select **Disabled** to take no action.
6. In the **Timeout** field, using an integer between 1 and 2147483647, 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.
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 left blank, you will receive an error message indicating which fields must be filled.

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 Alerts](#) on page 761.

To add an SNMP alert from the rule details:

ACCESS: Admin/Intrusion Admin

- ▶ 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](#) on page 762.

To add a comment from the rule details:

ACCESS: Admin/Intrusion Admin

1. Click **Add** next to Comments.
The Add Comment dialog box appears.
2. Type the rule comment.
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.

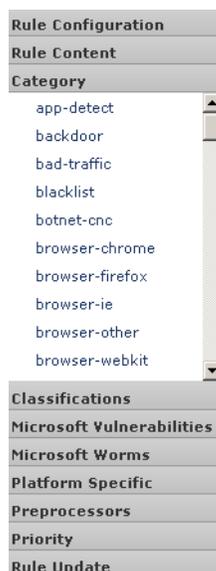
The filter you construct is shown in the Filter text box.

Filter:

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](#) on page 730
- [Setting a Rule Filter in an Intrusion Policy](#) on page 741

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](#) on page 731
- [Understanding Rule Configuration Filters](#) on page 734
- [Understanding Rule Content Filters](#) on page 737
- [Understanding Rule Categories](#) on page 739
- [Editing a Rule Filter Directly](#) on page 739

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 of the rule filters have multiple levels that you 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.

For example, if you click **Drop and Generate Events** under Rule Configuration > Recommendation in the filter panel, **Recommendation:"Drop and Generate Events"** is added to the filter text box. If you then click **Generate Events** under Rule Configuration > Recommendation, the filter changes to **Recommendation:"Generate Events"**.

- 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 hit 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.

IMPORTANT! The Sourcefire 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 [Rule Filter Groups](#) table describes the different rule filters.

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 on page 734.	No	A grouping	keywords
Rule Content	Finds rules according to the content of the rule. See Understanding Rule Content Filters on page 737.	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 on page 739.	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 Searching for Intrusion Events on page 644 and Defining the Intrusion Event Classification on page 1058.	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.

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 Understanding Preprocessors on page 780 and Setting Rule States on page 743 for more information.	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. For example, if you want to view the set of rules whose rule state does not match the recommended rule state, you can filter on rule state by selecting **Does not match recommendation**.

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 **Drop and Generate Events** under Rule Configuration > Recommendation in the filter panel, **Recommendation:"Drop and Generate Events"** is added to the filter text box. If you then click **Generate Events** under Rule

Configuration > Recommendation, the filter changes to **Recommendation:"Generate Events"**.

See the following procedures for more information on the rule configuration settings you can use to filter.

To use the Rule State filter:

ACCESS: Admin/Intrusion Admin

1. Under Rule Configuration, click **Rule State**.
2. Select the rule state to filter by:
 - To find rules that only generate events, select **Generate Events**, and click **OK**.
 - To find rules that are set to generate events and drop the matching packet, select **Drop and Generate Events**, and click **OK**.
 - To find disabled rules, select **Disabled**, and click **OK**.
 - To find rules whose rule state does not match the recommended state, select **Does not match recommendation**, and click **OK**.

The Rules page updates to display rules according to current rule state.

To use the Recommendation filter:

ACCESS: Admin/Intrusion Admin

1. Under **Rule Configuration**, click **Recommendation**.
2. Select the FireSIGHT rule state recommendation to filter by.
The Rules page updates to display rules according to recommended rule state.

To use the Threshold filter:

ACCESS: Admin/Intrusion Admin

1. Under Rule Configuration, click **Threshold**.
2. Select the threshold setting to filter by:
 - To find rules with a threshold type of **Limit**, select **Limit**, and click **OK**.
 - To find rules with a threshold type of **threshold**, select **Threshold**, and click **OK**.
 - To find rules with a threshold type of **both**, select **Both**, and click **OK**.
 - To find rules with thresholds tracked by **source**, select **Source**, and click **OK**.
 - To find rules with thresholds tracked by destination, select **Destination**, and click **OK**.
 - To find any rule with a threshold set, select **All**, and 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:

ACCESS: Admin/Intrusion Admin

1. Under **Rule Configuration**, click **Suppression**.
2. Select the suppression setting to filter by:
 - To find rules where events are suppressed for packets inspected by that rule, select **Rule**, and click **OK**.
 - To find rules where events are suppressed based on the source of the traffic, select **Source**, and click **OK**.
 - To find rules where events are suppressed based on the destination of the traffic, select **Destination**, and click **OK**.
 - To find any rule with suppression set, select **All**, and 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:

ACCESS: Admin/Intrusion Admin

1. Under **Rule Configuration**, click **Dynamic State**.
2. Select the suppression setting to filter by:
 - To find rules where a dynamic state is configured for packets inspected by that rule, select **Rule**, and click **OK**.
 - To find rules where a dynamic state is configured for packets based on the source of the traffic, select **Source**, and click **OK**.
 - To find rules where a dynamic state is configured based on the destination of the traffic, select **Destination**, and click **OK**.
 - To find rules where a dynamic state of **Generate Events** is configured, select **Generate Events**, and click **OK**.
 - To find rules where a dynamic state of **Drop and Generate Events** is configured, select **Drop and Generate Events**, and click **OK**.
 - To find where a dynamic state of **Disabled** is configured, select **Disabled**, and click **OK**.
 - To find any rule with suppression set, select **All**, and 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 Comment filter:

ACCESS: Admin/Intrusion Admin

1. Under **Rule Configuration**, click **Comment**.
2. Type the string of comment text to filter by.
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 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 [Rule Content Filters](#) table.

Rule Content Filters

TO USE THIS FILTER, CLICK...	THEN...	RESULT
Message	Type the message string to filter by, and click OK .	Finds rules that contain the supplied string in the message field.
SID	Type the SID number to filter by, and click OK .	Finds rules that have the specified SID.
GID	Type the GID number to filter by, and click OK .	Finds rules that have the specified GID.
Reference	Type the reference string to filter by, and click OK .	Finds rules that contain the supplied string in the reference field.

Rule Content Filters (Continued)

TO USE THIS FILTER, CLICK...	THEN...	RESULT
Action	Select the action to filter by: <ul style="list-style-type: none"> • To find alert rules, select Alert, and click OK. • To find pass rules, select Pass, and click OK. 	Finds rules that start with alert or pass .
Protocol	Select the protocol to filter by.	Finds rules that include the selected protocol.
Direction	Select a directional setting to filter by: <ul style="list-style-type: none"> • To find rules that inspect traffic moving in a specific direction, select Directional, and click OK. • To find rules that inspect traffic moving in either direction between a source and destination, select Bidirectional, and click OK. 	Finds rules based on whether the rule includes the indicated directional setting.
Source IP	Type the source IP address to filter by. 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. 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.
Source port	Type the source port to filter by. The port value must be an integer between 1 and 65535 or a port variable.	Finds rules that include the specified source port.

Rule Content Filters (Continued)

TO USE THIS FILTER, CLICK...	THEN...	RESULT
Destination port	Type the destination port to filter by. 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.	Finds rules with the selected rule overhead.
Metadata	Type the metadata key-value pair to filter by, separated by a space. For example, type <code>metadata:"service http"</code> 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 Sourcefire 3D System 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 might filter by the **os-linux** category and then disable all the rules showing to disable the entire **os-linux** category.

You can hover your pointer over a category name to display the number of rules in the category.

IMPORTANT! The Sourcefire 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](#) on page 734, [Understanding Rule Content Filters](#) on page 737, and [Understanding Rule Categories](#) on page 739. 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 the [Rule Types table](#) on page 719 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](#) on page 1058. For information on rule SIDs and GIDs, see [Reading Preprocessor Generator IDs](#) on page 784. 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, including selecting any of the features available in the context menu. 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 pre-defined 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](#) on page 730.

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.

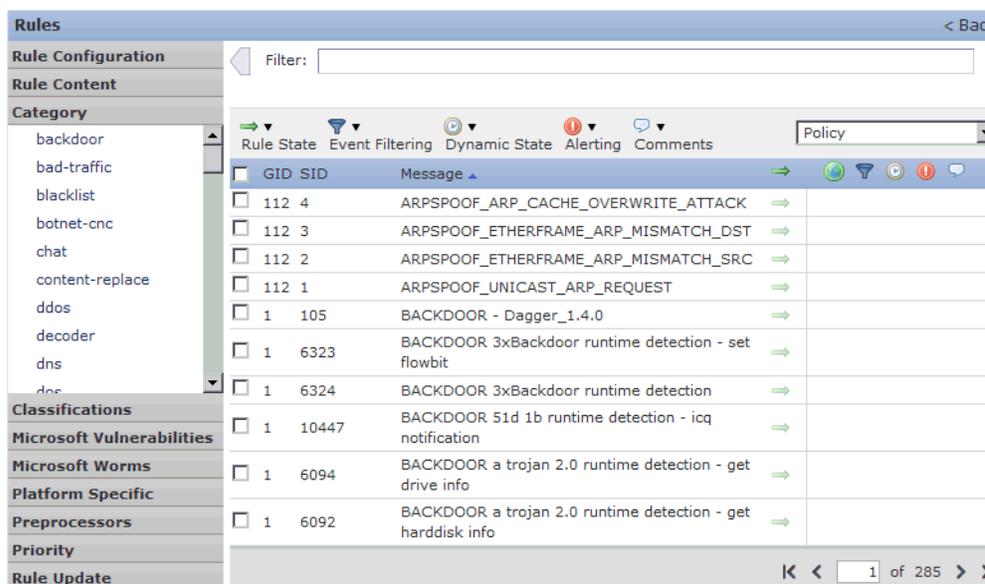
To filter for specific rules in an intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Manage Rules** in the Policy Information page.
The Rules page appears. By default, the page lists the rules alphabetically by message.



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](#) on page 731.
 - [Understanding Rule Configuration Filters](#) on page 734
 - [Understanding Rule Content Filters](#) on page 737
 - [Understanding Rule Categories](#) on page 739
 - [Editing a Rule Filter Directly](#) on page 739 for more information.The page refreshes to display all matching rules, and the number of rules matching the filter is displayed above the filter text box.
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.

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](#) on page 743 for information on enabling and disabling rules on the Rules page.
 - See [Filtering Intrusion Event Notification Per Policy](#) on page 746 for information on adding thresholding and suppression to rules.
 - See [Adding Dynamic Rule States](#) on page 756 for information on setting dynamic rule states that trigger when rate anomalies occur in matching traffic.
 - See [Adding Alerts](#) on page 761 for information on adding SNMP alerts to specific rules.
 - See [Adding Rule Comments](#) on page 762 for more information on adding rule comments to rules.
7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Setting Rule States

LICENSE: Protection

The Sourcefire 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, including when a 3D9900 or Series 3 device inline interface set is in tap mode. Setting a rule to Generate Events or 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 settings:

- 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, including when a 3D9900 or Series 3 device inline interface set is in tap mode.

Note that your intrusion policy must be set to drop rules in an inline deployment for the system to drop packets; see [Setting Drop Behavior in an Inline Deployment](#) on page 689 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 to a managed device that uses an inline set.

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](#) on page 729.

See [Understanding and Writing Intrusion Rules](#) on page 1046 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 will not override your change.

To change the rule state for one or more rules:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.

2. Click the edit icon (✎) next to the policy you want to edit.

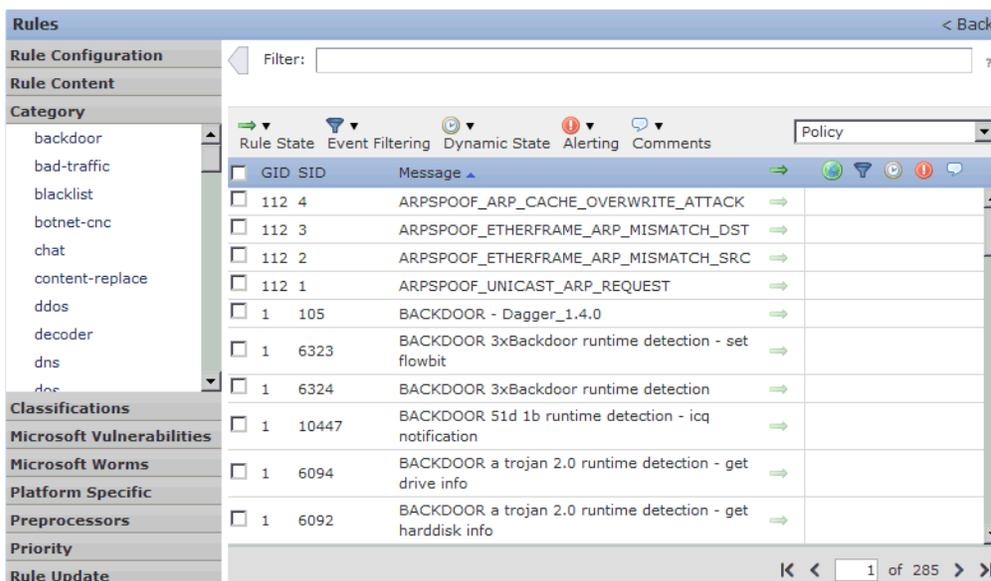
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Note that this page indicates the total number of enabled rules and 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.

3. Click **Manage Rules** on the Policy Information page.

The Rules page appears. By default, the page lists the rules alphabetically by message.



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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.

The page refreshes to display all matching rules.

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.
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**.

IMPORTANT! Sourcefire **strongly** recommends that you **do not** enable all the intrusion rules in an intrusion policy. The performance of your managed device is likely to degrade if all the rules are enabled. Instead, tune your rule set to match your network environment as closely as possible.

7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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 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](#) on page 747 explains how to set thresholds that dictate how often (based on the number of occurrences) an event is displayed. You can configure thresholding per event, per policy.
- [Configuring Suppression Per Intrusion Policy](#) on page 753 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](#) on page 747
- [Adding and Modifying Intrusion Event Thresholds](#) on page 749
- [Viewing and Deleting Intrusion Event Thresholds](#) on page 751
- [Setting a Threshold for a Rule](#) on page 725

Understanding Event Thresholding

LICENSE: Protection

First, you must specify the thresholding type. You can select from the options discussed in the [Thresholding Options](#) table.

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.

Thresholding Options (Continued)

OPTION	DESCRIPTION
Threshold	<p>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.</p>
Both	<p>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:</p> <ul style="list-style-type: none"> • 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, you must specify the tracking, which determines whether the event threshold is calculated per source or destination IP address. Select one of the options from the [Thresholding IP Options](#) table to specify how the system tracks event instances.

Thresholding IP Options

OPTION	DESCRIPTION
Source	<p>Calculates event instance count per source IP address.</p>
Destination	<p>Calculates the event instance count per destination IP address.</p>

Finally, you must specify the number of instances and time period that define the threshold.

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](#) on page 756, [Filtering Events](#) on page 1165, and [Configuring Suppression Per Intrusion Policy](#) on page 753 for more information.

See the following sections for more information:

- [Adding and Modifying Intrusion Event Thresholds](#) on page 749
- [Setting a Threshold for a Rule](#) on page 725
- [Viewing and Deleting Intrusion Event Thresholds](#) on page 751

TIP! You can also add thresholds from within the packet view of an intrusion event. See [Viewing Event Information](#) on page 625 for more information.

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 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](#) on page 751.

You can also modify the global threshold that applies by default to all rules and preprocessor-generated events. For more information, see [Using Global Rule Thresholding](#) on page 1009.

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.

TIP! A global or individual threshold on a managed device with multiple CPUs may result in a higher number of events than expected.

To add or modify event thresholds:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy.**

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Manage Rules in the Policy Information page.**

The Rules page appears. By default, the page lists the rules alphabetically by message.

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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.

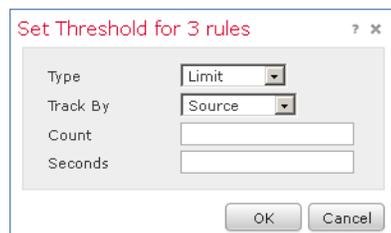
The page refreshes to display all matching rules.

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.

6. Select **Event Filtering > Threshold**.

The thresholding pop-up window appears.



7. 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.

8. Select the appropriate option for **Track By** to indicate whether you want the event instances tracked by source or destination IP address.

9. In the **Count** field, specify the number of event instances you want to use as your threshold.

10. In the **Seconds** field, specify the number of seconds that make up the time period for which event instances are tracked.

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.

12. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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 [Using Global Rule Thresholding](#) on page 1009 for more information.

To view or delete a threshold:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Manage Rules** in the Policy Information page.
The Rules page appears. By default, the page lists the rules alphabetically by message.
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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.
The page refreshes to display all matching rules.
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.
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 and click **Delete** next to the threshold settings. Click **OK** to confirm that you want to delete the configuration.

The page refreshes and the threshold is deleted.

7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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 trigger 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 can 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](#) on page 756, [Filtering Events](#) on page 1165, and [Configuring Event Thresholding](#) on page 747 for more information.

See the following sections for more information:

- [Suppressing Intrusion Events](#) on page 753
- [Viewing and Deleting Suppression Conditions](#) on page 755

TIP! You can also add suppressions from within the packet view of an intrusion event. See [Viewing Event Information](#) on page 625 for more information. You can also access suppression settings by using the right-click context menu on the Rule Editor page and on any intrusion event page (if the event was triggered by an intrusion rule).

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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Manage Rules** in the Policy Information page.
The Rules page appears. By default, the page lists the rules alphabetically by message.

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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.

The page refreshes to display all matching rules.

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 the rules in the current list, select the check box at the top of the column.

6. Select **Event Filtering > Suppression**.

The suppression pop-up window appears.



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.

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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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.

10. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Viewing and Deleting Suppression Conditions

LICENSE: Protection

You may want to view or delete an existing suppression condition. For example, you might 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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Manage Rules** in the Policy Information page.

The Rules page appears. By default, the page lists the rules alphabetically by message.

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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.The page refreshes to display all matching rules.
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 the rules in the current list, select the check box at the top of the column.
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. Click **OK** to confirm that you want to delete the configuration.The page refreshes and the suppression is deleted.
7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 757
- [Setting a Dynamic Rule State](#) on page 758

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 managed devices deployed inline to block rate-based attacks for a specified time and 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 that initially configured for the rule.

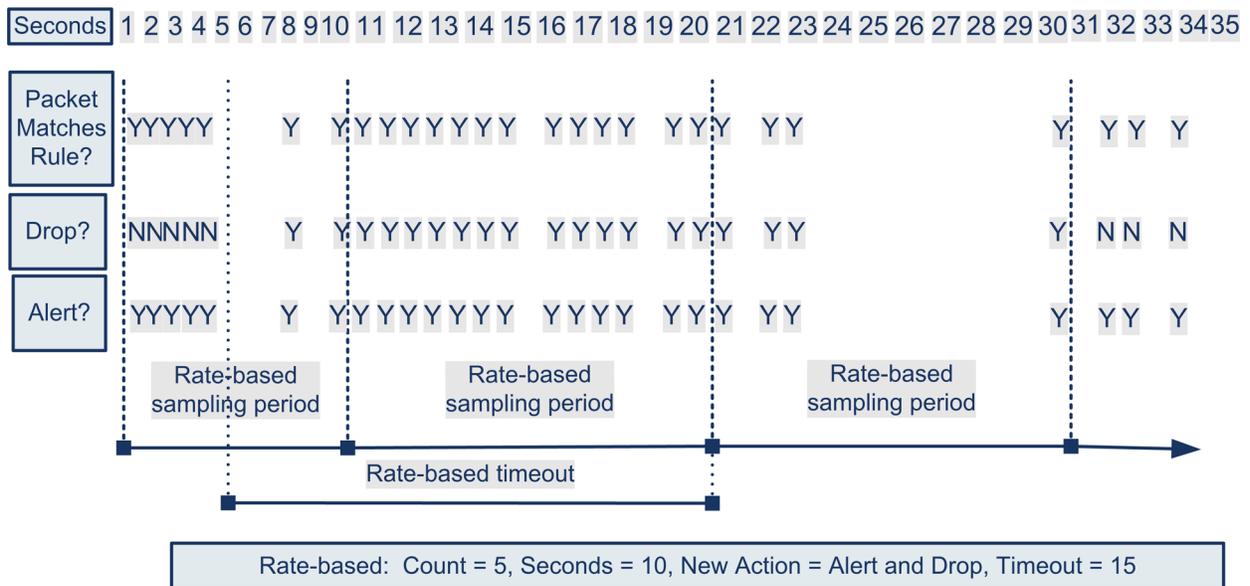
You can configure rate-based attack prevention in an inline deployment to block attacks, either temporarily or permanently. Without rate-based configuration, rules set to Generate Events do generate events, but the system does not drop packets for those rules. However, if the attack traffic matches rules 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.

IMPORTANT! 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 lets 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 those 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 on 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.

IMPORTANT! Dynamic rule states cannot enable disabled rules or drop traffic that matches disabled rules.

To add a dynamic rule state:

ACCESS: Admin/Intrusion Admin

1. Select Policies > Intrusion > Intrusion Policy.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click Rules in the Policy Information page.

The Rules page appears.

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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.

The page refreshes to display all matching rules.

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.

6. Select **Dynamic State > Add Rate-Based Rule State**.

The Add Rate-Based Rule State dialog box appears.



7. Select the appropriate **Track By** 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.

8. When 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

9. 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.

10. Select a **New State** radio button to 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.

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.

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 left blank, you will receive an error message indicating which fields must be filled.

TIP! To delete all dynamic rule settings for a set of rules, select the rules on the Rules page, 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**, and clicking **Delete** by the rate-based filter you want to remove.

13. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Adding Alerts

LICENSE: Protection

If you configure SNMP alerting for your Sourcefire 3D System, you can add an alert to specific rules in your intrusion policy. For more information, see [Adding SNMP Alerts](#) on page 761.

Adding SNMP Alerts

LICENSE: Protection

If you configure an SNMP alert for your Sourcefire 3D System, you can configure rules within an intrusion policy to use that alert when traffic matches the rule and an event is generated.

To set an SNMP alert:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Manage Rules** in the Policy Information page.
The Rules page appears.
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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.
The page refreshes to display all matching rules.
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.

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**.

7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 1184.

To add a comment to a rule:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Manage Rules** in the Policy Information page.

The Rules page appears.

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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.

The page refreshes to display all matching rules.

5. Select the rule or rules where you want to add a comment to a rule:

- 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.

6. Select **Comments > Add Rule Comment**.

The Add Comment dialog box appears.



7. Type the rule comment.

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**. 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.

9. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Managing FireSIGHT Rule State Recommendations

LICENSE: FireSIGHT + Protection

You can use the FireSIGHT Recommended Rules feature to associate the operating systems, servers, and client application protocols detected on your network (see [Introduction to Network Discovery](#) on page 1250) with rules written to protect those assets.

When you configure the FireSIGHT Recommended Rules feature, the system searches your base policy for rules that protect against vulnerabilities associated with your network assets, and identifies the current state of rules in your base policy. The system then recommends rule states and, optionally, sets the rules to the recommended states using the criteria in the following table.

FireSIGHT Rule State Recommendations Based on Vulnerabilities

BASE POLICY RULE STATE	RULE PROTECTS YOUR DISCOVERED ASSETS?	RECOMMEND RULE STATE
Generate Events or Disable	yes	Generate Events
Drop and Generate Events	yes	Drop and Generate Events
any	no	Disable

The Sourcefire Vulnerability Research Team (VRT) determines the appropriate state of each rule in the default policies provided by Sourcefire. Thus, when your base policy is a default policy provided by Sourcefire, the net effect of allowing the system to set your rules to the FireSIGHT recommended rule states is that

the rules in your intrusion policy match the settings recommended by Sourcefire for your network assets. See [Using Default Intrusion Policies](#) on page 692 for more information.

Generating rule state recommendations can be as simple as choosing whether to use the recommended rule states, either when you generate recommendations or at a later time. Advanced recommendations options allow you to tailor your configuration.

Note that while the system typically recommends rule state changes for standard text rules and shared object rules, changes can also be recommended for preprocessor and decoder rules.

You can schedule a task to generate recommendations automatically based on the most recently saved configuration settings in your intrusion policy. For information on scheduling a task to generate recommended rule states, see [Automating FireSIGHT Recommendations](#) on page 1943.

See the following sections for more information:

- [Understanding Basic Rule State Recommendations](#)
- [Understanding Advanced Rule State Recommendations](#)
- [Using FireSIGHT Recommendations](#)

Understanding Basic Rule State Recommendations

LICENSE: Protection + FireSIGHT

You can generate recommendations without using the recommended rule states in your policy. You can then display any of three filtered views of the Rules page to show rules that the system recommends you set to Generate Events, Drop and Generate Events, or Disable. This allows you to see beforehand which rules would be modified when you elect to use the recommended rule states. You can also choose to generate recommendations and immediately use them.

While displaying the recommendation-filtered Rules page, or after accessing the Rules page directly from the navigation panel or the Policy Information page, you can manually set rule states, sort rules, and take any of the other actions available on the Rules page such as suppressing rules, setting rule thresholds, and so on. See [Setting Rule States](#) on page 743 for information on manually changing the state of selected rules. See [Managing Rules in an Intrusion Policy](#) on page 717 for information on other actions available on the Rules page for tailoring the rules in your intrusion policy.

The system will not change rule states that you set manually. When you elect to use the recommended rule states while generating recommendations:

- manually setting the states of specified rules before you generate recommendations prevents the system from modifying the states of those rules in the future
- manually setting the states of specified rules after you generate recommendations overrides the recommended states of those rules

TIP! You can include a list in the intrusion policy report of rules whose rule state differs from the recommended state. See [Viewing an Intrusion Policy Report](#) on page 682 for more information.

Note that choosing to use recommended rule states adds a read-only FireSIGHT Recommendations layer to your intrusion policy, and subsequently choosing not to use recommended rule states removes the layer. See [Using Layers in an Intrusion Policy](#) on page 792 for information on using policy layers to more efficiently manage multiple intrusion policies.

Note also that when you generate recommendations without changing the advanced settings for FireSIGHT recommended rules, the system recommends rule state changes for hosts identified by the \$HOME_NET variable. Also by default, the system generates recommendations only for rules with low or medium overhead, and generates recommendations to disable rules. See [Understanding Advanced Rule State Recommendations](#) on page 766 for more information.

Understanding Advanced Rule State Recommendations

LICENSE: Protection or Protection + FireSIGHT

Advanced settings allow you to redefine which hosts on your network the system monitors for vulnerabilities, to influence which rules the system recommends based on rule overhead, and to specify whether to generate recommendations to disable rules.

If you want to dynamically adapt active rule processing for specific packets based on host information, you can also enable adaptive profiles. For more information, see [Adaptive Profiles and FireSIGHT Recommended Rules](#) on page 1005.

See the following sections for more information:

- [Understanding the Network to Examine](#) on page 766
- [Understanding Rule Overhead](#) on page 767

Understanding the Network to Examine

LICENSE: Protection + FireSIGHT

You configure the FireSIGHT Recommended Rules feature by identifying a network in the network map to examine. The system then recommends the rules

you can activate to protect your network. For information on the network map, see [Using the Network Map](#) on page 1314.

You can specify the network to examine for FireSIGHT recommendations using the following methods, separately or combined:

- by using the value defined for the \$HOME_NET variable in your intrusion policy
Note that the value of **any** is used if you do not explicitly set a policy-specific value. See [Managing Variables](#) on page 698 and [Understanding Existing Variables](#) on page 699 for more information.
- by specifying hosts to examine for FireSIGHT recommendations on your network
- by creating a policy-specific value for the \$HOME_NET variable
Note that a policy-specific \$HOME_NET value is used regardless of whether you specify the hosts to examine.

Lists of addresses within the \$HOME_NET variable and within hosts that you specify are linked with an **OR** operation except for negations, which are linked with an **AND** operation after all **OR** operations are calculated. When you combine the two addressing methods, the \$HOME_NET variable and the additional hosts you specify are also linked with an **OR** operation.

Understanding Rule Overhead

LICENSE: Protection

Sourcefire rates the overhead of each intrusion rule as none, low, medium, high, or very high based on the rule's potential impact on system performance and the likelihood that the rule might generate false positives. You can view the overhead rating for a rule in the rule detail view on the Rules page. See [Viewing Rule Details](#) on page 723 for more information.

You can set the system to make rule state recommendations based on all rules up to and including a specified overhead rating, except for very high. You must manually set the rule state for any rule with a very high overhead rating. For example, when you generate recommendations for rules with medium overhead, the system makes recommendations based on all rules with an overhead rating of none, low, or medium, and does not make any recommendations for rules with high or very high overhead.

Note that the system factors rule overhead into recommendations to generate events or to drop and generate events. The system does not factor rule overhead into recommendations to disable rules. Note also that local rules do not have any overhead, unless they are mapped to a third-party vulnerability. See [Importing Local Rule Files](#) on page 2080 and [Managing Third-Party Product Mappings](#) on page 1682 for more information.

Generating recommendations for rules with the overhead rating at a particular setting does not preclude you from generating recommendations with different overhead and then generating recommendations again for the original overhead

setting. You will get the same rule state recommendations for each overhead setting each time you generate recommendations for the same rule set, regardless of the number of times you generate recommendations or with how many different overhead settings you generate. For example, you can generate recommendations with overhead set to medium, then to high, then to very high, and then to medium again and, if the hosts and applications on your network have not changed, both sets of recommendations with overhead set to medium will be the same for that rule set.

Using FireSIGHT Recommendations

LICENSE: FireSIGHT + Protection

You can generate recommendations with or without using the recommended rule states, and with or without modifying the advanced settings for generating recommendations. See [Understanding Basic Rule State Recommendations](#) on page 765 and [Understanding Advanced Rule State Recommendations](#) on page 766 for more information.

After generating recommendations, you can use the recommended rule states; you can also view recommended states and use any features available on the Rules page.

To use FireSIGHT rule state recommendations:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

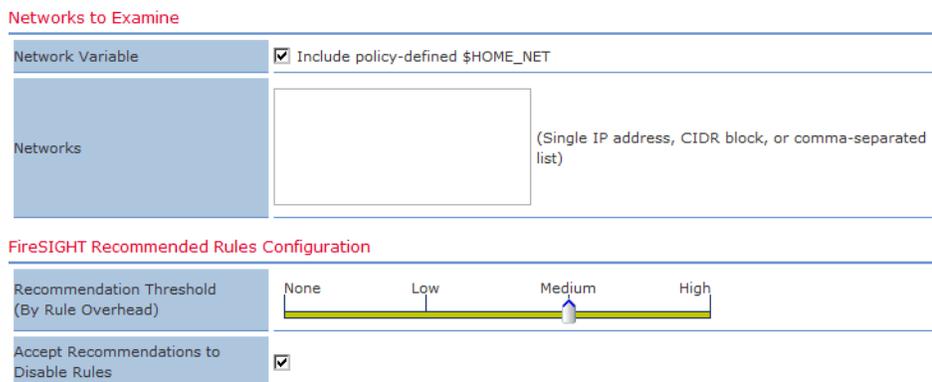
- You have two options:
 - If you have not generated recommendations, click **No recommendations have been generated. Click here to set up FireSIGHT recommendations.**
 - If you have generated recommendations, select **Click to change recommendations.**

The FireSIGHT Recommended Rules Configuration page appears.



- You have the following choices:
 - To have the corresponding intrusion policy report list the rule message, recommended state, and actual state for all rules whose actual state differs from the recommended state, select **Include all differences between recommendations and rule states in policy reports.** See [Viewing an Intrusion Policy Report](#) on page 682 for more information.
 - To generate recommendations using the default settings, go to step 9.
 - To modify the advanced recommendations options, go to step 5.
- Click the plus icon (**+**) to expand the Advanced Settings section.

The advanced FireSIGHT recommendations options appear.



6. You have three choices:
 - To specify the network to examine for FireSIGHT recommendations using the value of the `$HOME_NET` variable, ensure that the **Network Variable** check box is selected. See [Managing Variables](#) on page 698 for more information.

Note that this option is selected by default. When not selected, a default value of `any` is used. See [Modifying Variables](#) on page 701 for information on configuring the value of `$HOME_NET`.

- To manually specify the network to examine for recommendations, clear the **Network Variable** check box and identify the host addresses to examine for recommendations in the **Networks** field.

For information on using IP address notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

- To specify the network to examine using both the `$HOME_NET` variable and additional hosts you specify, ensure that the **Network Variable** check box is selected and, in the **Networks** field, identify additional hosts to examine for recommendations.

Note that lists of addresses and combined configurations of the `$HOME_NET` variable and **Networks** are linked with an **OR** operation except for negations, which are linked with an **AND** operation after all **OR** operations are calculated. See [Understanding the Network to Examine](#) on page 766 for more information.

7. Optionally, drag the **Recommendation Threshold (By Rule Overhead)** slide bar to specify the amount of overhead a rule must have to be included in the recommendations you generate.

Dragging the slide bar to the right includes rules with higher overhead and will likely result in more recommendations, but could increasingly affect system performance. See [Understanding Rule Overhead](#) on page 767 for more information.

8. You have the following options:
 - To generate recommendations to disable rules, select the **Accept Recommendations to Disable Rules** check box.

Note that accepting recommendations to disable rules restricts your rule coverage.

- To prevent generating recommendations to disable rules, clear the **Accept Recommendations to Disable Rules** check box.

Note that omitting recommendations to disable rules augments your rule coverage.

9. You have several options:
- Click **Generate and Use Recommendations** if you have not yet generated recommendations and want the system to change your rule states automatically to the recommended states while generating recommendations.
The system generates recommended rule state changes and automatically sets rules to the recommended states.
 - Click **Generate Recommendations** if you want the system to generate recommendations without changing your rule states automatically to the recommended states.
The system generates recommended rule state changes.
 - Click **Update Recommendations** to update existing recommendations.
The system generates recommended rule state changes and, if recommendations are in use, automatically sets rules to the recommended states. The status updates for the number of recommendations, the number of hosts with recommended rule state changes, and the number of recommendations to generate events, drop and generate events, or disable rules.
 - Click **Use Recommendations** to use recommendations that you have generated but have not used.
The system automatically sets rules to the recommended states.
 - Click **Do Not Use Recommendations** to stop using recommendations currently in use.
The system automatically resets rules to the default rule states unless a specific rule state was applied to the rule before using recommendations; in that case, the rule reverts to the specific rule state.

Note that the system does not recommend a rule state for an intrusion rule that is based on a vulnerability that you disable using the Impact Qualification feature. For more information, see [Setting the Vulnerability Impact Qualification](#) on page 1367.

Note also that updating the policy to use or not use recommendations may take several minutes, depending on the size of your network and rule set.

IMPORTANT! The system always recommends that you enable a local rule associated with a third-party vulnerability mapped to a host. The system does not make state recommendations for unmapped local rules. For more information, see [Managing Third-Party Product Mappings](#) on page 1682.

10. Optionally, click **View** next to a recommendation type to display a recommendations-filtered view of the Rules page for the type of recommendation you selected.

11. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 21

USING ADVANCED SETTINGS IN AN INTRUSION POLICY

Advanced settings are preprocessor and other intrusion policy detection and performance configurations that require specific expertise to configure. Advanced settings typically require little or no modification and are not common to every deployment.

You can enable, disable, and modify the configuration of advanced settings. The base policy for your intrusion policy determines which advanced settings are enabled by default and the default configuration for each. See [Understanding the Base Policy](#) on page 691 for more information.

Some advanced settings must be enabled for certain standard text rules, shared object rules, and preprocessor rules to function correctly. When you save an intrusion policy with a required advanced setting that is disabled, you are prompted whether you want the system to automatically enable the required advanced setting.

The web interface identifies some advanced configuration options as troubleshooting options that you should use only with the assistance of Sourcefire Support.

See the following sections for more information:

- [Modifying Advanced Settings](#) on page 774 explains how to access configuration pages for advanced settings and lists the advanced settings that you can enable, disable, and configure in an intrusion policy.
- [Understanding preprocessors](#) on page 780 explains how preprocessors normalize traffic for use by the rules engine.
- [Automatically Enabling Advanced Settings](#) on page 787 explains how you can automatically enable preprocessors and other advanced settings that are required by enabled rules or rule options.

- [Understanding Troubleshooting Options](#) on page 790 explains troubleshooting options that you should set only when asked to do so by Sourcefire Support.
- [Using Layers in an Intrusion Policy](#) on page 792 explains how you can more efficiently manage multiple intrusion policies in a complex network by adding intrusion policy layers comprised of individual configurations for rule attributes and advanced settings.

Modifying Advanced Settings

LICENSE: Protection

When you select **Advanced Settings** in the navigation panel, you go to the Advanced Settings page, where advanced settings are listed by type. On this page you can enable or disable advanced settings in your intrusion policy and access advanced setting configuration pages.

An advanced setting must be enabled for you to configure it. Your configuration is retained if you configure an advanced setting and then disable 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. When you disable an advanced setting, the advanced setting sublink and **Edit** link no longer appear.

TIP! You cannot disable the Performance Statistics Configuration advanced setting. This ensures that Sourcefire Support can troubleshoot your system.

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.

Application Layer Preprocessors

Application-layer protocol decoders normalize specific types of packet data into formats that the rules engine can analyze. See [Using Application Layer Preprocessors](#) on page 809 for more information.

Application Layer Preprocessor Settings

FOR INFORMATION ON...	SEE...
DCE/RPC Configuration	Configuring the DCE/RPC Preprocessor on page 821
DNS Configuration	Configuring the DNS Preprocessor on page 830
FTP and Telnet Configuration	Decoding FTP and Telnet Traffic on page 832
HTTP Configuration	Decoding HTTP Traffic on page 850
Sun RPC Configuration	Configuring the Sun RPC Preprocessor on page 870
SIP Configuration	Configuring the SIP Preprocessor on page 875
GTP Command Channel Configuration	Configuring the GTP Command Channel on page 878
IMAP Configuration	Configuring the IMAP Preprocessor on page 882
POP Configuration	Configuring the POP Preprocessor on page 887
SMTP Configuration	Configuring SMTP Decoding on page 895
SSH Configuration	Selecting SSH Preprocessor Options on page 901
SSL Configuration	Configuring the SSL Preprocessor on page 907

SCADA Preprocessors

The Modbus and DNP3 preprocessors detect traffic anomalies and provide data to the rules engine for inspection.

SCADA Preprocessor Settings

FOR INFORMATION ON...	SEE...
Modbus Configuration	Configuring the Modbus Preprocessor on page 909
DNP3 Configuration	Configuring the DNP3 Preprocessor on page 911

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 rules engine; it also detects various anomalous behaviors in packet headers.

Transport and Network Layer Preprocessor Settings

FOR INFORMATION ON...	SEE...
Checksum Verification	Verifying Checksums on page 915
Detection Settings	Ignoring VLAN Headers on page 917
Inline Normalization	Normalizing Inline Traffic on page 918
IP Defragmentation	Configuring IP Defragmentation on page 932
Packet Decoding	Configuring Packet Decoding on page 938
TCP Stream Configuration	Configuring TCP Stream Preprocessing on page 952
UDP Stream Configuration	Configuring UDP Stream Preprocessing on page 957

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. The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text.

Specific Threat Detection Settings

FOR INFORMATION ON...	SEE...
Back Orifice Detection	Detecting Back Orifice on page 959
Portscan Detection	Configuring Portscan Detection on page 965
Rate-Based Attack Prevention	Configuring Rate-Based Attack Prevention on page 981
Sensitive Data Detection	Configuring Sensitive Data Detection on page 989

Detection Enhancement

With adaptive profiles, the system can adapt to network traffic by associating traffic with host information from the network map and then processing the traffic accordingly.

Detection Enhancement Settings

FOR INFORMATION ON...	SEE...
Adaptive Profiles	Configuring Adaptive Profiles on page 1006

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.

Intrusion Rule Threshold Settings

FOR INFORMATION ON...	SEE...
Global Rule Thresholding	Configuring Global Thresholds on page 1012

Performance Settings

The system provides server settings for improving system performance.

Performance Settings

FOR INFORMATION ON...	SEE...
Event Queue Configuration	Event Queue Configuration on page 1016
Latency-Based Packet Handling	Configuring Packet Latency Thresholding on page 1020
Latency-Based Rule Handling	Configuring Rule Latency Thresholding on page 1025
Performance Statistics Configuration	Performance Statistics Configuration on page 1026
Regular Expression Limits	Constraining Regular Expressions on page 1028
Rule Processing Configuration	Rule Processing Configuration on page 1030

External Responses

In addition to the various views of intrusion events within the web interface, you can enable logging to syslog facilities or send event data to an SNMP trap server. You can specify intrusion event notification limits, set up intrusion event

notification to external logging facilities, and configure external responses to intrusion events.

External Response Settings

FOR INFORMATION ON...	SEE...
SNMP Alerting	Configuring SNMP Responses on page 1037
Syslog Alerting	Configuring Syslog Responses on page 1040

To modify advanced settings:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. To select the advanced setting you want to modify, click **Advanced Settings** in the navigation panel on the left, enable the configuration if it is disabled, then click **Edit**.
The configuration page appears. You can modify any of the configuration options for the advanced setting you selected.
To access the configuration page for an advanced setting that is enabled, you can also expand **Advanced Settings** in the navigation panel on the left, then click the name of the advanced setting.

TIP! You cannot disable the Performance Statistics Configuration advanced setting. This ensures that Sourcefire Support can troubleshoot your system.

4. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Understanding Preprocessors

LICENSE: Protection

Preprocessors reformat traffic to make sure the rules engine reads the traffic in the same format it will be received by the host. Without preprocessing, the system cannot appropriately evaluate traffic because protocol differences make pattern matching impossible. Sourcefire preprocessors normalize traffic and help identify network layer and transport layer protocol anomalies by identifying inappropriate header options, defragmenting IP datagrams, providing TCP stateful inspection and stream reassembly, providing UDP stream preprocessing, resolving application protocol command syntax, and validating checksums.

You can configure these preprocessors to ensure that the packets the system analyzes resemble, as closely as possible, the packets processed by the hosts on your network. Each preprocessor has a variety of options and settings that you can configure to meet the needs of your network environment, allowing you to minimize both false positives and false negatives and to optimize performance by executing only those preprocessors appropriate to your network traffic.

In general, as intrusion detection and prevention systems become important components in securing networks, the systems themselves become targets for attackers. For example, attackers sometimes attempt to purposefully create denial of service attacks by sending SYN packets with spoofed source IP addresses, causing the recipient server to allocate memory for the pending TCP connection. The server then sends a SYN-ACK to the originating IP address to establish a TCP session. Because attackers do not use legitimate IP addresses, the SYN-ACK message times out and the server resends it, keeping memory allocated for a longer period of time. These half-open TCP connections drain system resources. Because most systems attempt to perform stateful inspection on TCP sessions, the system may go into a denial-of-service condition while attempting to establish the state of these open TCP sessions. However, the transport layer preprocessor, included as part of the system, detects the state of a TCP connection, and can dispense with half-open connections and prevent overloading the rules engine with false connections.

Preprocessor options can protect you from attacks against the managed device itself, ensuring higher availability and better security for your network. Many preprocessor options are associated with preprocessor rules that you can enable to generate events when triggered. If you deploy your Sourcefire 3D System inline, you can set the rule state for preprocessor rules in your inline intrusion policy to drop malicious packets. For more information on configuring rules to generate events and, in an inline deployment, to drop packets, see [Setting Rule States](#) on page 743.

You can configure rule state, thresholding, suppression, rate-based rule state, alerting, and rule comments for preprocessor rules. Preprocessor rules are listed by preprocessor in the Preprocessors filter group on the intrusion policy Rules page, and also in the preprocessor and packet decoder sub-groupings in the Category filter group. You must set the rule state of preprocessor and decoder rules to Generate Events or, optionally, to Drop and Generate events in an inline

deployment, if you want the preprocessor or packet decoder to log intrusion events. Note that a status message appears at the bottom of the Policy Information page when you enable preprocessor rules and your policy contains unsaved changes. See [Editing an Intrusion Policy](#) on page 675, [Managing Rules in an Intrusion Policy](#) on page 717, and [Setting Rule States](#) on page 743 for more information.

In addition to preprocessors, the system also provides advanced settings for detecting anomalous traffic, enhancing detection, applying a global rule threshold, tuning performance, and configuring external SNMP, and syslog alerting.

See the following sections for more information:

- [Meeting Traffic Challenges with Preprocessors](#) on page 781 describes both normal traffic and the inspection challenges experienced at the network layer, transport layer, and application layer.
- [Understanding Preprocessor Execution Order](#) on page 782 explains the order of execution in Sourcefire 3D System preprocessors.
- [Reading Preprocessor Events](#) on page 784 describes preprocessor events and the information they contain.

Meeting Traffic Challenges with Preprocessors

LICENSE: Protection

The system is responsible for inspecting the traffic that traverses the segment of your network that you want to monitor. Although this seems straightforward, variations in the way data is represented and the characteristics inherent in the way data is transmitted can make the inspection of any traffic more complex. The Sourcefire 3D System mitigates the challenges inherent in normal traffic, as well as those inherent in packets designed to cause damage or to evade inspection.

Each layer of TCP/IP provides challenges:

- Network and Link Layers
Normal traffic at the network layer can be fragmented. That is, IP datagrams can exceed the maximum transmission unit and must be transported in smaller fragments. IP Datagrams that are fragmented must be reconstructed before meaningful attack analysis can occur. Additionally, attackers can use malicious IP fragmentation, including overlapping fragments, multiple zero-offset fragments (the Jolt2 denial of service, or DoS, attack), and fragmented protocol headers, all of which mask traffic you might not normally allow on your network. Additionally, the network layer can be attacked by crafting packets with invalid, zero-length IP options, used to cause DoS attacks.
- Transport Layer
The transport layer is subject to TCP stream-based attacks, such as sending TCP packets with overlapping sequence numbers to force the system to determine which sequence number is valid. The transport layer can be open to TCP header option attacks such as spoofing a TCP packet and changing

header values to choke the TCP connection and propagate further attacks. Additionally, TCP is subject to state-related attacks such as those produced by stick or snot, which generate TCP packets that are not part of an established connection and which can trigger a large volume of rules, creating a DoS attack against both the system and the analyst. Attackers can also launch subterfuge attacks by transmitting TCP, UDP and ICMP packets with invalid checksums in an attempt to cause the system to inspect packets that the destination host never receives. Reassembling TCP sessions provides context for each packet, supporting effective analysis of traffic.

Additionally, tracking associated UDP user datagrams allows the system greater specificity in detecting attacks.

- Application Layer

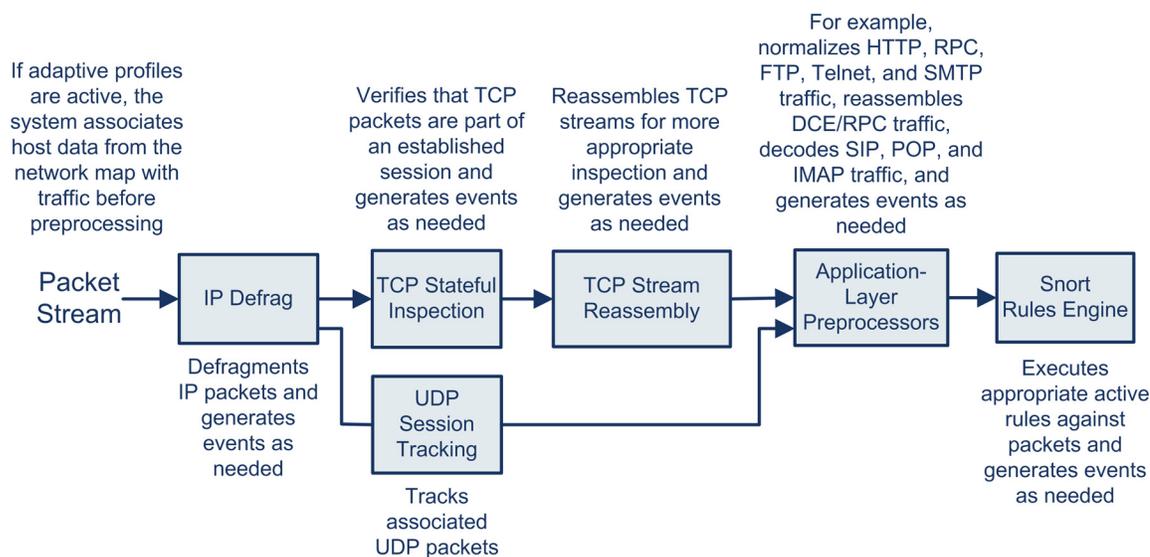
Application layer protocols like HTTP, Telnet, FTP, SMTP, and RPC may have multiple ways of representing the same data. This causes rules designed to check for specific packet payload content to fail because the payload is represented differently in a packet than in the rule. Decoding HTTP, Telnet, FTP, SMTP, and RPC packets and then normalizing their data to a standard representation mitigates this challenge.

Understanding Preprocessor Execution Order

LICENSE: Protection

Protocol decoders, preprocessors, and rules run in a specific order so that they can perform IP transfer layer protocol decoding first, then perform data normalization if needed, and then evaluate the resulting packets against the currently enabled rules. The default policy configuration sets the preprocessors to

perform IP transfer layer protocol decoding first, then perform data normalization as needed.



This approach provides the following benefits:

- The system can generate an intrusion event against fragmented IP datagrams that cannot be defragmented, and then stop inspecting those packets.
- The system can generate an event against TCP packets whose state cannot be validated, and then stop inspecting those packets.
- The system can generate events against related UDP packets.
- Only packets that can be appropriately tested by rules are normalized, optimizing performance by ignoring TCP packets that cannot be reassembled and are not part of a valid TCP session.
- The system can adapt IP defragmentation and stream preprocessing behavior to fit the operating system formats on the target host using adaptive profiles, target-based policies, or both adaptive profiles and target-based policies.
- After preprocessing, traffic can be analyzed by the rules engine in the same way that it is analyzed by the receiving host.

WARNING! Preprocessors are executed based on your configuration. If you change the default configuration, the system will not execute the preprocessors you disabled. If, for example, you disable transport layer protocol preprocessors, the system runs content rules against packets that may have been logged and removed from inspection by transport layer protocol preprocessors had they inspected the packets. Note this does not change the order of execution.

Reading Preprocessor Events

LICENSE: Protection

Preprocessors provide two functions: performing the specified action on the packet (for example, decoding and normalizing HTTP traffic) and reporting the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled (for example, you can enable the **double Encoding** HTTP Inspect option and the associated preprocessor rule with the HTTP Inspect generator (GID) 119 and the Snort ID (SID) 2 to generate an event when the preprocessor encounters IIS double-encoded traffic). Generating events to report the execution of preprocessors helps you detect anomalous protocol exploits. For example, attackers can craft overlapping IP fragments to cause a DoS attack on a host. The IP defragmentation preprocessor can detect this type of attack and generate an intrusion event for it.

See the following sections for more information:

- [Understanding the Preprocessor Event Packet Display](#) on page 784 describes the information contained in a preprocessor-generated event.
- [Reading Preprocessor Generator IDs](#) on page 784 details the information provided by the preprocessor generator ID.

Understanding the Preprocessor Event Packet Display

LICENSE: Protection

Preprocessor events differ from rule events in that the packet display does not include a detailed rule description for the event. Instead, the packet display shows the event message, the generator ID, Snort ID, the packet header data, and the packet payload. This allows you to analyze the packet's header information, determine if its header options are being used and if they can exploit your system, and inspect the packet payload. After the preprocessors analyze each packet, the rules engine executes appropriate rules against it (if the preprocessor was able to defragment it and establish it as part of a valid session) to further analyze potential content-level threats and report on them.

Reading Preprocessor Generator IDs

LICENSE: Protection

Each preprocessor has its own Generator ID number, or GID, that indicates which preprocessor was triggered by the packet. Some of the preprocessors also have related SIDs, which are ID numbers that classify potential attacks. This helps you analyze events more effectively by categorizing the type of event much the way a rule's Snort ID (SID) can offer context for packets triggering rules. You can list preprocessor rules by preprocessor in the Preprocessors filter group on the intrusion policy Rules page; you can also list preprocessor rules in the preprocessor and packet decoder sub-groupings in the Category filter group. See [Managing Rules in an Intrusion Policy](#) on page 717 and the [Rule Types table](#) on

page 719 for more information.

IMPORTANT! Events generated by standard text rules have a generator ID of 1. The event's SID indicates which specific rule triggered. For shared object rules, the events have a generator ID of 3 and a SID that indicates which specific rule was triggered.

The [Generator IDs](#) table describes the types of events that generate each GID.

Generator IDs

ID	COMPONENT	DESCRIPTION
1	Standard Text Rule	The event was generated when the packet triggered a standard text rule. See the Rule Types table on page 719 for more information.
2	Tagged Packets	The event was generated by the Tag generator, which generates packets from a tagged session. This occurs when the <code>tag</code> rule option is used. For more information, see Evaluating Post-Attack Traffic on page 1166.
3	Shared Object Rule	The event was generated when the packet triggered a shared object rule. See the Rule Types table on page 719 for more information.
102	HTTP Decoder	The decoder engine decoded HTTP data within the packet.
105	Back Orifice Detector	The Back Orifice Detector identified a Back Orifice attack associated with the packet. See Detecting Back Orifice on page 959 for more information.
106	RPC Decoder	The RPC decoder decoded the packet. See Using the Sun RPC Preprocessor on page 869 for more information.
116	Packet Decoder	The event was generated by the packet decoder. See Understanding Packet Decoding on page 934 for more information.
119, 120	HTTP Inspect Preprocessor	The event was generated by the HTTP Inspect preprocessor. GID 120 rules relate to server-specific HTTP traffic. See Decoding HTTP Traffic on page 850 for more information.
122	Portscan Detector	The event was generated by the portscan flow detector. See Detecting Portscans on page 961 for more information.
123	IP Defragmentor	The event was generated when a fragmented IP datagram could not be properly reassembled. See Defragmenting IP Packets on page 927 for more information.

Generator IDs (Continued)

ID	COMPONENT	DESCRIPTION
124	SMTP Decoder	The event was generated when the SMTP preprocessor detected an exploit against an SMTP verb. See Understanding SMTP Decoding on page 890 for more information.
125	FTP Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within FTP traffic. See Understanding Server-Level FTP Options on page 838 and Understanding Client-Level FTP Options on page 846 for more information.
126	Telnet Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within telnet traffic. See Decoding FTP and Telnet Traffic on page 832 for more information.
128	SSH Preprocessor	The event was generated when the SSH preprocessor detected an exploit within SSH traffic. See Detecting Exploits Using the SSH Preprocessor on page 899 for more information.
129	Stream Preprocessor	The event was generated during stream preprocessing by the stream preprocessor. See Using TCP Stream Preprocessing on page 940 for more information.
131	DNS Preprocessor	The event was generated by the DNS preprocessor. See Detecting Exploits in DNS Name Server Responses on page 827 for more information.
133	DCE/RPC Preprocessor	The event was generated by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 810 for more information.
134	Rule Latency, Packet Latency	The event was generated when rule latency suspended (134:1) or re-enabled (134:2) a group of intrusion rules, or when the system stopped inspecting a packet because the packet latency threshold was exceeded (134:3). For more information, see Understanding Packet Latency Thresholding on page 1017, Understanding Troubleshooting Options on page 790, and Understanding Rule Latency Thresholding on page 1021.
135	Rate-Based Attack Detector	The event was generated when a rate-based attack detector identified excessive connections to hosts on the network. See Preventing Rate-Based Attacks on page 971 for more information.
138, 139	Sensitive Data Preprocessor	The event was generated by the sensitive data preprocessor. See Detecting Sensitive Data on page 983 for more information.
140	SIP Preprocessor	The event was generated by the SIP preprocessor. See Decoding the Session Initiation Protocol on page 872 for more information.
141	IMAP Preprocessor	The event was generated by the IMAP preprocessor. See Decoding IMAP Traffic on page 880 for more information.

Generator IDs (Continued)

ID	COMPONENT	DESCRIPTION
142	POP Preprocessor	The event was generated by the POP preprocessor. See Decoding POP Traffic on page 884 for more information.
143	GTP Preprocessor	The event was generated by the GTP preprocessor. See Configuring the GTP Command Channel on page 878 for more information.
144	Modbus Preprocessor	The event was generated by the Modbus SCADA preprocessor. See Configuring the Modbus Preprocessor on page 909 for more information.
145	DNP3 Preprocessor	The event was generated by the DNP3 SCADA preprocessor. See Configuring the DNP3 Preprocessor on page 911 for more information.

Automatically Enabling Advanced Settings

LICENSE: Protection

The system can enable advanced settings when they are required by a standard text rule, shared object rule, preprocessor rule, or another advanced setting. When you save an intrusion policy with a disabled advanced setting that is required by a rule, rule option, or other advanced setting, you are prompted whether you want the system to automatically enable the required advanced setting. Before you can save the policy, you must either manually enable the required advanced setting configuration, allow the system to automatically enable the required advanced setting, or disable any rule or other advanced setting that requires the advanced setting.

Note that the system uses the default configuration for an automatically enabled advanced setting that you have not configured.

The [Automatically Enabled Advanced Settings](#) table lists the rules and rule options required by different advanced settings.

Automatically Enabled Advanced Settings

ADVANCED SETTING TYPE	ADVANCED SETTING	RULE AND RULE OPTIONS CAUSING AUTO-ENABLE PROMPT
Application Layer Preprocessors	DCE/RPC Configuration	Keyword: <ul style="list-style-type: none"> byte_jump (if DCE/RPC option is enabled) byte_test (if DCE/RPC option is enabled) byte_extract (if DCE/RPC option is enabled) dce_iface dce_opnum dce_stub_data
Application Layer Preprocessors	HTTP Configuration	Keyword: <ul style="list-style-type: none"> content (if an HTTP content option is enabled) urilen http_encode pcre (if the P, I, C, K, Y, M, U, S, H, or D option is used in the rule)
Application Layer Preprocessors	SIP Configuration	Keyword: <ul style="list-style-type: none"> sip_header sip_body sip_method sip_status_code
Application Layer Preprocessors	GTP Command Channel Configuration	Keyword: <ul style="list-style-type: none"> gtp_version gtp_type gtp_info
Application Layer Preprocessors	SSL Configuration	Keyword: <ul style="list-style-type: none"> ssl_state ssl_version
SCADA Preprocessors	Modbus Configuration	Keyword: <ul style="list-style-type: none"> modbus_data modbus_func modbus_unit
SCADA Preprocessors	DNP3 Configuration	Keyword: <ul style="list-style-type: none"> dnp3_data dnp3_func dnp3_ind dnp3_obj

Automatically Enabled Advanced Settings (Continued)

ADVANCED SETTING TYPE	ADVANCED SETTING	RULE AND RULE OPTIONS CAUSING AUTO-ENABLE PROMPT
Transport/Network Layer Preprocessors	TCP or UDP Stream Configuration	Keyword: <ul style="list-style-type: none"> • flow • flowbits • stream_size
Transport/Network Layer Preprocessors	TCP Stream Configuration	Keyword: stream_reassemble
Specific Threat Detection	Sensitive Data Detection	Generator ID: <ul style="list-style-type: none"> • 138 • 139
Performance Settings	Regular Expression Limits	Keyword: pcre

When you enable a preprocessor that requires stream preprocessing, you are prompted when you save the policy whether to enable stream preprocessing for the appropriate protocol if stream preprocessing is disabled.

You are prompted whether to enable TCP stream preprocessing when it is disabled and you enable the following preprocessors:

- the DCE/RPC preprocessor when the RPC over HTTP proxy, RPC over HTTP server, TCP, or SMB transport protocol is selected
- the DNS preprocessor
- the FTP/Telnet preprocessor
- the HTTP Inspect preprocessor
- the IMAP preprocessor
- the POP preprocessor
- the SMTP preprocessor
- the SSL preprocessor
- the Modbus preprocessor
- the DNP3 preprocessor
- portscan detection when the TCP protocol is selected
- rate-based attack prevention
- sensitive data detection

You are prompted whether to enable UDP stream preprocessing when it is disabled and you enable any of the following preprocessors:

- the DCE/RPC preprocessor with the UDP transport protocol selected
- the SIP preprocessor
- the GTP preprocessor

Understanding Troubleshooting Options

LICENSE: Protection

Sourcefire Support might ask you to modify one or more troubleshooting options during a troubleshooting call. Troubleshooting options appear on the configuration page for the advanced setting to which they are related. Although these options can be used in conjunction with the other options related to the advanced setting, changing the settings for these options will affect performance and should be done only with Support guidance.

The [Troubleshooting Options](#) table describes these troubleshooting options.

Troubleshooting Options

ADVANCED SETTING	OPTION	DESCRIPTION
FTP and Telnet Configuration (Policy)	Log FTP Command Validation Configuration	This FTP/Telnet target-based policy option enables or disables printing of the configuration information for each FTP command listed for the server. See Understanding Server-Level FTP Options on page 838 for more information.
Performance Statistics Configuration (Global)	Log Session/Protocol Distribution	This global performance option logs protocol distribution, packet length, and port statistics. See Performance Statistics Configuration on page 1026 for more information.
Performance Statistics Configuration (Global)	Summary	This global performance option instructs the system to calculate the performance statistics only when the Snort® process is shut down or restarted. Note that this option is only available when the Log Session/Protocol Distribution troubleshooting option is enabled. See Performance Statistics Configuration on page 1026 for more information.

Troubleshooting Options (Continued)

ADVANCED SETTING	OPTION	DESCRIPTION
TCP Stream Configuration (Global)	Session Termination Logging Threshold	<p>This global TCP option logs a message when an individual connection exceeds the specified threshold.</p> <p>A value of 0 turns off the message.</p> <p>The upper limit of 1GB is also restricted by the amount of memory on the managed device allocated for stream processing.</p> <p>See Using TCP Stream Preprocessing on page 940 for more information.</p>
TCP Stream Configuration (Policy)	Maximum Queued Bytes	<p>This TCP target-based policy option specifies the amount of data that can be queued on one side of a TCP connection.</p> <p>A value of 0 specifies an unlimited number of bytes.</p> <p>See Using TCP Stream Preprocessing on page 940 for more information.</p>
TCP Stream Configuration (Policy)	Maximum Queued Segments	<p>This TCP target-based policy option specifies the maximum number of bytes of data segments that can be queued on one side of a TCP connection.</p> <p>A value of 0 specifies an unlimited number of data segment bytes.</p> <p>See Using TCP Stream Preprocessing on page 940 for more information.</p>

CHAPTER 22

USING LAYERS IN AN INTRUSION POLICY

Larger organizations with many managed devices may have many intrusion policies to support the unique needs of different departments, business units or, in some instances, different companies. The rule settings and advanced settings in an intrusion policy are contained in building blocks called policy *layers*, which you can use to more efficiently manage multiple policies.

You can create and edit a policy without consciously using layers. You can modify rule settings and advanced settings and, if you have not added user layers to your policy, the system automatically includes your changes in a single configurable layer. Optionally, you can also add up to 200 layers where you can configure any combination of rule settings and advanced settings. You can copy, merge, move, and delete user layers and, most importantly, share individual user layers with other policies.

See the following sections for more information:

- [Understanding Intrusion Policy Layers](#) on page 792 explains the layers that comprise a basic policy and how you can use them.
- [Configuring User Layers](#) on page 803 explains how you can add, copy, merge, and share user-configurable layers, and how to view and access the configuration pages for rules and advanced settings.

Understanding Intrusion Policy Layers

LICENSE: Protection

A policy where you do not add layers includes the read-only base policy layer and a single user-configurable layer that is initially named My Changes. If you generate and use rule state recommendations based on network discovery data,

the system automatically inserts a read-only FireSIGHT Recommendations layer immediately above the base policy. You can copy, merge, move, or delete any user-configurable layer and set any user-configurable layer to be shared by other intrusion policies; note that the My Changes layer is a user-configurable layer.

Each policy layer contains complete settings for all intrusion rules, preprocessor rules, and advanced settings. The layer at the bottom of the stack includes all the settings from the base policy you selected when you created the policy. A setting in a higher layer in the policy layer stack 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 below where they are explicitly set.

The following figure shows an example intrusion policy layer stack that, in addition to the base policy layer and the initial My Changes layer, also includes two additional user-configurable layers and the FireSIGHT Recommendations layer. 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.

User Layer 2
User Layer 1
User Layer (My Changes)
FireSIGHT Recommendations Layer
Base Policy Layer

When the highest layer in your policy is a read-only layer, or a shared layer as described in [Sharing Layers](#) on page 794, 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) from the intrusion policy Rules page. See [Managing Rules in an Intrusion Policy](#) on page 717 for more information.
- enable, disable, or modify an advanced setting. See [Modifying Advanced Settings](#) on page 774 for more information.

All settings in the system-added layer are inherited except for the rule or advanced setting changes that resulted in the new layer.

Note that in the case where the highest layer is a shared layer, the system adds a layer when you have set the highest layer to be shared by other policies or you have added a shared layer to your policy.

When the system applies a policy to traffic, it *flattens* the layers; that is, it applies only one configuration for each option. If you configure, for example, a rule state for the same rule within more than one layer in an intrusion policy, the system applies the setting that is configured at the highest layer.

Note that 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 defaults in your base policy layer; your changes are made in a higher layer, so they override any changes that a rule update makes to your default policy. See [Importing Rule Updates and Local Rule Files](#) on page 2072 for more information.

TIP! You can create an intrusion policy based solely on the default settings in the base policy and, optionally, using rule state recommendations.

See the following sections for more information on using policy layers:

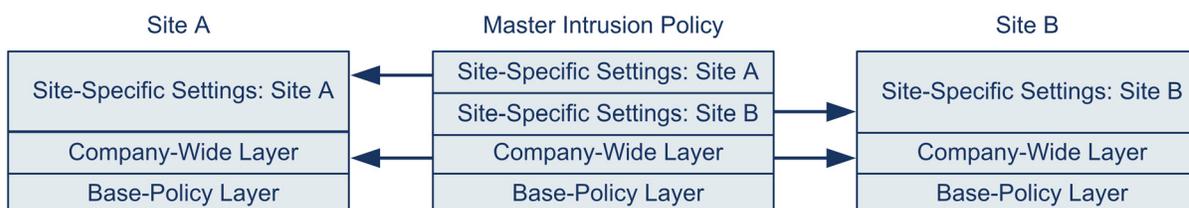
- [Sharing Layers](#) on page 794 provides an example intrusion policy that shows how you can share the settings in a layer with other intrusion policies.
- [Using Rules in Layers](#) on page 795 explains how you can work with rules in an intrusion policy layer.
- [Removing Multi-Layer Rule Settings](#) on page 797 explains how you can remove settings for event filters, dynamic states, and alerting from multiple layers using the intrusion policy Rules page.
- [Using the FireSIGHT Recommendations Layer](#) on page 799 explains how you can view and delete rule attributes in layers.
- [Using Layers with Advanced Settings](#) on page 800 explains how you can work with advanced settings in an intrusion policy layer.

Sharing Layers

LICENSE: Protection

You can share any user-configurable layer with other intrusion policies. When you share a layer and then edit a configuration within that layer, the system updates all policies that use the shared layer when you commit your changes and provides you with a list of all affected policies. A shared layer can only be modified in the policy where it is created.

The following figure shows an example master intrusion 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 intrusion policies at Site A and Site B. It also includes site-specific layers for each policy. For example, 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 settings in a higher layer in the site-specific policies, you could also further tune the policy for each site if necessary with any setting adjustments.

It is unlikely that the flattened net settings in the example master policy would be useful for motoring traffic, but the time saved in configuring and updating the site-specific policies makes this a useful application of policy layers.

Many other advanced layer configurations are possible. For example, you could define policy layers by company, by department, by network, or even by user. You could also include preprocessor settings in one layer, other advanced settings in a second layer, and rule settings in a third.

See the [Policy Layer Configuration Actions](#) table on page 804 for instructions on configuring shared layers.

TIP! You cannot add a shared layer to an intrusion policy where 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 messages indicates that the policy includes a circular dependency. See [Using a Custom Base Policy](#) on page 693 for more information.

Using Rules in Layers

LICENSE: Protection

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. You can view individual settings on the Rules page for the layer, or view the effective 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.

Note that you can switch to another layer at any time using the layer drop-down list.



The [Rule Settings in Multiple Layers](#) table describes the effects of configuring the same type of setting in multiple layers.

Rule Settings in Multiple Layers

YOU CAN SET...	OF THIS SETTING TYPE...	TO...
one	Rule State	<p>override a rule state set for the rule in a lower layer, and ignore all threshold, suppression, rate-based rule states, and alerts for that rule configured in lower layers. See Setting Rule States on page 743 for more information.</p> <p>If you want a rule to inherit the rule state for the rule from the base policy or a lower layer, set the rule state to Inherit. Note that you cannot set a rule state to Inherit when you are working on the intrusion policy Rules page, and the Inherit state does not appear in the Rule State column.</p> <p>Note also that rules with rule states set in a lower layer are highlighted in yellow and rules with states set in a higher layer are highlighted in red when you view them on the Rules page for a specific layer. Because the intrusion policy Rules page is a composite view of all rule settings, rule states are not color-coded on the policy view of the Rules page.</p>
one	Threshold SNMP Alert	<p>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 on page 747 and Adding Alerts on page 761 for more information.</p>
one or more	Suppression Rate-Based Rule State	<p>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 on page 753 and Adding Dynamic Rule States on page 756 for more information.</p>
one or more	Comment	<p>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 on page 729 for more information.</p>

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 view:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
2. Click the edit icon () next to the policy you want to view or edit.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Expand **Policy Layers** in the navigation panel and expand the policy layer you want to view or edit.
4. Click **Rules** under the policy layer you want to view or edit.
The Rules page for the layer appears.
You can modify any of the settings in the [Rule Settings in Multiple Layers table](#) on page 796.
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.

Removing Multi-Layer Rule Settings

LICENSE: Protection

You can select one or more rules on the intrusion policy view of the Rules page and then simultaneously remove a specific type of event filter, dynamic state, or alerting from multiple layers in your 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 ignores all layers below it.

When the system encounters the setting type in a shared layer or in the base policy, 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.

IMPORTANT! Removing rule settings 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** in the topmost layer. See [Setting Rule States](#) on page 743 for more information.

To remove settings in multiple layers using the Rules page:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
2. Click the edit icon (✎) next to the intrusion policy where you want to remove multiple settings.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. To access the intrusion policy Rules page, click **Rules** in the top of the navigation panel above the dividing line.

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. By default, the page lists the rules alphabetically by message.

4. Locate the rule or rules where you want to remove multiple settings. 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](#) on page 730 and [Setting a Rule Filter in an Intrusion Policy](#) on page 741.
The page refreshes to display all matching rules.

5. Select the rule or rules for 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.
6. You have the following options:
 - To remove all thresholds for a rule, select **Event Filtering > Remove Thresholds**. Click **OK** in the confirmation pop-up window that appears.
 - To remove all suppression for a rule, select **Event Filtering > Remove Suppressions**. Click **OK** in the confirmation pop-up window that appears.
 - To remove all rate-based rule states for a rule, select **Dynamic State > Remove Rate-Based Rule States**. Click **OK** in the confirmation pop-up window that appears.
 - To remove all SNMP alert settings for a rule, select **Alerting > Remove SNMP Alerts**. Click **OK** in the confirmation pop-up window that appears.

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.

IMPORTANT! Removing rule settings 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** in the topmost layer. See [Setting Rule States](#) on page 743 for more information.

7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Using the FireSIGHT Recommendations Layer

LICENSE: Protection

When you have generated rule state recommendations, you can choose whether to automatically modify rule states based on the recommendations.

Choosing to use the recommended rule states adds or updates a read-only, built-in FireSIGHT Recommendations system layer immediately above the base layer in your intrusion policy. Subsequently choosing not to use the recommended rule states removes the FireSIGHT Recommendations system layer. Note that you can repeatedly remove and restore the FireSIGHT Recommendations layer by choosing to use or not use recommendations, but you cannot delete the layer manually.

Adding the FireSIGHT Recommendations layer adds a FireSIGHT Recommendations link under Policy Layers in the navigation panel. That link leads you to a read-only view of the FireSIGHT Recommendations layer page. From the FireSIGHT Recommendations layer page, you can display recommendation-filtered views of the Rules page in read-only mode. On the Rules page, you can further filter the read-only recommendations, sort the display by column, and show details of individual rules. See [Managing Rules in an Intrusion Policy](#) on page 717 for more information on working with rules on the Rules page.

Adding the FireSIGHT Recommendations layer also adds a Rules sublink beneath the FireSIGHT Recommendations link in the navigation panel. The Rules sublink provides access to a read-only display of the Rules page in the FireSIGHT Recommendations layer. Note the following in this view:

- When there is no rule state icon in the state column, the state is inherited from the base policy.
- When there is no rule state icon in the FireSIGHT Recommendations column in this or other Rules page views, there is no recommendation for this rule.

Note that when a rule in the FireSIGHT Recommendations layer has no recommendation, its rule overhead rating was higher than the setting for **Recommendation Threshold (By Rule Overhead)** when recommendations were last generated. See [Understanding Rule Overhead](#) on page 767 for more information.

See [Managing FireSIGHT Rule State Recommendations](#) on page 764 for more information.

Using Layers with Advanced Settings

LICENSE: Protection

When you select **Advanced Settings** in the navigation panel, you go to the Advanced Settings page. On this page you can enable or disable advanced settings in your intrusion policy and access advanced setting configuration pages. The Advanced Settings page provides a summary of the effective states for all advanced settings in your intrusion policy. For example, if SSL Configuration is set to Disabled in one layer, then set to Enabled in a higher layer, the Advanced Settings page shows SSL Configuration as set to Enabled. Changes made in the Advanced Settings page appear in the top layer of the policy. See [Modifying Advanced Settings](#) on page 774 for more information on working with advanced settings on the Advanced Settings page.

When you expand **Policy Layers** in the navigation panel and then select any user-configurable layer, you go to the Layer summary page for the layer. On this page you can enable or disable advanced settings and access advanced setting configuration pages for the layer. You can also modify the layer name and description and configure whether to share the layer with other intrusion policies. See [Sharing Layers](#) on page 794 for more information.

If you want an advanced setting to inherit its state and configuration from the base policy or a lower layer, set the state to **Inherit**. Note that the Inherit state does not appear when you are working in the Advanced Settings page. You can switch to the Layer summary page for another layer at any time by selecting the layer name beneath **Policy Layers** in the navigation panel.

When you enable an advanced setting, a sublink to the configuration page for the advanced setting appears beneath the layer name in the navigation panel, and an **Edit** link to the configuration page for the advanced setting appears on the Layer summary page for the advanced setting you enabled. When you disable an advanced setting within a layer or set it to **Inherit**, the advanced setting sublink and **Edit** link no longer appear.

You can display the configuration page for an advanced setting from the Layer summary page by first enabling the configuration if it is disabled and then clicking on **Edit**. When the advanced setting is enabled in the layer, you can also display its configuration page by clicking on the sublink named for the advanced setting in the navigation panel under **Policy Layers**.

You can set the state (enabled or disabled) of advanced settings in the current layer or in a layer above or below that layer. Setting the state for an advanced setting in a layer overrides the state for that advanced setting in lower layers. When the advanced setting is enabled in a layer, the configuration in that layer also overrides the configuration of the advanced setting in lower layers.

Advanced setting states set in a different layer are color-coded to show whether they are set in a layer above or below. Note that because the Advanced Settings page is a composite view of all state settings, it does not use color coding to indicate where an advanced setting state is set in the layer order.

The system uses the configuration for an advanced setting in the highest layer where the configuration is enabled. Unless you explicitly modify the configuration, the system uses the default configuration. For example, if you enable and modify the DCE/RPC configuration in a layer, and you enable the DCE/RPC configuration but do not modify it in a higher layer, the system will use the default configuration in the higher layer.

You can view the layers where advanced settings are enabled, disabled, and inherited by clicking **Policy Layers** in the navigation panel. See [Configuring User Layers](#) on page 803 for more information.

The [Layer Summary Page Actions](#) table describes the actions available on the Layer summary page for user-configurable layers in your intrusion policy.

Layer Summary Page Actions

To...	YOU CAN...
modify the layer name or description	type a new value for Name or Description . Note that this action is not available on the Advanced Settings page.
share the layer with other intrusion policies	select Allow this layer to be used by other policies . See Using Layers with Advanced Settings on page 800 and Configuring User Layers on page 803 for more information. Note that this action is not available on the Advanced Settings page.
enable an advanced setting in the current layer	click Enabled next to the advanced setting you want to enable. The page refreshes, a sublink to the configuration page for the advanced setting appears beneath the layer name in the navigation panel, and an Edit link appears for the advanced setting you enabled. Optionally, click the Edit link or the advanced setting sublink to modify the current configuration. See Modifying Advanced Settings on page 774 for links to the configuration pages for all advanced settings. Note that the Back Orifice preprocessor has no user-configurable options.
disable the advanced setting in the current layer	click Disabled . The page refreshes and, if the advanced setting was enabled, the advanced setting sublink and Edit link no longer appear.
inherit the advanced setting state and configuration from the settings in the highest layer below the current layer	click Inherit . The page refreshes and, if the advanced setting was enabled, the advanced setting sublink and Edit link no longer appear.

To view or modify advanced settings in a layer view:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
2. Click the edit icon () next to the policy you want to view or edit.
If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.

3. Expand **Policy Layers** in the navigation panel and click the name of the layer you want to view or edit.
The Layer summary page for the layer appears.
4. Optionally, you can take any of the actions in the [Layer Summary Page Actions table](#) on page 802.
5. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Configuring User Layers

LICENSE: Protection

The Policy Layers page provides a single-page summary of all of the layers in your intrusion policy. For each layer, you can view whether an advanced setting is enabled or disabled in the layer or in a layer above or below it in the stack. You can also view the number of rules whose states are set in the layer, and the number of rules set to each rule state. You can also see a summary of the net effect of all enabled rules and advanced settings throughout the layers in the policy.

On this page you can also add shared and unshared layers, access rules and advanced settings to edit them within a layer, and copy, merge, move, and delete layers.

The [Policy Layer Configuration Actions](#) table explains how to view and interpret the policy layer summary and describes the layer configuration actions available on the Policy Layers summary page.

Policy Layer Configuration Actions

To...	YOU CAN...
add a shared layer from another policy	<p>click Add Shared Layer, then select the layer you want to add from the drop-down list in the Add Shared Layer pop-up window and click OK, or click Cancel if you decide not to add a shared layer.</p> <p>The Policy Layers summary page appears. If you selected a shared layer, the screen refreshes and the shared layer you selected appears as the highest layer in your policy.</p> <p>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.</p>
add a layer to your policy	<p>click Add Layer. Type a unique Name for the layer in the Add Layer pop-up window and click OK, or click Cancel if you decide not to add a layer. You can add up to 200 layers to an intrusion policy.</p> <p>The Policy Layers summary page appears. If you added a layer, the screen refreshes and the layer you added appears as the highest layer in your policy. Note that, in the new layer, the state of all advanced settings and rules is initially set to Inherit, and no event filtering, dynamic state, or alerting rule actions are set.</p>
enable or disable sharing a layer in your policy with other policies	<p>click the name of the layer in the navigation panel and select or clear the Sharing check box, then click Back to return to the Policy Layer summary page.</p> <p>Note that to disable sharing a layer that is in use in another policy, you must first delete the layer from the other policy or delete the other policy.</p>
move a layer above or below another layer	<p>click anywhere inside 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.</p> <p>The screen refreshes and the layer appears in the new location.</p>
manage rules or modify advanced setting configurations in a layer	<p>click the edit icon () for the layer.</p> <p>The Layer summary page for the layer appears. From this page you can display a layer-filtered view of the intrusion policy Rule page, enable, disable, or inherit advanced settings in the layer, and access advanced setting configuration pages in the layer. See Managing Rules in an Intrusion Policy on page 717 and Modifying Advanced Settings on page 774 for more information.</p> <p>Note that when you add a layer and enable an advanced setting in the new layer, the advanced setting configuration options are initially set to the default settings in the base policy.</p>

Policy Layer Configuration Actions (Continued)

To...	YOU CAN...
merge a layer into the next layer beneath it	<p>click the merge icon () for the layer you want to merge, then click OK when prompted or click Cancel to abandon the merge.</p> <p>The page refreshes and the layer is merged with the layer beneath it.</p> <p>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 rule or advanced setting. The merged layer retains the name of the lower layer.</p> <p>In the policy where you created a shared layer that you have added 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.</p> <p>In a policy where you have added 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 above the shared layer into the shared layer.</p>
copy a layer	<p>click the copy icon () of the layer you want to copy.</p> <p>The page refreshes and a copy of the layer appears as the highest layer. Note that copying a shared layer creates an unshared copy which, optionally, you can then identify as a layer that can be shared with other policies.</p>
delete a layer	<p>click the delete icon () for the layer you want to delete and then click OK at the prompt, or click Cancel if you decide not to delete the layer.</p> <p>The page refreshes and the layer is deleted.</p> <p>Note that you cannot delete a layer with sharing enabled if the layer is in use by another policy. Note also that you can delete the initial My Changes layer if it is unshared or if sharing is allowed but it has not been added to any other intrusion policies.</p>
display the Policy Information page	<p>click Policy Summary.</p> <p>See Managing Intrusion Policies on page 670 for an explanation of the actions you can take from the Policy Information page.</p>

Policy Layer Configuration Actions (Continued)

To...	YOU CAN...
display the Layer summary page for a layer	<p>click the layer name in the summary for the layer.</p> <p>The Layer summary page for the layer appears.</p> <p>From this page you can modify the layer name and description, set the layer to be shared by other intrusion policies, configure advanced settings states, and access advanced settings configuration pages. You can also display filtered Rules page views of rules whose states are set in the layer; you can display filtered views for all rules or by rule state. See Sharing Layers on page 794, Using Layers with Advanced Settings on page 800, and Using Rules in Layers on page 795 for more information.</p> <p>Note that, alternatively, you can click the view icon () to access the Layer summary page a shared layer. Note also that the Layer summary page for a shared layer is read-only.</p>
display the Layer summary page for the base policy	<p>click the base policy name in the base policy summary.</p> <p>The Layer summary page for the base policy appears.</p> <p>From this page you can select a different base policy for your intrusion policy and specify whether changes in an imported rule update your intrusion policy. You can view which advanced settings are enabled or disabled in your base policy and access read-only configuration pages showing the default configurations of advanced settings in the policy. Status messages give the number of rules enabled in the policy and the number set to generate events and to drop packets and generate events. From this page you can access a read-only view of the Rules page showing the settings for all rules in the base policy. See Understanding the Base Policy on page 691, Allowing Rule Updates to Modify the Base Policy on page 694, Using Layers with Advanced Settings on page 800, and Using Rules in Layers on page 795 for more information.</p>
display a layer-level advanced setting configuration page	<p>click the advanced setting name in the summary for the layer.</p> <p>Note that configuration pages are read-only in the base policy and in shared layers. See Sharing Layers on page 794 and Understanding the Base Policy on page 691, and for more information.</p>
display rules in a layer by rule state type	<p>click the icon for drop and generate events (), generate events (), or disabled () in the summary for the layer, or on the description next to the icon for the rule state type you want to display.</p> <p>Note that disabled rules are not displayed for the base policy or the Policy Summary. Note also that each layer summary provides the total number of rules enabled (that is, the total set to generate events or to drop and generate events) in the layer, and the total for each enabled rule state. Also, note that the base policy rule state totals are the default enabled rule state settings in the policy, and the Policy Summary totals are the effective sum of all enabled rule states for all layers in the policy.</p>

To configure layers in your intrusion policy:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Policy Layers** in the navigation panel.

The Policy Layers summary page appears, displaying a summary of the rule states and advanced settings in each layer, and a flattened view of the policy showing the net effect of the states in all layers.

Note that the advanced setting name in the summary for each layer indicates which advanced setting configurations are enabled, disabled, overridden, or inherited in the layer, as follows:

WHEN THE ADVANCED SETTING CONFIGURATION IS...	THE ADVANCED SETTING NAME IS...
enabled in the layer	written in plain text

WHEN THE ADVANCED SETTING CONFIGURATION IS...	THE ADVANCED SETTING NAME IS...
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

4. Optionally, you can take any of the actions in the [Policy Layer Configuration Actions](#) on page 804.
5. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 23

USING APPLICATION LAYER PREPROCESSORS

Application-layer protocols can represent the same data in a variety of ways. Sourcefire provides application-layer protocol decoders that normalize specific types of packet data into formats that the 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. See [Setting Rule States](#) on page 743 for more information.

See the following sections for more information:

- [Decoding DCE/RPC Traffic](#) on page 810 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](#) on page 827 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](#) on page 832 describes the FTP/Telnet decoder and explains how to configure it to normalize and decode FTP and Telnet traffic.
- [Decoding HTTP Traffic](#) on page 850 describes the HTTP decoder and explains how to configure it to normalize HTTP traffic.
- [Using the Sun RPC Preprocessor](#) on page 869 describes the RPC decoder and explains how to configure it to normalize RPC traffic.
- [Decoding the Session Initiation Protocol](#) on page 872 explains how you can use the SIP preprocessor to decode and detect anomalies in SIP traffic.

- [Configuring the GTP Command Channel](#) on page 878 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](#) on page 880 explains how you can use the IMAP preprocessor to decode and detect anomalies in IMAP traffic.
- [Decoding POP Traffic](#) on page 884 explains how you can use the POP preprocessor to decode and detect anomalies in POP traffic.
- [Decoding SMTP Traffic](#) on page 889 describes the SMTP decoder and explains how to configure it to decode and normalize SMTP traffic.
- [Detecting Exploits Using the SSH Preprocessor](#) on page 899 explains how to identify and process exploits in SSH-encrypted traffic.
- [Using the SSL Preprocessor](#) on page 905 explains how you can use the SSL preprocessor to identify encrypted traffic and eliminate false positives by stopping inspection of that traffic.
- [Working with SCADA Preprocessors](#) on page 909 explains how you can use the Modbus and DNP3 preprocessors to detect anomalies in corresponding traffic and provide data to the 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 IP defragmentation and TCP stream reassembly. Note that TCP stream preprocessing must be enabled to detect TCP-transported DCE/RPC, including SMB and RPC over HTTP, and IP defragmentation must be enabled when you enable the DCE/RPC preprocessor because, ultimately, IP transports all DCE/RPC traffic. See [Using TCP Stream Preprocessing](#) on page 940 and [Defragmenting IP Packets](#) on page 927.

Finally, the DCE/RPC preprocessor normalizes DCE/RPC traffic for processing by the rules engine. See [DCE/RPC Keywords](#) on page 1119 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. A link on the configuration page takes you to a filtered view of DCE/RPC preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.
- When a shared object rule or standard text rule that requires this preprocessor is enabled in an intrusion policy where the preprocessor is disabled, you must enable the preprocessor or choose to allow the system to enable it automatically before you can save the policy. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

See the following sections for more information:

- [Selecting Global DCE/RPC Options](#) on page 811
- [Understanding Target-Based DCE/RPC Server Policies](#) on page 813
- [Understanding DCE/RPC Transports](#) on page 814
- [Selecting DCE/RPC Target-Based Policy Options](#) on page 818
- [Configuring the DCE/RPC Preprocessor](#) on page 821

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](#) on page 1065 and [Using](#)

[Byte_Jump and Byte_Test](#) on page 1079.

If no preprocessor rule is mentioned, 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](#) on page 743 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](#) on page 813 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](#) on page 814 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 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.

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](#) on page 812.

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.

Sourcefire 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](#) on page 818 for recommendations for enabling or disabling auto-detection ports for each transport.

Note that any port configured for the **TCP Ports** or **TCP Auto-Detect Ports** option is automatically activated as a TCP stream preprocessor client or server reassembly port for the duration of a DCE/RPC session over the configured TCP port. Only TCP ports are activated, and TCP ports are automatically deactivated at the end of the session. See [Reassembling TCP Streams](#) on page 949 and [Selecting Stream Reassembly Options](#) on page 950 for more information.

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](#) on page 815
- [Understanding the RPC over HTTP Transport](#) on page 817

Understanding Connectionless and Connection-Oriented DCE/RPC Traffic

LICENSE: Protection

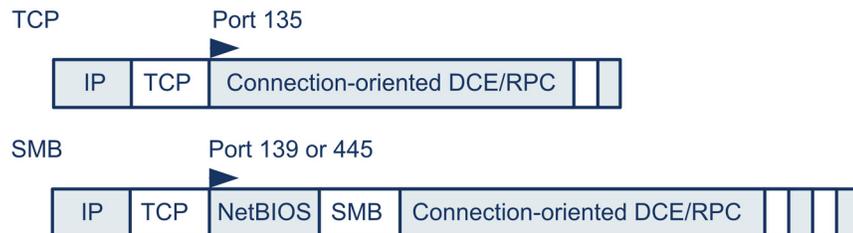
DCE/RPC messages comply with one of two distinct DCE/RPC Protocol Data Unit (PDU) protocols:

- 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.

Connection-oriented DCE/RPC



Connectionless DCE/RPC



▶ = DCE/RPC preprocessor starts decoding

Note the following in the figure:

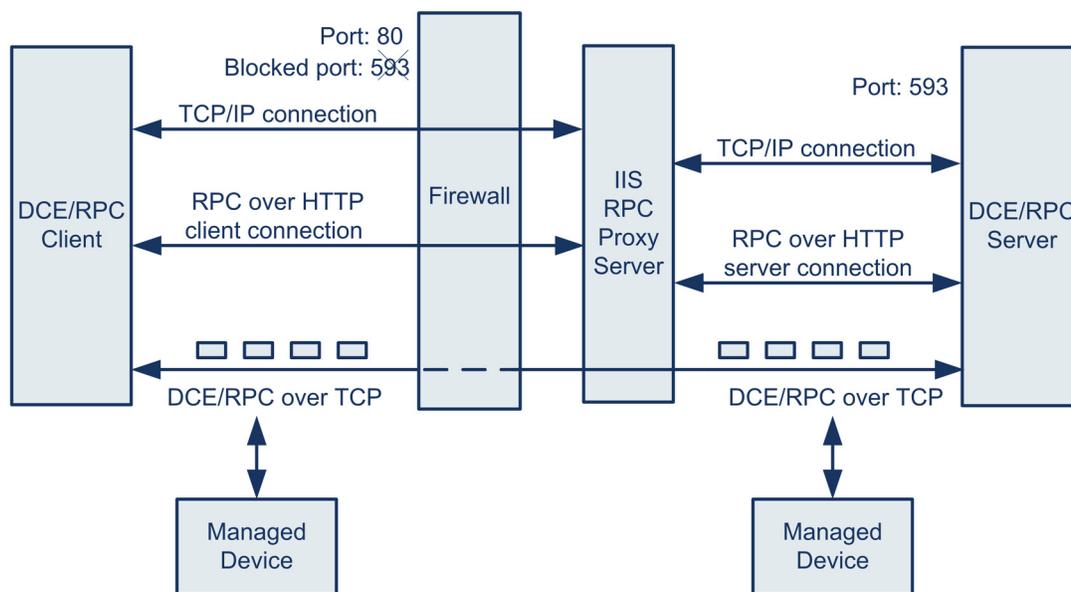
- The well-known TCP or UDP port 135 identifies DCE/RPC traffic in the TCP and UDP transports.
- The figure does not include RPC over HTTP.
For RPC over HTTP, connection-oriented DCE/RPC is transported directly over TCP as shown in the figure after an initial setup sequence over HTTP. See [Understanding the RPC over HTTP Transport](#) on page 817 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.
You must ensure that IP defragmentation is enabled when you enable the DCE/RPC preprocessor. See [Defragmenting IP Packets](#) on page 927 for more information.

- TCP transports all connection-oriented DCE/RPC.
You must ensure that TCP stream preprocessing is enabled when you enable the TCP, SMB, or RPC over HTTP transport. See [Using TCP Stream Preprocessing](#) on page 940 for more information.
- UDP transports connectionless DCE/RPC.
You must ensure that UDP stream preprocessing is enabled when you enable the UDP transport. See [Using UDP Stream Preprocessing](#) on page 956 for more information.

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.



Example 1: RPC over HTTP proxy

Example 2: RPC over HTTP server

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 network addresses where the policy is targeted, the type of Windows or Samba implementation used by the targeted host or hosts, whether to detect specified SMB invalid shares, a maximum number of chained SMB AndX commands to permit, the DCE/RPC transports in which you want to detect DCE/RPC traffic, and the ports to detect or auto-detect for each transport.

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, 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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).

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](#) on page 813 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](#) on page 812.

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](#) on page 743 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.

IMPORTANT! 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](#) on page 743 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 managed device is positioned between the DCE/RPC client and the Microsoft IIS RPC proxy server. See [Understanding the RPC over HTTP Transport](#) on page 817.

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](#) on page 817.

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 managed device is positioned between the DCE/RPC client and the MicroSoft IIS RPC proxy server. See [Understanding the RPC over HTTP Transport](#) on page 817.

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](#) on page 817.

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.

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. A link on the configuration page takes you to a filtered view of DCE/RPC preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Selecting Global DCE/RPC Options](#) on page 811 and [Selecting DCE/RPC Target-Based Policy Options](#) on page 818 for rules associated with specific detection options; see also [Setting Rule States](#) on page 743.

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 [Traffic-Associated DCE/RPC Rules](#) table identifies the rules that you can enable for each type of traffic.

Traffic-Associated DCE/RPC Rules

TRAFFIC	PREPROCESSOR RULE GID:SID
SMB	133:2 through 133:26, and 133:48 through 133:56
Connection-Oriented DCE/RPC	133:27 through 133:39
Detect Connectionless DCE/RPC	133:40 through 133:43

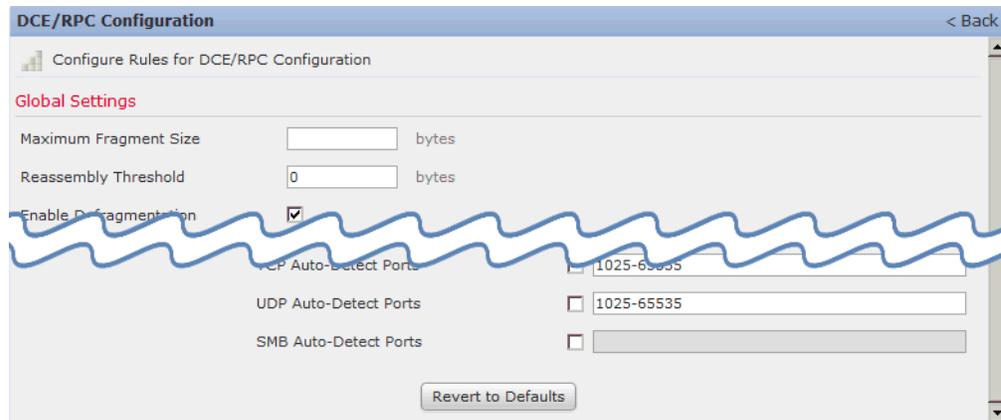
To configure the DCE/RPC preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can modify any of the options described in [Selecting Global DCE/RPC Options](#) on page 811.



6. 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

You can configure up to 255 policies, including the default policy.

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.

7. 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](#) on page 812.

- 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 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.

IMPORTANT! 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](#) on page 814, [Understanding the RPC over HTTP Transport](#) on page 817, and [Selecting DCE/RPC Target-Based Policy Options](#) on page 818 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](#) on page 814, [Understanding the RPC over HTTP Transport](#) on page 817, and [Selecting DCE/RPC Target-Based Policy Options](#) on page 818 for more information.

See [Selecting DCE/RPC Target-Based Policy Options](#) on page 818 for more information.

8. Optionally, click **Configure Rules for DCE/RPC Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the DCE/RPC Configuration page.
9. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 827
- [Detecting Overflow Attempts in RData Text Fields](#) on page 829
- [Detecting Obsolete DNS Resource Record Types](#) on page 829
- [Detecting Experimental DNS Resource Record Types](#) on page 830
- [Configuring the DNS Preprocessor](#) on page 830

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 [DNS Name Server RR Responses](#) table describes these three sections.

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

DNS Name Server RR Responses (Continued)

THIS SECTION...	INCLUDES...	FOR EXAMPLE...
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:

Name	
Type	Class
TTL	
RData Length	
RData	

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. TCP stream preprocessing must be enabled to enable the DNS preprocessor. However, you do not have to enable UDP session tracking because the DNS preprocessor inspects UDP traffic on a packet-by-packet basis. For more information, see [Using TCP Stream Preprocessing](#) on page 940 and [Using UDP Stream Preprocessing](#) on page 956.

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](#) on page 743 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 [Obsolete DNS Resource Record Types](#) table lists and describes these record types.

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](#) on page 743 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 [Experimental DNS Resource Record Types](#) table lists and describes these record types.

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](#) on page 743 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](#) on page 829, [Detecting Obsolete DNS Resource Record Types](#) on page 829, and [Detecting Experimental DNS Resource Record Types](#) on page 830.

To configure the DNS preprocessor:

ACCESS: Admin/Intrusion Admin

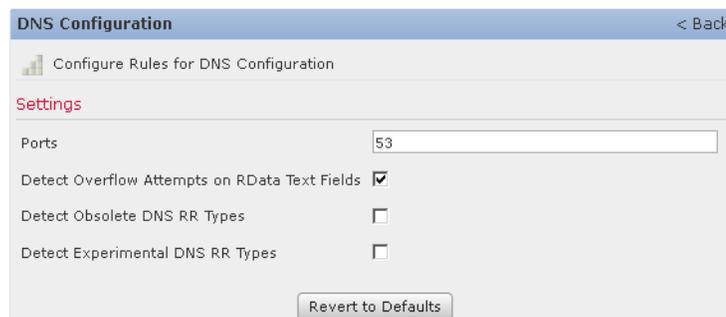
1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.



The screenshot shows the 'DNS Configuration' page with a '< Back' button in the top right. Below the title bar is a section 'Configure Rules for DNS Configuration'. Underneath is a 'Settings' section with the following fields and checkboxes:

Ports	<input type="text" value="53"/>
Detect Overflow Attempts on RData Text Fields	<input checked="" type="checkbox"/>
Detect Obsolete DNS RR Types	<input type="checkbox"/>
Detect Experimental DNS RR Types	<input type="checkbox"/>

At the bottom of the settings area is a 'Revert to Defaults' button.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. 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.
6. Optionally, click **Configure Rules for DNS Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the DNS Configuration page.

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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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.

Note the following when using the FTP/Telnet decoder:

- The FTP/Telnet decoder requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- 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. A link on the configuration page takes you to a filtered view of FTP and telnet preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

For more information, see the following topics:

- [Understanding Global FTP and Telnet Options](#) on page 832
- [Configuring Global FTP/Telnet Options](#) on page 833
- [Understanding Telnet Options](#) on page 835
- [Configuring Telnet Options](#) on page 836
- [Understanding Server-Level FTP Options](#) on page 838
- [Configuring Server-Level FTP Options](#) on page 842
- [Understanding Client-Level FTP Options](#) on page 846
- [Configuring Client-Level FTP Options](#) on page 847

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, 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](#) on page 743 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](#) on page 832.

To configure global options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

TIP! For more information on configuring the other options on this page, see [Configuring Telnet Options](#) on page 836, [Configuring Server-Level FTP Options](#) on page 842, and [Configuring Client-Level FTP Options](#) on page 847.



5. 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.

WARNING! If you disable **TCP Stream Configuration** in an intrusion policy (not recommended), FTP and telnet processing becomes implicitly stateless even if you select **Stateful Inspection** here, because the TCP layer does not pass on any state information. You can determine whether TCP Stream Configuration is enabled by expanding Advanced Settings on the left side of the page; TCP Stream Configuration is enabled if it appears as a sublink beneath Advanced Settings. For more information on stateful inspection and stream reassembly settings, see [Using TCP Stream Preprocessing](#) on page 940 and [Reassembling TCP Streams](#) on page 949.

- 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.
6. Optionally, click **Configure Rules for FTP and Telnet Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the FTP and Telnet Configuration page.
 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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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, 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.

IMPORTANT! Any port you add to the telnet **Ports** list should also be added in each TCP policy to the appropriate list of TCP reassembly ports, depending on whether you are monitoring client or server traffic, or both. Note, however, that reassembling additional traffic types (client, server, both) increases resource demands. For more information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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](#) on page 743 for more information.

Are You There Attack Threshold Number

Detects when the number of consecutive AYT commands exceeds the specified threshold. Sourcefire 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](#) on page 743 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](#) on page 835.

To configure telnet options:

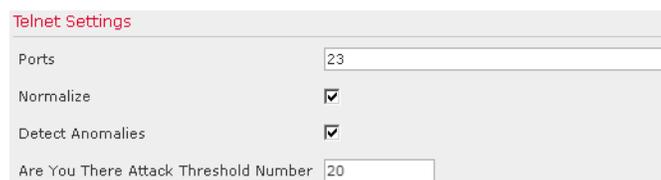
ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

TIP! For more information on configuring the other options on this page, see [Configuring Global FTP/Telnet Options](#) on page 833, [Configuring Server-Level FTP Options](#) on page 842, and [Configuring Client-Level FTP Options](#) on page 847.



Telnet Settings	
Ports	<input type="text" value="23"/>
Normalize	<input checked="" type="checkbox"/>
Detect Anomalies	<input checked="" type="checkbox"/>
Are You There Attack Threshold Number	<input type="text" value="20"/>

5. 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.

Add the same list of ports indicated here to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see [Reassembling TCP Streams](#) on page 949.

WARNING! 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! Sourcefire recommends that you set the AYT threshold to a value no higher than the default value.

6. Optionally, click **Configure Rules for FTP and Telnet Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the FTP and Telnet Configuration page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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, 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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).

Ports

Use this option to specify the ports on the FTP server where the managed device should monitor traffic. In the interface, list multiple ports separated by commas.

IMPORTANT! Any port you add to the server-level FTP **Ports** list should also be added in each TCP policy to the appropriate list of TCP reassembly ports, depending on whether you are monitoring client or server traffic, or both. Note, however, that reassembling additional traffic types (client, server, both) increases resource demands. For more information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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 Sourcefire 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 Sourcefire 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](#) on page 743 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](#) on page 743 for more information.

Command Validity

Use this option to enter a valid format for a specific command. See [Creating FTP Command Parameter Validation Statements](#) on page 841 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](#) on page 743 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](#) on page 743 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.

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](#) on page 838 for more information.

Any of the parameters listed in the following table can be used in FTP command parameter validation statements.

FTP Command Parameters

IF YOU USE...	THE FOLLOWING VALIDATION OCCURS...
<code>int</code>	The represented parameter must be an integer.
<code>number</code>	The represented parameter must be an integer between 1 and 255.
<code>char _chars</code>	<p>The represented parameter must be a single character and a member of the characters specified in the <code>_chars</code> argument.</p> <p>For example, defining the command validity for <code>MODE</code> with the validation statement <code>char SBC</code> checks that the parameter for the <code>MODE</code> command comprises the character <code>S</code> (representing Stream mode), the character <code>B</code> (representing Block mode), or the character <code>C</code> (representing Compressed mode).</p>
<code>date _datefmt</code>	<p>If <code>_datefmt</code> contains <code>#</code>, the represented parameter must be a number.</p> <p>If <code>_datefmt</code> contains <code>C</code>, the represented parameter must be a character.</p> <p>If <code>_datefmt</code> contains literal strings, the represented parameter must match the literal string.</p>
<code>string</code>	The represented parameter must be a string.
<code>host_port</code>	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.

IMPORTANT! 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](#) on page 838.

To configure server-level FTP options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

TIP! For more information on configuring the other options on this page, see [Configuring Global FTP/Telnet Options](#) on page 833, [Configuring Telnet Options](#) on page 836, and [Configuring Client-Level FTP Options](#) on page 847.



5. You have two options:

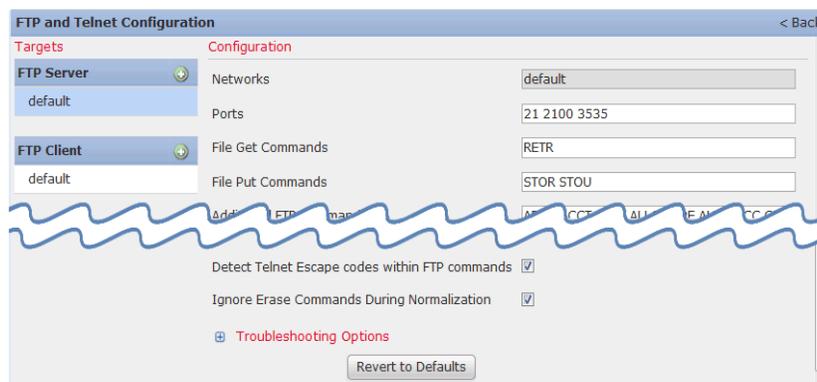
- Add a new server profile. Click the add icon (⊕) next to **FTP Server** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**.

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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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.



6. 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.

IMPORTANT! Add the same list of ports indicated here to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see [Reassembling TCP Streams](#) on page 949.

- 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.

IMPORTANT! Do not change the values in the **File Get Commands** and **File Put Commands** field unless directed to do so by Sourcefire 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.

IMPORTANT! 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](#) on page 838.

- To improve performance on FTP data transfers by disabling all inspection other than state inspection on the data transfer channel, enable **Ignore FTP Transfers**.

IMPORTANT! 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](#) on page 832.

- 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**.
7. Optionally, click **Configure Rules for FTP and Telnet Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the FTP and Telnet Configuration page.
 8. Optionally, modify the related troubleshooting option only if asked to do so by Sourcefire Support; click the **+** sign next to **Troubleshooting Options** to expand the troubleshooting options section. See [Understanding Troubleshooting Options](#) on page 790 for more information.
 9. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Understanding Client-Level FTP Options

LICENSE: Protection

You can create profiles for FTP clients. Within each profile, you can specify the maximum response length for an FTP response from a client. You can also configure whether the decoder detects bounce attacks and use of telnet commands on the FTP command channel for a particular client.

If no preprocessor rule is mentioned, 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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).

Max Response Length

Use this option to specify the maximum length of a response string from the FTP client.

You can enable rule 125:6 to generate events for this option. See [Setting Rule States](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 846. For more information on telnet options, see [Understanding Telnet Options](#) on page 835. For more information on additional FTP options, see [Understanding Server-Level FTP Options](#) on page 838 and [Understanding Global FTP and Telnet Options](#) on

page 832.

To configure client-level FTP options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. 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.



5. You have two options:

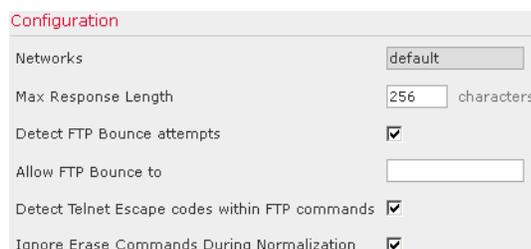
- Add a new client profile. Click the add icon () next to **FTP Client** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Client Address** field and click **OK**.

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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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.



Configuration	
Networks	default
Max Response Length	256 characters
Detect FTP Bounce attempts	<input checked="" type="checkbox"/>
Allow FTP Bounce to	<input type="text"/>
Detect Telnet Escape codes within FTP commands	<input checked="" type="checkbox"/>
Ignore Erase Commands During Normalization	<input checked="" type="checkbox"/>

6. 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

IMPORTANT! 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**.
7. Optionally, click **Configure Rules for FTP and Telnet Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the FTP and Telnet Configuration page.
 8. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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. The HTTP Inspect preprocessor requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- You must enable HTTP preprocessor rules, which have a generator ID (GID) of 119, if you want these rules to generate events. A link on the configuration page takes you to a filtered view of HTTP preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.
- When a rule that requires this preprocessor is enabled in an intrusion policy where the preprocessor is disabled, you must enable the preprocessor or choose to allow the system to enable it automatically before you can save the policy. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

See the following sections for more information:

- [Selecting Global HTTP Normalization Options](#) on page 851
- [Configuring Global HTTP Configuration Options](#) on page 853
- [Selecting Server-Level HTTP Normalization Options](#) on page 854
- [Selecting Server-Level HTTP Normalization Encoding Options](#) on page 862
- [Configuring HTTP Server Options](#) on page 865
- [Enabling Additional HTTP Inspect Preprocessor Rules](#) on page 868

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](#) on page 854 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. See [Setting the Default Action](#) on page 428, and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511 for more information.

If no preprocessor rule is mentioned, 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.

IMPORTANT! If you turn this option on, make sure to list all ports that do receive HTTP traffic in a server profile on the HTTP Configuration page. If you do not, and you have enabled 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](#) on page 743 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](#) on page 743 for more information.

Maximum Compressed Data Depth

Sets the maximum size of compressed data to decompress when **Inspect Compressed Data** 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** is enabled. You can specify from 1 to 65535 bytes.

Configuring Global HTTP Configuration Options

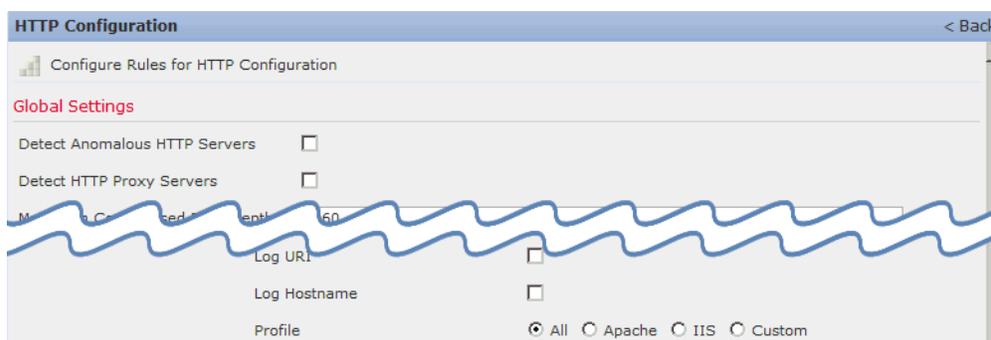
LICENSE: Protection

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](#) on page 851.

To configure global HTTP configuration options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.



5. You can modify any of the global options described in [Selecting Global HTTP Normalization Options](#) on page 851.
6. Optionally, click **Configure Rules for HTTP Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the HTTP Configuration page.

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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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, the option is not associated with a preprocessor rule.

Networks

Use this option to specify the IP address of one or more servers.

Note that 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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).

Ports

The ports whose HTTP traffic the preprocessor engine normalizes. Separate multiple port numbers with commas.

IMPORTANT! Any port you add to the HTTP **Ports** list should also be added in each TCP policy to the appropriate list of TCP reassembly ports, depending on whether you are monitoring client or server traffic, or both. Note, however, that reassembling additional traffic types (client, server, both) increases resource demands. For more information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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](#) on page 743 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](#) on page 1069 for more information.

You can specify a value from -1 to 1460. Sourcefire 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](#) on page 1069 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. Sourcefire 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](#) on page 743 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](#) on page 743 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](#) on page 743 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** keyword **HTTP Client Body** option. See [HTTP Content Options](#) on page 1069 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](#) on page 743 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** keyword with its **HTTP Method** argument to search for content in HTTP methods. See [HTTP Content Options](#). 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.

IMPORTANT! 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](#) on page 1069, [Generating Events on the HTTP Encoding Type and Location](#) on page 1175, and [Pointing to a Specific Payload Type](#) on page 1177.

You can enable rules 120:2 and 120:3 to generate events for this option. See [Setting Rule States](#) on page 743 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](#) on page 743 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](#) on page 1177 for more information.

Unlimited Decompression

When **Inspect Compressed Data** 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](#) on page 851.

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](#) on page 1177 for more information.

You can enable rules 120:9, 120:10, and 120:11 to generate events for this option, as follows:

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](#) on page 743 for more information.

Extract Original Client IP Address

Enables extraction of the original client IP address from the X-Forwarded-For (XFF) or True-Client-IP HTTP header. You can display the extracted original client IP address in the intrusion events table view. See [Understanding Intrusion Events](#) on page 604 for more information.

You can enable rules 119:23, 119:29 and 119:30 to generate events for this option. See [Setting Rule States](#) on page 743 for more information.

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 [Understanding Intrusion Events](#) on page 604 and [Viewing Event Information](#) on page 625 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 [Understanding Intrusion Events](#) on page 604 and [Viewing Event Information](#) on page 625 for more information.

You can enable rule 119:25 to generate events for this option. See [Setting Rule States](#) on page 743 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](#) on page 868 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](#) on page 862 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.

Note that the base36 encoding type has been deprecated. For backward compatibility, the base36 option is allowed in existing intrusion policies, but it does not cause the system to detect base36 traffic.

If no preprocessor rule is mentioned, 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](#) on page 743 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](#) on page 743 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 Sourcefire recommends decoding HTTP traffic encoded with %u encodings.

You can enable rule 119:3 to generate events for this option. See [Setting Rule States](#) on page 743 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. Sourcefire 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](#) on page 743 for more information.

Microsoft IIS Encoding

Decodes using Unicode codepoint mapping.

TIP! Sourcefire 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](#) on page 743 for more information.

Double Encoding

Decodes IIS double encoded traffic by making two passes through the request URI performing decodes in each one. Sourcefire 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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.

IMPORTANT! 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](#) on page 743 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](#) on page 743 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](#) on page 743 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.

IMPORTANT! 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](#) on page 743 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](#) on page 743 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](#) on page 854 and [Selecting Server-Level HTTP Normalization Encoding Options](#) on page 862.

To configure server-level HTTP configuration options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You have two options:

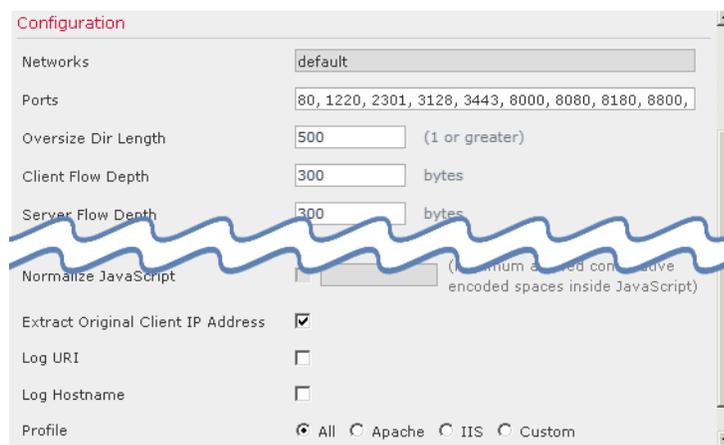
- Add a new server profile. Click the add icon () next to **Servers** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**.

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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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.



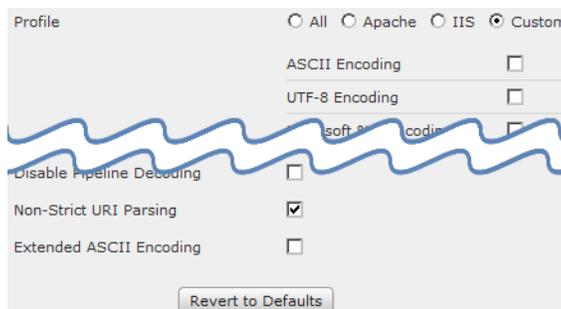
6. Optionally, modify the address or addresses listed in the **Networks** field and click any other area of the page.

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 servers on your network that are not identified in another profile.

7. In the **Ports** field, list the ports whose traffic you want to inspect with HTTP Inspect. Separate multiple ports with commas.

8. You can modify any of the other options described in [Selecting Server-Level HTTP Normalization Options](#) on page 854.
9. Select a server profile as follows:
 - Select **Custom** to create your own server profile (see [Selecting Server-Level HTTP Normalization Encoding Options](#) on page 862 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.
10. If you selected **Custom**, the custom options appear.



11. Configure the HTTP decoding options you want in your profile.
See [Selecting Server-Level HTTP Normalization Options](#) on page 854 for details on available normalization options.
12. Optionally, click **Configure Rules for HTTP Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the HTTP Configuration page.
13. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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 [Additional HTTP Inspect Preprocessor Rules](#) table to generate events for HTTP Inspect

preprocessor rules that are not associated with specific configuration options. See [Setting Rule States](#) on page 743 for more information.

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 <code>content-length</code> 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 on page 940.
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 `admin` runs. Some UNIX hosts use RPC `admin` 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.

IMPORTANT! Any port you add to the RPC **Ports** list should also be added in each TCP policy to the appropriate list of TCP reassembly ports, depending on whether you are monitoring client or server traffic, or both. Note, however, that reassembling additional traffic types (client, server, both) increases resource demands. For more information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 869.

To configure the Sun RPC preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.

Settings	
Ports	111, 32770, 32771, 32772, 32773
Detect fragmented RPC records	<input type="checkbox"/>
Detect multiple records in one packet	<input type="checkbox"/>
Detect fragmented record sums which exceed one packet	<input type="checkbox"/>
Detect single fragment records which exceed the size of one packet	<input type="checkbox"/>

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. In the **Ports** field, type the port numbers where you want to decode RPC traffic. Separate multiple ports with commas.
6. 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**

7. Optionally, click **Configure Rules for Sun RPC Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the Sun RPC Configuration page.
8. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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. UDP session tracking must be enabled before you can save a policy with the SIP preprocessor enabled. See [Using UDP Stream Preprocessing](#) on page 956

and [Automatically Enabling Advanced Settings](#) on page 787 for more information.

- 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](#) on page 1125.
- 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. A link on the configuration page takes you to a filtered view of SIP preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.
- When a shared object rule or standard text rule that requires this preprocessor is enabled in an intrusion policy where the preprocessor is disabled, you must enable the preprocessor or choose to allow the system to enable it automatically before you can save the policy. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

See the following sections for more information:

- [Selecting SIP Preprocessor Options](#) on page 873
- [Configuring the SIP Preprocessor](#) on page 875
- [Enabling Additional SIP Preprocessor Rules](#) on page 876

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, 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](#) on page 1126 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

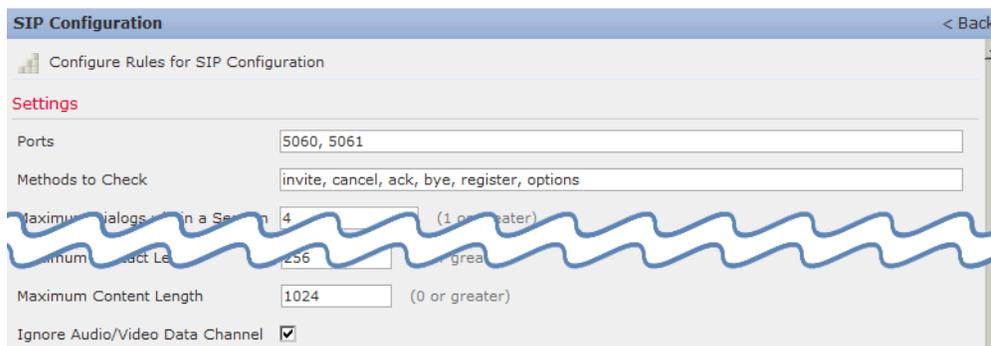
Use the following procedure to configure the SIP preprocessor.

To configure the SIP preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can modify any of the options described in [Selecting SIP Preprocessor Options](#) on page 873.
6. Optionally, click **Configure Rules for SIP Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the SIP Configuration page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on

page 743 for information on enabling rules.

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.
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.

Additional SIP Preprocessor Rules (Continued)

PREPROCESSOR RULE GID:SID	DESCRIPTION
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 signalling 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.

Note the following information regarding the use of the GTP preprocessor:

- The GTP preprocessor requires UDP stream configuration. When you enable the GTP preprocessor and UDP stream configuration is disabled, you are prompted whether to enable UDP stream configuration when you save the policy.
- Both the GTP command channel configuration and UDP stream configuration advanced settings must be enabled to allow processing of rules using GTP keywords. When either is disabled and you enable rules that use GTP keywords, you are prompted whether to enable the advanced setting when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

You must enable the GTP preprocessor rules in the following table if you want them to generate events. See [Setting Rule States](#) on page 743 for information on enabling rules.

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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.



5. 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.
6. Optionally, click **Configure Rules for GTP Command Channel Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the GTP Command Channel Configuration page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 1177 for more information.

Extraction and decoding include multiple attachments, when present, and large attachments that span multiple packets.

Note the following when using the IMAP preprocessor:

- Because IMAP traffic is carried over TCP/IP connections, the IMAP preprocessor requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the IMAP preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- If you want IMAP preprocessor rules to generate events, you must enable the rules. IMAP preprocessor rules have a generator ID (GID) of 141. A link on the configuration page takes you to a filtered view of IMAP preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

See the following sections for more information:

- [Selecting IMAP Preprocessor Options](#) on page 881
- [Configuring the IMAP Preprocessor](#) on page 882
- [Enabling Additional IMAP Preprocessor Rules](#) on page 884

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 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. See [Setting the Default Action](#) on page 428, and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511 for more information.

If no preprocessor rule is mentioned, 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.

IMPORTANT! Any port you add to the IMAP port list should also be added to the TCP client reassembly list for each TCP policy. For information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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](#) on page 881.

To configure the IMAP preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

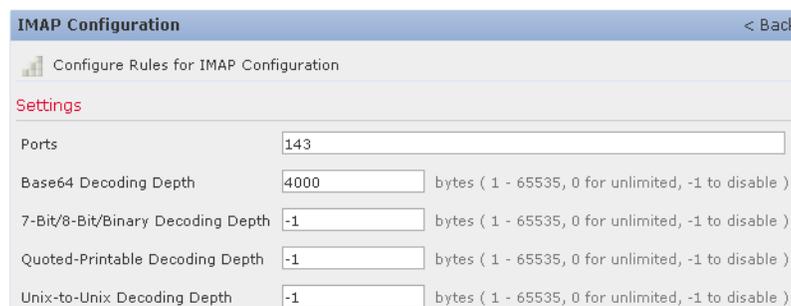
2. Click the edit icon () next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.



The screenshot shows the 'IMAP Configuration' page with a '< Back' button in the top right. Below the title is a sub-header 'Configure Rules for IMAP Configuration'. Underneath is a 'Settings' section with a table of configuration options:

Setting	Value	Notes
Ports	143	
Base64 Decoding Depth	4000	bytes (1 - 65535, 0 for unlimited, -1 to disable)
7-Bit/8-Bit/Binary Decoding Depth	-1	bytes (1 - 65535, 0 for unlimited, -1 to disable)
Quoted-Printable Decoding Depth	-1	bytes (1 - 65535, 0 for unlimited, -1 to disable)
Unix-to-Unix Decoding Depth	-1	bytes (1 - 65535, 0 for unlimited, -1 to disable)

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Specify the **Ports** where IMAP traffic should be decoded. Separate multiple port numbers with commas.

IMPORTANT! Any port you add to the IMAP port list should also be added to the TCP client reassembly list for each TCP policy. For information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

6. 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](#) on page 1177 for more information.

7. Optionally, click **Configure Rules for IMAP Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the IMAP Configuration page.
8. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 743 for information on enabling rules.

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](#) on page 1177 for more information.

Extraction and decoding include multiple attachments, when present, and large attachments that span multiple packets.

Note the following when using the POP preprocessor:

- Because POP traffic is carried over TCP/IP connections, the POP preprocessor requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the POP preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- If you want POP preprocessor rules to generate events, you must enable the rules. POP preprocessor rules have a generator ID (GID) of 142. A link on the configuration page takes you to a filtered view of POP preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

See the following sections for more information:

- [Selecting POP Preprocessor Options](#) on page 885
- [Configuring the POP Preprocessor](#) on page 887
- [Enabling Additional POP Preprocessor Rules](#) on page 889

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. See [Setting the Default Action](#) on page 428, and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511 for more information.

If no preprocessor rule is mentioned, 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.

IMPORTANT! Any port you add to the POP port list should also be added to the TCP client reassembly list for each TCP policy. For information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 885.

To configure the POP preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
2. Click the edit icon () next to the policy you want to edit.
If you have unsaved advanced editor changes in another policy, click **OK** to discard those changes and continue. See [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

- 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.

The screenshot shows the 'POP Configuration' page with a '< Back' link in the top right. Below the title is a 'Configure Rules for POP Configuration' section. Underneath is a 'Settings' section with five rows of configuration options, each with a text input field and a help text: 'Ports' (110), 'Base64 Decoding Depth' (4000), '7-Bit/8-Bit/Binary Decoding Depth' (-1), 'Quoted-Printable Decoding Depth' (-1), and 'Unix-to-Unix Decoding Depth' (-1). The help text for all depth settings is 'bytes (1 - 65535, 0 for unlimited, -1 to disable)'.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

- Specify the **Ports** where IMAP traffic should be decoded. Separate multiple port numbers with commas.

IMPORTANT! Any port you add to the POP port list should also be added to the TCP client reassembly list for each TCP policy. For information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

- 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](#) on page 1177 for more information.

- Optionally, click **Configure Rules for POP Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the POP Configuration page.

8. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 743 for information on enabling rules.

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:

- The SMTP preprocessor requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the SMTP preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- You must enable SMTP preprocessor rules, which have a generator ID (GID) of 124, if you want these rules to generate events. A link on the configuration page takes you to a filtered view of SMTP preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

For more information, see the following sections:

- [Understanding SMTP Decoding](#) on page 890
- [Configuring SMTP Decoding](#) on page 895
- [Enabling SMTP Maximum Decoding Memory Alerting](#) on page 899

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. See [Setting the Default Action](#) on page 428, and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511 for more information.

If no preprocessor rule is mentioned, 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.

IMPORTANT! Any port you add to the SMTP **Ports** list should also be added to the TCP client reassembly list for each TCP policy. For more information on configuring TCP reassembly ports, see [Selecting Stream Reassembly Options](#) on page 950.

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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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

IMPORTANT! 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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 [Understanding Intrusion Events](#) on page 604 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 [Understanding Intrusion Events](#) on page 604 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 [Understanding Intrusion Events](#) on page 604 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** 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 [Constraining Content Matches](#) on page 1065 and [Using the Packet View](#) on page 622 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

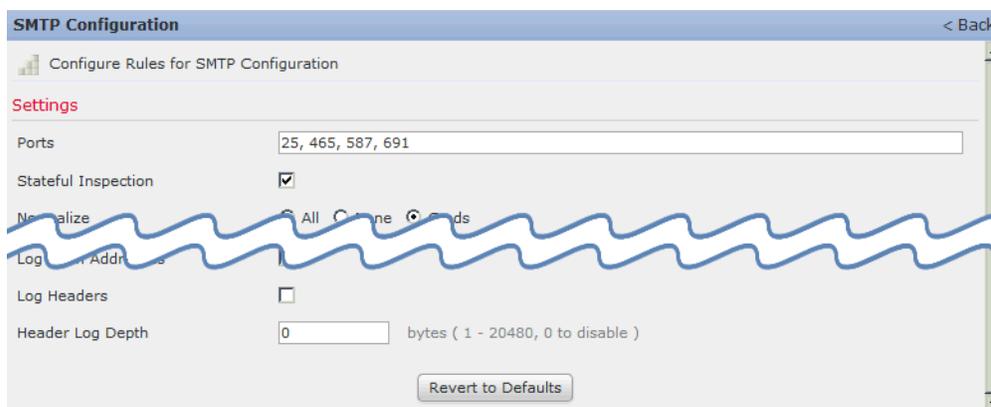
You can use the SMTP Configuration page of an intrusion policy to configure SMTP normalization. For more information on SMTP preprocessor configuration options, see [Understanding SMTP Decoding](#) on page 890.

To configure SMTP decoding options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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. The following graphic shows the Defense Center packet view.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Specify the **Ports** where SMTP traffic should be decoded, separated by commas.
6. Select **Stateful Inspection** to examine reassembled TCP streams containing SMTP packets. Clear **Stateful Inspection** to inspect only unreassembled SMTP packets.

7. 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**.
8. Specify a maximum command line length in the **Max Command Line Len** field.
9. Specify a maximum data header line length in the **Max Header Line Len** field.
10. Specify a maximum response line length in the **Max Response Line Len** field.

IMPORTANT! 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.

11. 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.
12. Specify any commands that you want to treat as invalid and detect in the **Invalid Commands** field. Separate commands with spaces.
13. Specify any commands that you want to treat as valid in the **Valid Commands** field. Separate commands with spaces.

IMPORTANT! 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, QUEUE, STARTTLS, TICK, TIME, TURN, TURNME, VERB, VRFY, X-EXPS, X-LINK2STATE, XADR, XAUTH, XCIR, XEXCH50, XGEN, XLICENSE, XQUE, XSTA, XTRN, or XUSR.

14. 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.
15. 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.
16. Specify any commands that initiate an authentication exchange between client and server in the **Authentication Commands** field. Separate commands with spaces.
17. To detect packets that are part of X-Link2State Microsoft Exchange buffer data overflow attacks, select **Detect xlink2state**.
18. 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](#) on page 1177 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.

19. 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** keyword with email header data as a pattern. See [Viewing Event Information](#) on page 625 and [Searching for Content Matches](#) on page 1063 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**.

20. Optionally, click **Configure Rules for SMTP Configuration** at the top of the page to display rules associated with individual options.

Click **Back** to return to the SMTP Configuration page.

21. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 743 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.

Both 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. A link on the configuration page takes you to a filtered view of SSH preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.
- The SSH preprocessor requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the SSH preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- The SSH preprocessor does not handle brute force attacks. For information on brute force attempts, see [Adding Dynamic Rule States](#) on page 756.

See the following sections for more information:

- [Selecting SSH Preprocessor Options](#) on page 901
- [Configuring the SSH Preprocessor](#) on page 903

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, 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 for more information.

Configuring the SSH Preprocessor

LICENSE: Protection

This section explains how to configure the SSH preprocessor.

To configure the SSH preprocessor:

ACCESS: Admin/Intrusion Admin

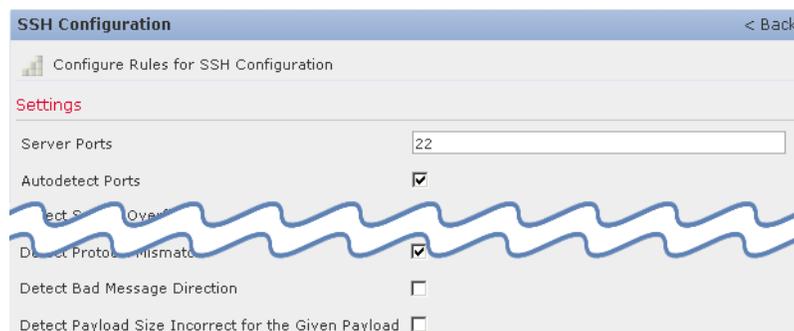
1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can modify any of the options on the SSH Configuration preprocessor page. See [Selecting SSH Preprocessor Options](#) on page 901 for more information.
6. Optionally, click **Configure Rules for SSH Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the SSH Configuration page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Using the SSL Preprocessor

LICENSE: Protection

Although the system cannot analyze the contents of encrypted traffic, an SSL preprocessor option can be set to continue to attempt to inspect the traffic, occasionally generating false positives and wasting detection resources. 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.

Note the following when using the SSL preprocessor:

- The SSL preprocessor requires TCP stream preprocessing. If TCP stream preprocessing is disabled and you enable the SSL preprocessor, you are prompted when you save the policy whether to enable TCP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 and [Using TCP Stream Preprocessing](#) on page 940 for more information.
- When an intrusion rule that requires this preprocessor is enabled in an intrusion policy where the preprocessor is disabled, you must enable the preprocessor or choose to allow the system to enable it automatically before you can save the policy. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

For more information, see the following sections:

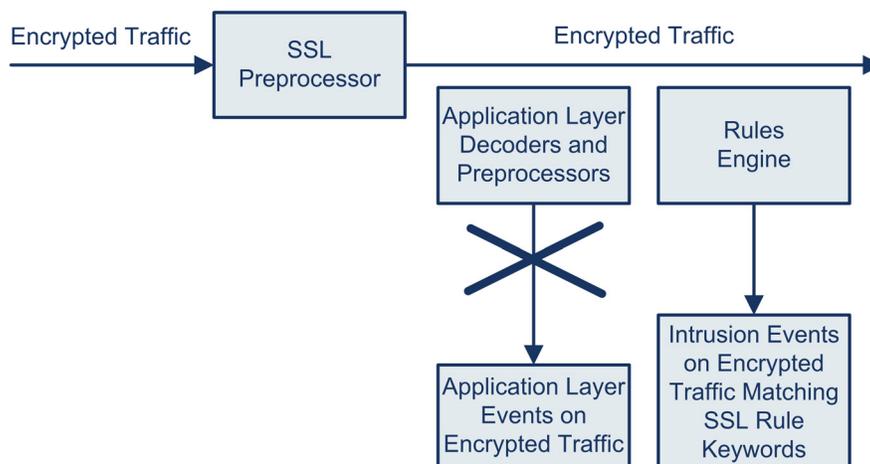
- [Understanding SSL Preprocessing](#) on page 905
- [Enabling SSL Preprocessor Rules](#) on page 907
- [Configuring the SSL Preprocessor](#) on page 907

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.

IMPORTANT! 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](#) on page 1113. Note that the SSL preprocessor must be enabled to allow processing of rules that contain SSL keywords.

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. A link on the configuration page takes you to a filtered view of SSL preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

The [SSL Preprocessor Rules](#) table describes the SSL preprocessor rules you can enable.

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 on page 907 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.

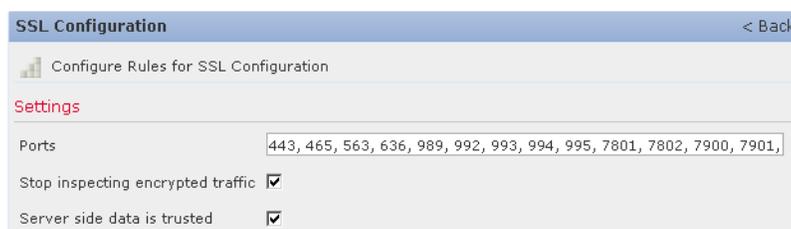
IMPORTANT! If the SSL preprocessor detects non-SSL traffic over the ports specified for SSL monitoring, it tries to decode the traffic as SSL traffic, and then flags it as corrupt.

To configure the SSL preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. 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.

6. 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.
7. Click the **Server side data is trusted** check box to enable or disable identification of encrypted traffic based only on the client-side traffic.
8. Optionally, click **Configure Rules for SSL Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the SSH Configuration page.
9. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Working with SCADA Preprocessors

LICENSE: Protection

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. The Sourcefire 3D System provides preprocessors for the Modbus and DNP3 SCADA protocols.

See the following sections for more information:

- [Configuring the Modbus Preprocessor](#) on page 909
- [Configuring the DNP3 Preprocessor](#) on page 911

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](#) on page 1145 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 predecessor rules in the following table if you want these rules to generate events. See [Setting Rule States](#) on page 743 for information on enabling rules.

Modbus Preprocessor Rules

PREPROCESSOR RULE GID:SID	DESCRIPTION
144:1	<p>Generates an event when the length in the Modbus header does not match the length required by the Modbus function code.</p> <p>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.</p>
144:2	<p>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 predecessor does not process these other protocols, this event is generated instead.</p>
144:3	<p>Generates an event when the predecessor detects a reserved Modbus function code.</p>

Note the following information regarding the use of the Modbus predecessor:

- If your network does not contain any Modbus-enabled devices, you should not enable this predecessor in an intrusion policy that you apply to traffic.
- The Modbus predecessor requires TCP stream configuration. When you enable the Modbus predecessor and TCP stream configuration is disabled, you are prompted whether to enable the advanced setting when you save the policy.
See [Configuring TCP Stream Preprocessing](#) on page 952 and [Automatically Enabling Advanced Settings](#) on page 787 for more information.
- Both TCP stream configuration and the Modbus predecessor must be enabled to allow processing of rules using Modbus keywords. When either is disabled and you enable rules that use Modbus keywords, you are prompted whether to enable the disabled advanced setting when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

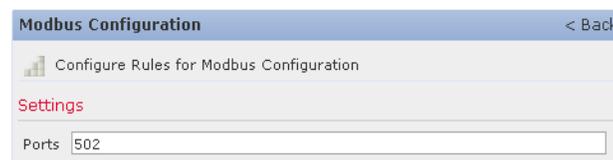
You can use the following procedure to modify the ports the Modbus predecessor monitors.

To configure the Modbus preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.



5. 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.
6. Optionally, click **Configure Rules for Modbus Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the Modbus Configuration page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 1148 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](#) on page 743 for information on enabling rules.

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 the following information regarding the use of the DNP3 preprocessor:

- If your network does not contain any DNP3-enabled devices, you should not enable this preprocessor in an intrusion policy that you apply to traffic.
- The DNP3 preprocessor requires TCP stream configuration. When you enable the DNP3 preprocessor and TCP stream configuration is disabled, you are prompted whether to enable the advanced setting when you save the policy.
See [Configuring TCP Stream Preprocessing](#) on page 952 and [Automatically Enabling Advanced Settings](#) on page 787 for more information.
- Both TCP stream configuration and the DNP3 preprocessor must be enabled to allow processing of rules using DNP3 keywords. When either is disabled and you enable rules that use DNP3 keywords, you are prompted whether to enable the disabled advanced setting when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787 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.

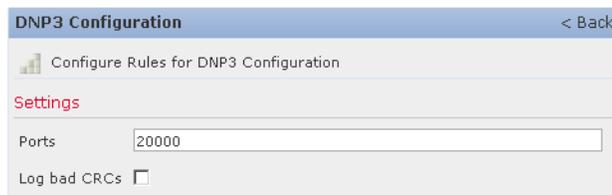
To configure the DNP3 preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. 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.



DNP3 Configuration < Back

Configure Rules for DNP3 Configuration

Settings

Ports

Log bad CRCs

5. Optionally, modify the **Ports** that the predecessor inspects for DNP3 traffic. You can specify an integer from 0 to 65535. Use commas to separate multiple ports.
6. 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.
7. Optionally, click **Configure Rules for DNP3 Configuration** at the top of the page to display rules associated with individual options.
Click **Back** to return to the DNP3 Configuration page.
8. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 24

USING TRANSPORT & NETWORK LAYER PREPROCESSORS

Sourcefire provides preprocessors that detect exploits at the network and transport layers. These 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 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.

See the following sections for more information:

- [Verifying Checksums](#) on page 915
- [Ignoring VLAN Headers](#) on page 917
- [Normalizing Inline Traffic](#) on page 918
- [Defragmenting IP Packets](#) on page 927
- [Understanding Packet Decoding](#) on page 934
- [Using TCP Stream Preprocessing](#) on page 940
- [Using UDP Stream Preprocessing](#) on page 956

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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

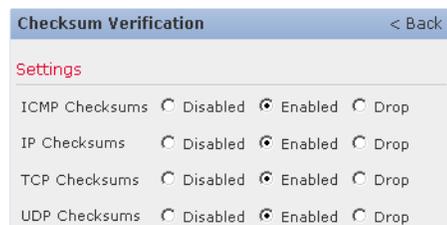
3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can set any of the options in the Checksum Verification section to **Enable** or **Disable** in a passive or inline deployment, or to **Drop** in an inline deployment:

- **ICMP Checksums**
- **IP Checksums**

- **TCP Checksums**
- **UDP Checksums**

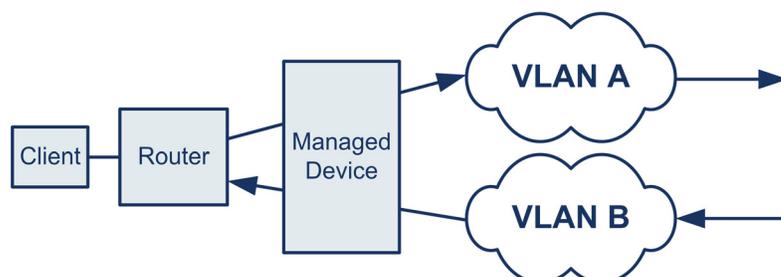
Note that to drop offending packets you must also enable **Drop when Inline** in addition to setting an option to **Drop** in the policy. See [Setting Drop Behavior in an Inline Deployment](#) on page 689 for more information. Note also that setting these options to **Drop** in a passive deployment is the same as setting them to **Enable**.

6. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Ignoring VLAN Headers

LICENSE: Protection

Different VLAN tags in traffic traveling in different directions for the same connection can affect traffic reassembly and rule processing. For example, in the following graphic traffic for the same connection could be transmitted over VLAN A and received over VLAN B.



When you enable the **Ignore VLAN Header** detection setting, the system ignores the VLAN header so packets can be correctly processed for your deployment.

To ignore VLAN headers:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. You have two choices:
 - If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Detection Settings page appears.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You have the following choices:
 - For deployed devices that might detect different VLAN tags for the same connection in traffic traveling in different directions, select the **Ignore VLAN Header** check box to ignore VLAN headers when identifying traffic.
 - For deployed devices that will not detect different VLAN tags for the same connection traffic traveling in different directions, clear the **Ignore VLAN Header** check box to include VLAN headers when identifying traffic.
6. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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. When you apply an intrusion policy as part of an access control policy and the inline normalization preprocessor is enabled, the system tests the following two conditions to ensure that you are using an inline deployment:

- **Drop when Inline** is enabled
- The policy is applied to a device using an inline set

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, so you must ensure that the TCP stream preprocessor is enabled when you enable normalization of TCP traffic.

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.

Note that 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, Sourcefire recommends you configure the inline normalization preprocessor, with the Normalize TCP and Normalize TCP Payload options enabled. In a passive deployment, Sourcefire recommends you configure adaptive profiles. For more information, see [Using Adaptive Profiles](#) on page 1003.

See the following sections for more information:

- [Understanding Protocol Normalization](#) on page 919
- [Configuring Inline Normalization](#) on page 922

Understanding Protocol Normalization

LICENSE: Protection

Normalization of each protocol includes one or more *base* normalizations, which occur automatically when you enable normalization of the protocol. Some protocols also have optional normalizations.

See [Configuring Inline Normalization](#) on page 922 for information on configuring normalization of traffic for different protocols. The following sections list the base normalizations and any optional normalizations for each protocol type:

- [IPv4 Normalization](#) on page 920
- [IPv6 Normalization](#) on page 920
- [ICMPv4 and ICMPv6 Normalization](#) on page 920
- [TCP Normalization](#) on page 921

IPv4 Normalization

LICENSE: Protection

When you enable **Normalize IPv4**, the system performs the following base 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)

In addition, the system performs the following optional normalizations when you enable IPv4 normalization and select the corresponding option:

- enabling the **Normalize Don't Fragment Bit** option clears the single-bit Don't Fragment subfield of the IPv4 Flags header field
- enabling the **Normalize Reserved Bit** option clears the single-bit Reserved subfield of the IPv4 Flags header field
- enabling the **Normalize TOS Bit** option clears the one byte Differentiated Services header field, formerly known as Type of Service (ToS)
- enabling the **Normalize Excess Payload** option trims excess payload to the datagram length specified in the IP header plus the Layer 2 header
- enabling the **Reset TTL** and **Minimize TTL** options sets the Time to Live (TTL) field as needed to a specified minimum value

See [Configuring Inline Normalization](#) on page 922 for more information.

IPv6 Normalization

LICENSE: Protection

When you enable **Normalize IPv6**, the system sets all Option Type fields in the Hop-by-Hop Options and Destination Options extension headers to 00 (Skip and continue processing).

Optionally, and as needed, the system also sets the Hop Limit field to a specified minimum value. See the **Reset TTL** and **Minimize TTL** options in [Configuring Inline Normalization](#) on page 922 for more information.

ICMPv4 and ICMPv6 Normalization

LICENSE: Protection

When you enable **Normalize ICMPv4**, **Normalize ICMPv6**, or both, the system clears the 8-bit Code field in Echo (Request) and Echo Reply messages in the corresponding ICMP traffic.

TCP Normalization

LICENSE: Protection

The following sections describe base TCP normalizations, including traffic that is dropped when you enable TCP normalization. It also explains normalizations associated with specific TPC normalization options.

Base TCP Normalizations

When you enable **Normalize TCP**, the system performs the following base normalizations:

- clears the 3-bit Reserved field in the TCP header
- clears the 16-bit Urgent Pointer field if the urgent (URG) control bit is not set
- clears the Urgent Pointer field and the URG control bit if there is no payload
- clears the urgent control bit if the urgent pointer is not set
- clears any option padding bytes
- blocks a subsequent SYN that does not have the same sequence number as the original SYN

Dropped TCP Packets

When you enable **Normalize TCP**, the system drops the following without generating an event:

- retransmitted copies of previously dropped packets
- traffic that attempts to continue a previously dropped session
- any packet that matches any of the following TCP stream preprocessor rules, regardless of whether the rules are enabled:

THESE PREPROCESSOR RULES DROP PACKETS WHEN NORMALIZE TCP IS ENABLED...

129:1, 129:3, 129:4, 129:6, 129:8, 129:11, 129: 14 through 129:19

The Blocked Packets performance graph tracks the number of packets dropped as the result of this options being enabled. See [Generating Intrusion Event Performance Statistics Graphs](#) on page 599 for more information.

Automatically Allowed TCP Options

When you enable **Normalize TCP** and do not specify any for **Allow These TCP Options**, the system performs the following normalizations:

- 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

- drops the packet if Time Stamp is negotiated but not present
- clears the Time Stamp Echo Reply (TSecr) option field if the Acknowledgement (ACK) control bit is not set
- sets the MSS and Window Scale options to No Operation (TCP Option 1) if the Synchronization (SYN) control bit is not set

See [Configuring Inline Normalization](#) on page 922 for more information.

Normalizations Associated with Specific TCP Options

The system performs the following optional normalizations when you enable **Normalize TCP** and select the corresponding option:

- enabling the **Normalize Urgent Pointer** option sets the two-byte Urgent Pointer header field to the payload length if the pointer is greater than the payload length
- enabling the **Normalize TCP Payload** option normalizes the TCP Data field to ensure consistency in retransmitted data and drops any segments that cannot be properly reassembled
- enabling the **Normalize TCP Excess Payload** option removes data in SYN and RST packets, and trims the Data field to the size specified in the Window field, or to the Maximum Segment Size (MSS) if the payload is longer than MSS
- enabling the **Explicit Congestion Notification** option clears ECN flags on a per-packet basis regardless of negotiation, or on a per-stream basis if usage was not negotiated

See [Configuring Inline Normalization](#) on page 922 for more information.

Configuring Inline Normalization

LICENSE: Protection

You can configure the inline normalization preprocessor to normalize IPv4, IPv6, ICMPv4, ICMPv6, and TCP traffic in any combination. In addition to the base normalizations provided when you enable normalization of each traffic type, specific optional normalizations are available for all protocols except ICMP; this includes using the **Reset TTL** option to enable TTL normalization when IPv4 normalization is enabled and IPv6 Hop Limit normalization when IPv6 normalization is enabled.

In addition to enabling and configuring the inline normalization preprocessor, you must also ensure the following or the preprocessor will not normalize traffic:

- your policy must be set to drop traffic in inline deployments; see [Setting Drop Behavior in an Inline Deployment](#) on page 689
- you must apply your policy to an inline set; see [Applying an Access Control Policy](#) on page 466

You must also ensure that the TCP stream preprocessor is enabled when you enable TCP normalization; see [Modifying Advanced Settings](#) on page 774.

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](#) on page 938 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. See [IPv4 Normalization](#) on page 920 for information on specific IPv4 normalizations. 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.

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 TCP

Enables normalization of TCP traffic. See [TCP Normalization](#) on page 921 for information on specific TCP normalizations. When this option is enabled, you can also enable **Normalize Urgent Pointer**, **Normalize TCP Payload**, **Normalize TCP Excess Payload**, and **Exploit Congestion Payload**, and configure **Allow These TCP Options**. You should ensure that the TCP stream preprocessor is enabled when you enable this option; see [Modifying Advanced Settings](#) on page 774.

Normalize Urgent Pointer

Sets the two-byte Urgent Pointer header field to the payload length if the pointer is greater than the payload length. You must enable **Normalize TCP** to select this option.

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. You must enable **Normalize TCP** to select this option.

Normalize TCP Excess Payload

Enables the following normalizations:

- removes data in synchronization (SYN) packets
- removes any data from a reset (RST) packet
- trims the Data field to the size specified in the Window field
- trims the Data field to the Maximum Segment Size (MSS) if the payload is longer than MSS

You must enable **Normalize TCP** to select this option.

Explicit Congestion Notification

Enables per-packet or per-stream normalization of Explicit Congestion Notification (ECN) flags as follows:

- select **Packet** to clear ECN flags regardless of negotiation
- select **Stream** to clear ECN flags if ECN use was not negotiated

You must enable **Normalize TCP** to select this option. 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](#) on page 945 for more information.

Allow These TCP Options

Disables normalization of specific TCP options you allow in traffic. You must enable **Normalize TCP** to select this option.

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 as described in [TCP Normalization](#) on page 921 regardless of the configuration of **Allow These TCP Options**. The system does not 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 Acknowledgement Permitted) and 5 (Selective Acknowledgement). Option keywords are not case sensitive.

You can also specify **any**, which allows all TCP options and effectively disables normalization of all TCP options. See [TCP Normalization](#) on page 921 for additional normalizations performed when you do not specify **any**.

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 Acknowledgement Permitted) and 5 (Selective Acknowledgement)
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)
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

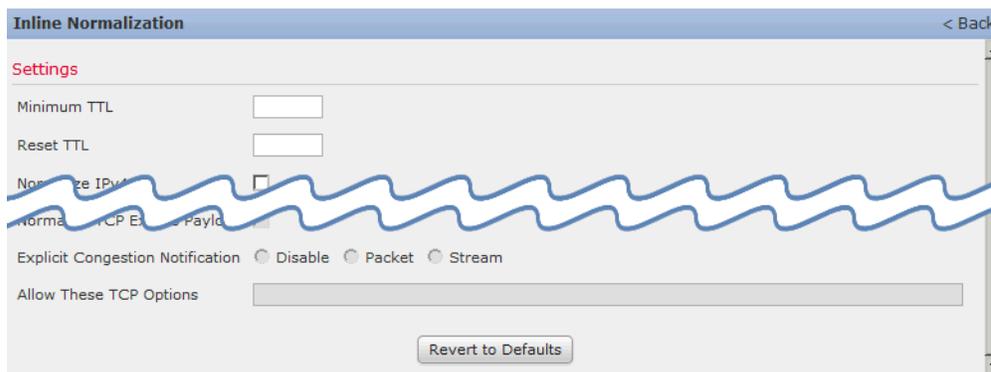
To configure the inline normalizations preprocessor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. You have two choices:
 - 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can set any of the options described in [Configuring Inline Normalization](#) on page 922.
6. 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 [Common Intrusion Policy Editing Actions](#) table on page 675 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

that 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. A link on the configuration page takes you to a filtered view of IP defragmentation preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

See the following sections for more information:

- [Understanding IP Fragmentation Exploits](#) on page 928
- [Target-Based Defragmentation Policies](#) on page 929
- [Selecting Defragmentation Options](#) on page 930
- [Configuring IP Defragmentation](#) on page 932

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](#) on page 929 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 [Using Adaptive Profiles](#) on page 1003.

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 [Target-Based Defragmentation Policies](#) 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.

Target-Based Defragmentation Policies

POLICY	OPERATING SYSTEMS
BSD	AIX FreeBSD IRIX VAX/VMS
BSD-right	HP JetDirect
First	Mac OS HP-UX

Target-Based Defragmentation Policies (Continued)

POLICY	OPERATING SYSTEMS
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, Sourcefire recommends that you specify the behavior of the enabled IP defragmentation preprocessor at a more granular level.

If no preprocessor rule is mentioned, 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.

WARNING! 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 managed device, the memory limit for the device takes precedence.

You can configure the following options for each IP defragmentation policy:

Network

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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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).

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](#) on page 929 for detailed information on these policies.

Timeout

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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 930.

To configure IP defragmentation:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

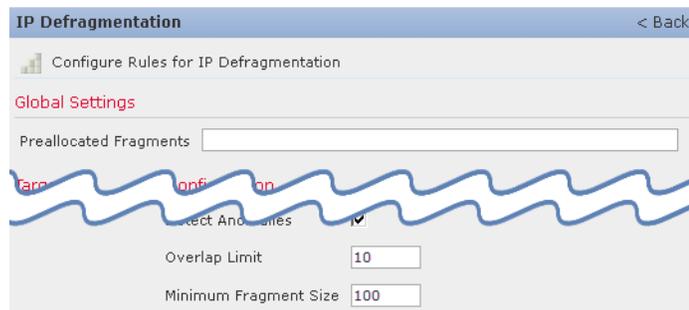
The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. 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.

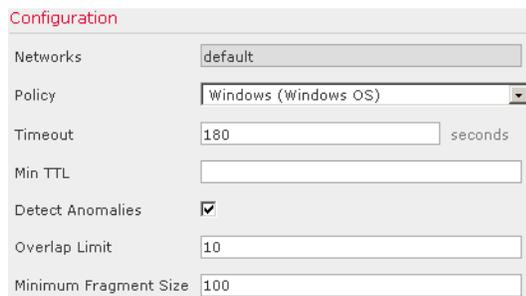


A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Optionally, you can modify the setting for **Preallocated Fragments** in the Global Settings page area.



6. 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, 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
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.



The screenshot shows a configuration window titled "Configuration". It contains several fields and a checkbox:

Networks	default
Policy	Windows (Windows OS)
Timeout	180 seconds
Min TTL	
Detect Anomalies	<input checked="" type="checkbox"/>
Overlap Limit	10
Minimum Fragment Size	100

7. Optionally, you can modify any of the options in the Configuration page area.
8. Optionally, click **Configure Rules for IP Defragmentation** at the top of the page to display rules associated with individual options.
Click **Back** to return to the IP Defragmentation page.
9. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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. For more information on packet decoding, see [Capturing and Decoding Packets](#) on page 584.

Note that you must enable packet decoder rules, which have a generator ID (GID) of 116, if you want these rules to generate events. A link on the configuration page takes you to a filtered view of packet decoder rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

If no preprocessor rule is mentioned, 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` system default variable to modify the ports that identify encapsulated GTP traffic. See [Understanding Existing Variables](#) on page 699 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](#) on page 729 and [Setting Rule States](#) on page 743 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](#) on page 743 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)

TCP OPTION	DESCRIPTION
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.

IMPORTANT! 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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 [Configuring Inline Normalization](#) on page 922 for more information.

Configuring Packet Decoding

LICENSE: Protection

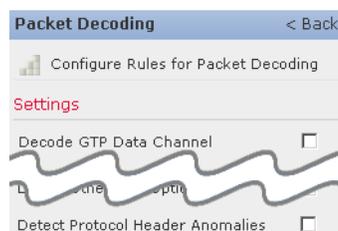
You can configure packet decoding on the Packet Decoding configuration page. For more information packet decoding configuration options, see [Understanding Packet Decoding](#) on page 934.

To configure packet decoding:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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 intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can enable or disable any of the detection options on the Packet Decoding page. See [Understanding Packet Decoding](#) on page 934 for more information.
6. Optionally, click **Configure Rules for Packet Decoding** at the top of the page to display rules associated with individual options.
Click **Back** to return to the Packet Decoding page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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. A link on the configuration page takes you to a filtered view of TCP stream preprocessor rules on the intrusion policy Rules page, where you can enable and disable rules and configure other rule actions. See [Setting Rule States](#) on page 743 for more information.

Note also that when a rule that requires this preprocessor is enabled in an intrusion policy, you must enable the preprocessor or choose to allow the system to enable it automatically before you can save the policy. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

If you enable any of the following, TCP stream preprocessing must be enabled:

- the DCE/RPC preprocessor when the RPC over HTTP proxy, RPC over HTTP server, TCP, or SMB transport protocol is selected
- the DNS preprocessor
- the FTP/Telnet preprocessor
- the HTTP Inspect preprocessor
- the IMAP preprocessor
- the POP preprocessor
- the SMTP preprocessor
- the SSL preprocessor
- the Modbus preprocessor
- the DNP3 preprocessor
- portscan detection when the TCP protocol is selected
- TCP intrusion rules that use the `flow`, `flowbits`, `stream-size`, or `stream-reassemble` keyword

See the following sections for more information:

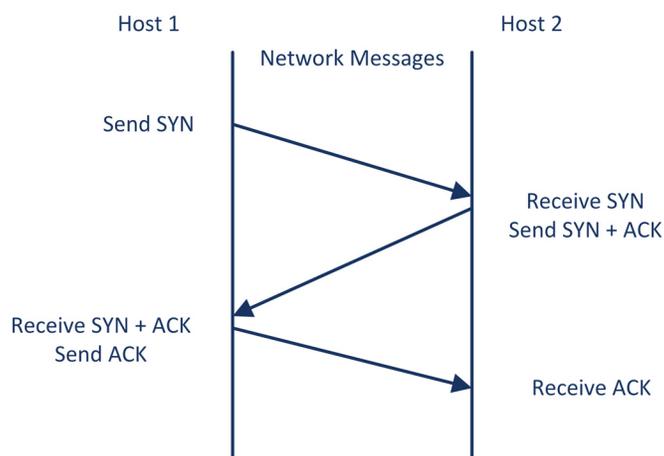
- [Understanding State-Related TCP Exploits](#) on page 941.
- [Initiating Active Responses with Drop Rules](#) on page 941.
- [Selecting TCP Global Options](#) on page 943.
- [Understanding Target-Based TCP Policies](#) on page 943.
- [Selecting TCP Policy Options](#) on page 945.

- [Reassembling TCP Streams](#) on page 949.
- [Configuring TCP Stream Preprocessing](#) on page 952.

Understanding State-Related TCP Exploits

LICENSE: Protection

If you add the `flow` keyword with the `established` argument to an intrusion rule, the 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`.

Initiating Active Responses with Drop Rules

LICENSE: Protection

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](#) on page 956 for further explanation of how the stream processor 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. When active responses are enabled in a passive deployment, the system cannot drop the packet or insert resets but sends 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 a UDP 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. Note that to initiate additional TCP resets you must ensure that TCP Stream Configuration is enabled, and to initiate additional ICMP unreachable packets you must ensure that UDP Stream Configuration is enabled. See [Modifying Advanced Settings](#) on page 774 for more information.

See [Selecting TCP Global Options](#) on page 943 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](#) on page 1160 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](#) on page 1164 for more information.

Selecting TCP Global Options

LICENSE: Protection

This section describes the options that control how the TCP stream preprocessor functions.

If no preprocessor rule is mentioned, the option is not associated with a preprocessor rule.

Packet Type Performance Boost

Enables ignoring TCP traffic for all ports and application protocols that are not specified in enabled 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.

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 Drop Rules](#) on page 941 and [Initiating Active Responses with Rule Keywords](#) on page 1160.

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.

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 [Using Adaptive Profiles](#) on page 1003.

The [TCP Operating System Policies](#) table identifies the operating system policies and the host operating systems that use each.

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.

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

TCP Operating System Policies (Continued)

POLICY	OPERATING SYSTEMS
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)

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, 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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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).

Policy

Identifies the TCP policy operating system of the target host or hosts. For more information, see [Understanding Target-Based TCP Policies](#) on page 943.

Timeout

The number of seconds between 1 and 86400 the rules engine keeps an inactive stream in the state table. If the stream is not reassembled in the specified time, the rules engine deletes it from the state table.

IMPORTANT! If your managed 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.

WARNING! 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](#) on page 743 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](#) on page 743 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 you must enable the Inline Normalization **Normalize TCP** option for this option to be effective. See [Normalizing Inline Traffic](#) on page 918 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 743 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](#) on page 950.

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](#) on page 950.

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.

IMPORTANT! Any port you add to the server-level FTP port list, or the DCE/RPC, HTTP, SMTP, Session Initiation Protocol, POP, IMAP, or SSL port list should also be added in each TCP policy to the appropriate list of TCP reassembly ports, depending on whether you are monitoring client or server traffic, or both. Note, however, that reassembling additional traffic types (client, server, both) increases resource demands. For more information, [Configuring the DCE/RPC Preprocessor](#) on page 821, see [Understanding Server-Level FTP Options](#) on page 838, [Selecting Server-Level HTTP Normalization Options](#) on page 854, [Decoding the Session Initiation Protocol](#) on page 872, [Decoding IMAP Traffic](#) on page 880, [Decoding POP Traffic](#) on page 884, [Decoding SMTP Traffic](#) on page 889, [Using the SSL Preprocessor](#) on page 905, and [Understanding Target-Based TCP Policies](#) on page 943.

See the following sections for more information:

- [Understanding Stream-Based Attacks](#) on page 950
- [Selecting Stream Reassembly Options](#) on page 950

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 [Using Adaptive Profiles](#) on page 1003 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 `a11` as the argument to provide reassembly for all ports, Sourcefire does **not** recommend setting ports to `a11` because it may increase the amount of traffic inspected by this preprocessor and slow performance unnecessarily.

If no preprocessor rule is mentioned, 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.

At least one client detector must be enabled (see [Activating and Deactivating Detectors](#) on page 1678) for each client service you select. By default, all Sourcefire-provided detectors are activated. If no detector is enabled for an associated client application, the system automatically enables all Sourcefire-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application.

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.

At least one detector must be enabled (see [Activating and Deactivating Detectors](#) on page 1678) for each service you select. By default, all Sourcefire-provided detectors are activated. If no detector is enabled for a service, the system automatically enables all Sourcefire-provided detectors for the associated application protocol; if none exist, the system enables the most recently modified user-defined detector for the application protocol.

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.

At least one detector must be enabled (see [Activating and Deactivating Detectors](#) on page 1678) for each service you select. By default, all Sourcefire-provided detectors are activated. If no detector is enabled for an associated client application or application protocol, the system automatically enables all Sourcefire-provided detectors for the application or application protocol; if none exist, the system enables the most recently modified user-defined detector for the application or application protocol.

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](#) on page 945.

To configure the stream preprocessor to track TCP sessions:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

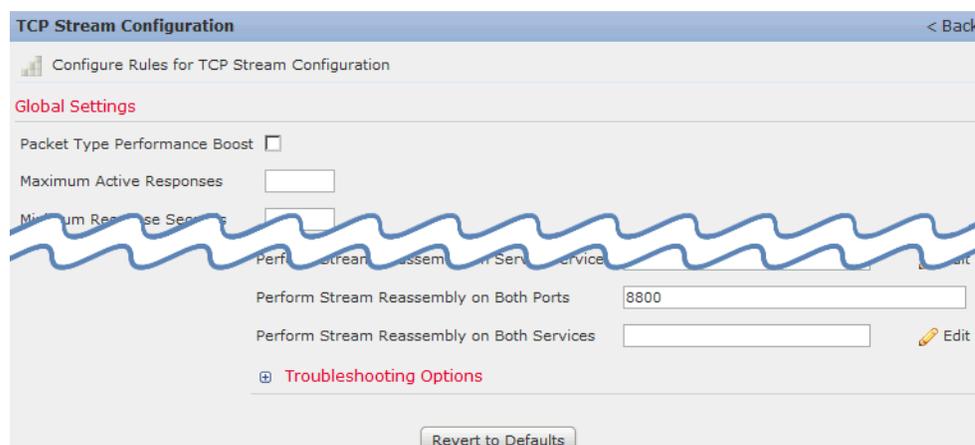
3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. 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**.

IMPORTANT! You cannot disable TCP stream preprocessing when the DNS, FTP/Telnet, HTTP Inspection, SMTP, or SSL preprocessor is enabled, or when the DCE/RPC preprocessor is enabled with the RPC over HTTP proxy, RPC over HTTP server, TCP, or SMB transport protocol selected, or when portscan detection is enabled with the TCP protocol selected. Also, you should not disable TCP stream preprocessing when you have TCP rules enabled that use the `flow` or `flowbits` keyword because these rules will not trigger unless TCP stream preprocessing is enabled.

The TCP Stream Configuration page appears.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Optionally, modify any of the options under Global Settings. See [Selecting TCP Global Options](#) on page 943 for more information.



6. You have two options:

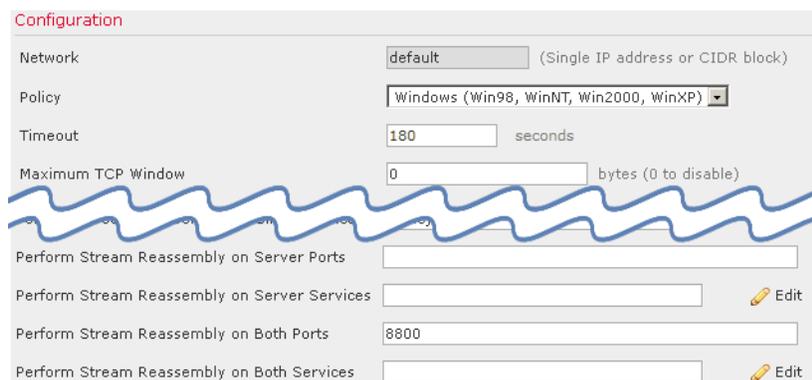
- Add a new target-based policy. Click the add icon () next to **Hosts** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses in the **Host Address** field and click **OK**.

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 Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

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.



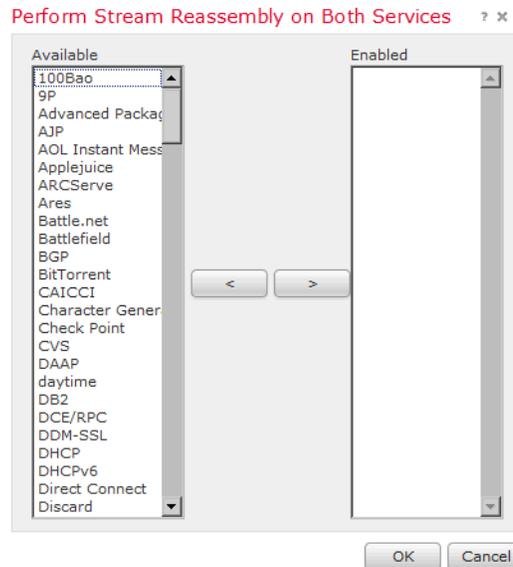
Configuration	
Network	default (Single IP address or CIDR block)
Policy	Windows (Win98, WinNT, Win2000, WinXP)
Timeout	180 seconds
Maximum TCP Window	0 bytes (0 to disable)
Perform Stream Reassembly on Server Ports	<input type="checkbox"/>
Perform Stream Reassembly on Server Services	<input type="checkbox"/> Edit
Perform Stream Reassembly on Both Ports	8800
Perform Stream Reassembly on Both Services	<input type="checkbox"/> Edit

7. 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 8; otherwise, go to step 11.

For more information, see [Selecting TCP Policy Options](#) on page 945, and [Selecting Stream Reassembly Options](#) on page 950.

8. 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 [Working with Servers](#) on page 1418 and [Using Adaptive Profiles](#) on page 1003 for more information.

9. 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.

10. Click **OK** to add the selections. The TCP Stream Configuration page is displayed and the services are updated.
11. Optionally, click **Configure Rules for TCP Stream Configuration** at the top of the page to display rules associated with individual TCP policy options. Click **Back** to return to the TCP Stream Configuration page.

12. Optionally, modify any of the TCP stream preprocessing global or policy troubleshooting options only if asked to do so by Sourcefire Support; click the **+** sign next to **Troubleshooting options** to expand the troubleshooting options section. For more information, see [Understanding Troubleshooting Options](#) on page 790.
13. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 1108) 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](#) on page 934.

Note also that UDP stream preprocessing can be automatically enabled when a rule that requires UDP stream preprocessing is enabled. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

The following configurations require UDP stream preprocessing to be enabled:

- DNS preprocessor
- SIP preprocessor

- DCE/RPC preprocessor with the UDP transport protocol selected
- UDP intrusion rules that use the `flow`, `flowbits`, or `stream-size` keyword

Configuring UDP Stream Preprocessing

LICENSE: Protection

You can configure UDP stream preprocessing.

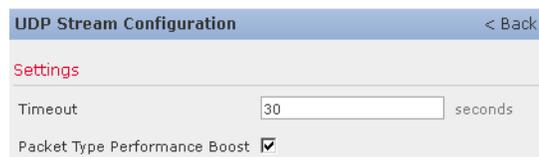
To configure the stream preprocessor to track UDP sessions:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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**.

IMPORTANT! You cannot disable UDP stream preprocessing when the DCE/RPC preprocessor is enabled with the UDP transport protocol selected, or when portscan detection is enabled with the UDP protocol selected. Also, you should not disable UDP stream preprocessing when you have UDP intrusion rules enabled that use the `flow` or `flowbits` keyword because these rules will not trigger unless UDP stream preprocessing is enabled.

The UDP Stream Configuration page appears.



UDP Stream Configuration < Back

Settings

Timeout 30 seconds

Packet Type Performance Boost

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Optionally, configure a **Timeout** value to specify the number of seconds between 1 and 86400 the predecessor keeps an inactive stream in the state table. If additional datagrams are not seen in the specified time, the predecessor deletes the stream from the state table.
6. 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.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 25

DETECTING SPECIFIC THREATS

You can use some of the advanced configuration options in an intrusion policy to detect specific threats, such as back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic. See the following sections for more information:

- [Detecting Back Orifice](#) on page 959 explains detection of Back Orifice attacks.
- [Detecting Portscans](#) on page 961 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](#) on page 971 explains how to limit denial of service (DoS) and SYN flood attacks.
- [Detecting Sensitive Data](#) on page 983 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 Sourcefire 3D System 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

[Back Orifice GID:SDs](#) table for the preprocessor to generate corresponding events. A link on the configuration page takes you to a filtered view of Back Orifice preprocessor rules on the Rules page, where you can enable and disable rules and configure other rule attributes. See [Setting Rule States](#) on page 743 for more information.

Back Orifice GID:SDs

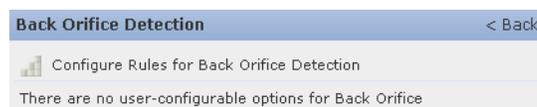
PREPROCESSOR RULE GID:SID	DESCRIPTION
105:1	Back Orifice traffic detected
105:2	Back Orifice client traffic detected
105:3	Back Orifice server traffic detected
105:4	Back Orifice snort buffer attack detected

To view the [Back Orifice Detection](#) page:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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 an Intrusion Policy](#) on page 792 for more information.

5. Optionally, click **Configure Rules for Back Orifice Detection** at the top of the page. A filtered view appears of Back Orifice preprocessor rules on the Rules page, where you can enable and disable rules and configure other rule attributes. See [Setting Rule States](#) on page 743 for more information.
Note that you must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline policy, if you want the preprocessor to log intrusion events.
Click **Back** to return to the Back Orifice Detection page.
6. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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 Rules page with generator ID (GID) 122 for enabled portscan types for the portscan detector to generate portscan events. A link on the configuration page takes you to a filtered view of portscan detection rules on the Rules page, where you can enable and disable rules and configure other rule attributes. See [Setting Rule States](#) on page 743 and the [Portscan Detection SIDs \(GID:122\) table](#) on page 968 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. Sourcefire'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 [Protocol Types](#) table describes the protocols you can activate in the portscan detector.

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.

IMPORTANT! 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 [Portscan Types](#) table describes the kinds of portscan activity you can detect.

Portscan Types

TYPE	DESCRIPTION
Portscan Detection	<p>A one-to-one portscan in which an attacker uses one or a few hosts to scan multiple ports on a single target host.</p> <p>One-to-one portscans are characterized by:</p> <ul style="list-style-type: none">• a low number of scanning hosts• a single host that is scanned• a high number of ports scanned <p>This option detects TCP, UDP, and IP portscans.</p>
Port Sweep	<p>A one-to-many portsweep in which an attacker uses one or a few hosts to scan a single port on multiple target hosts.</p> <p>Portsweeps are characterized by:</p> <ul style="list-style-type: none">• a low number of scanning hosts• a high number of scanned hosts• a low number of unique ports scanned <p>This option detects TCP, UDP, ICMP, and IP portsweeps.</p>
Decoy Portscan	<p>A one-to-one portscan in which the attacker mixes spoofed source IP addresses with the actual scanning IP address.</p> <p>Decoy portscans are characterized by:</p> <ul style="list-style-type: none">• 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 <p>The decoy portscan option detects TCP, UDP, and IP protocol portscans.</p>
Distributed Portscan	<p>A many-to-one portscan in which multiple hosts query a single host for open ports.</p> <p>Distributed portscans are characterized by:</p> <ul style="list-style-type: none">• 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 <p>The distributed portscan option detects TCP, UDP, and IP protocol portscans.</p>

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 [Sensitivity Levels](#) table describes the three different sensitivity levels you can choose from.

Sensitivity Levels

LEVEL	DESCRIPTION
Low	<p>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.</p> <p>This level uses the shortest time window for portscan detection.</p>
Medium	<p>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.</p> <p>Note that you can add the IP addresses of these active hosts to the Ignore Scanned field to mitigate this type of false positive.</p> <p>This level uses a longer time window for portscan detection.</p>
High	<p>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.</p> <p>This level uses a much longer time window for portscan detection.</p>

See the following sections for more information:

- [Configuring Portscan Detection](#) on page 965
- [Understanding Portscan Events](#) on page 968

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](#) on page 743 and the [Portscan Detection SIDs \(GID:122\) table](#) on page 968 for more information.

To configure portscan detection:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. 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.

Portscan Detection < Back

Configure Rules for Portscan Detection

Settings

Protocol: TCP, UDP, ICMP, IP

Scan Type: Portscan Detection, Port Sweep, Decoy Portscan, Distributed Portscan

Sensitivity Level: Low

Watch IP: (Single IP address, CIDR block, or comma-separated list)

Ignore Scanners: (Single IP address, CIDR block, or comma-separated list)

Ignore Scanned: (Single IP address, CIDR block, or comma-separated list)

Detect Ack Scans:

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. 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](#) on page 962 for more information.

Note that you must ensure that TCP stream processing is enabled to detect scans over TCP, and that UDP stream processing is enabled to detect scans over UDP.

6. 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](#) on page 963 for more information.

7. In the **Sensitivity Level** list, select the level you want to use: low, medium, or high.

See the [Sensitivity Levels table](#) on page 964 for more information.

8. 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

9. 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

10. 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

11. Optionally, clear the **Detect Ack Scans** check box to discontinue monitoring of sessions picked up in mid-stream.

IMPORTANT! Detection of mid-stream sessions helps to identify ACK scans, but may cause false events, particularly on networks with heavy traffic and dropped packets.

12. Set the portscan detection rules for each enabled portscan type to Generate Events; click **Configure Rules for Portscan Detection** at the top of the page to display rules associated with individual TCP policy options.

Note that although you can set portscan rules to Drop and Generate Events, the portscan detector does not drop packets, including in an inline deployment.

See [Setting Rule States](#) on page 743 for information on setting rule states.

To identify the rules associated with different portscan types, see the [Portscan Detection SIDs \(GID:122\) table](#) on page 968.

Click **Back** to return to the Portscan Detection page.

13. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 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](#) on page 743 for more information. The **Preprocessor Rule SID** column in the [Portscan Detection SIDs \(GID:122\)](#) table lists the SID for the preprocessor rule you must enable for each portscan type.

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

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.
IP	Medium or High	Does not generate events.	
	Low	10	
		Medium or High	14
	<hr/>		
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. You can follow the procedures in [Working with Intrusion Events](#) on page 593.

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.

IMPORTANT! 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 [Portscan Packet View](#) table describes the information provided in the packet view for portscan events. For any IP address, you can click the address to view

the context menu and select **whois** to perform a lookup on the IP address or **View Host Profile** to view the host profile for that host.

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.
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.

Portscan Packet View (Continued)

INFORMATION	DESCRIPTION
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](#) on page 971
- [Rate-Based Attack Prevention and Other Filters](#) on page 974
- [Configuring Rate-Based Attack Prevention](#) on page 981
- [Understanding Dynamic Rule States](#) on page 757
- [Setting a Dynamic Rule State](#) on page 758

Understanding Rate-Based Attack Prevention

LICENSE: Protection

You can configure your intrusion policy to include rate-based filters that detect excessive activity directed at hosts on your network. You can use this feature on managed devices deployed in inline mode to block rate-based attacks for a specified time and 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](#) on page 974.

- 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](#) on page 974.
- 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](#) on page 758.
- excessive matches for a particular rule across all traffic.
To configure rule-based dynamic rule states, see [Setting a Dynamic Rule State](#) on page 758.

In an intrusion policy, you can either configure SYN flood or TCP/IP connection flood detection for the entire policy, or 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. See [Preventing SYN Attacks](#) on page 974 and [Controlling Simultaneous Connections](#) on page 974 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

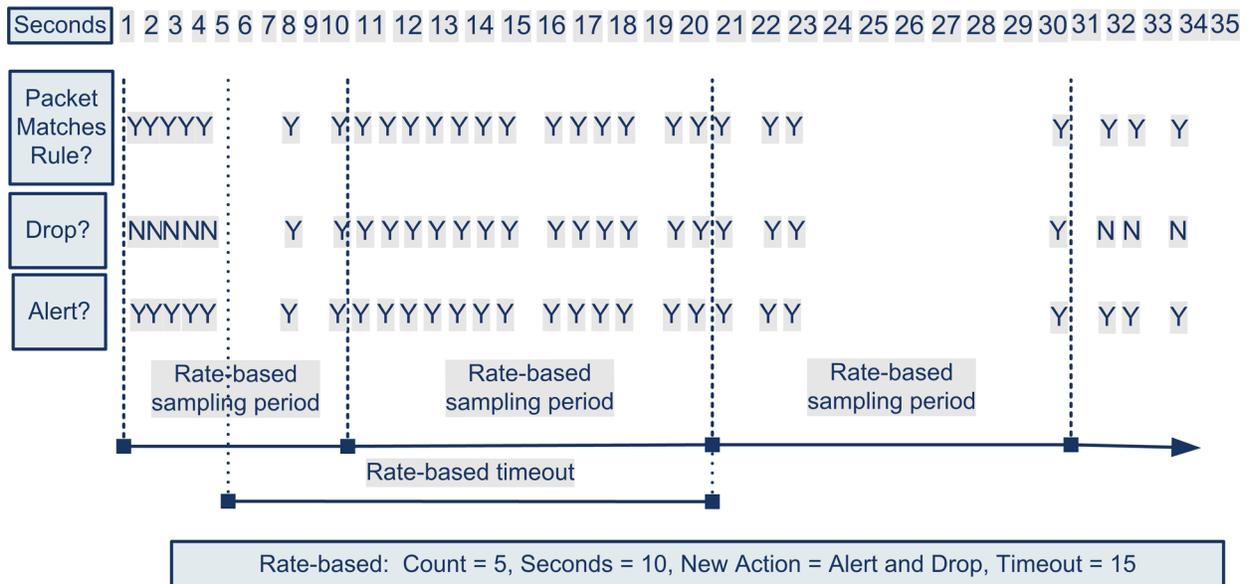
period of time that the rate action is active, even if those rules are not initially set to Drop and Generate Events.

IMPORTANT! 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 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](#) on page 975
- [Dynamic Rule States and Thresholding or Suppression](#) on page 976

- [Policy-Wide Rate-Based Detection and Thresholding or Suppression](#) on page 978
- [Rate-Based Detection with Multiple Filtering Methods](#) on page 979

Rate-Based Attack Prevention and Detection Filtering

LICENSE: Protection

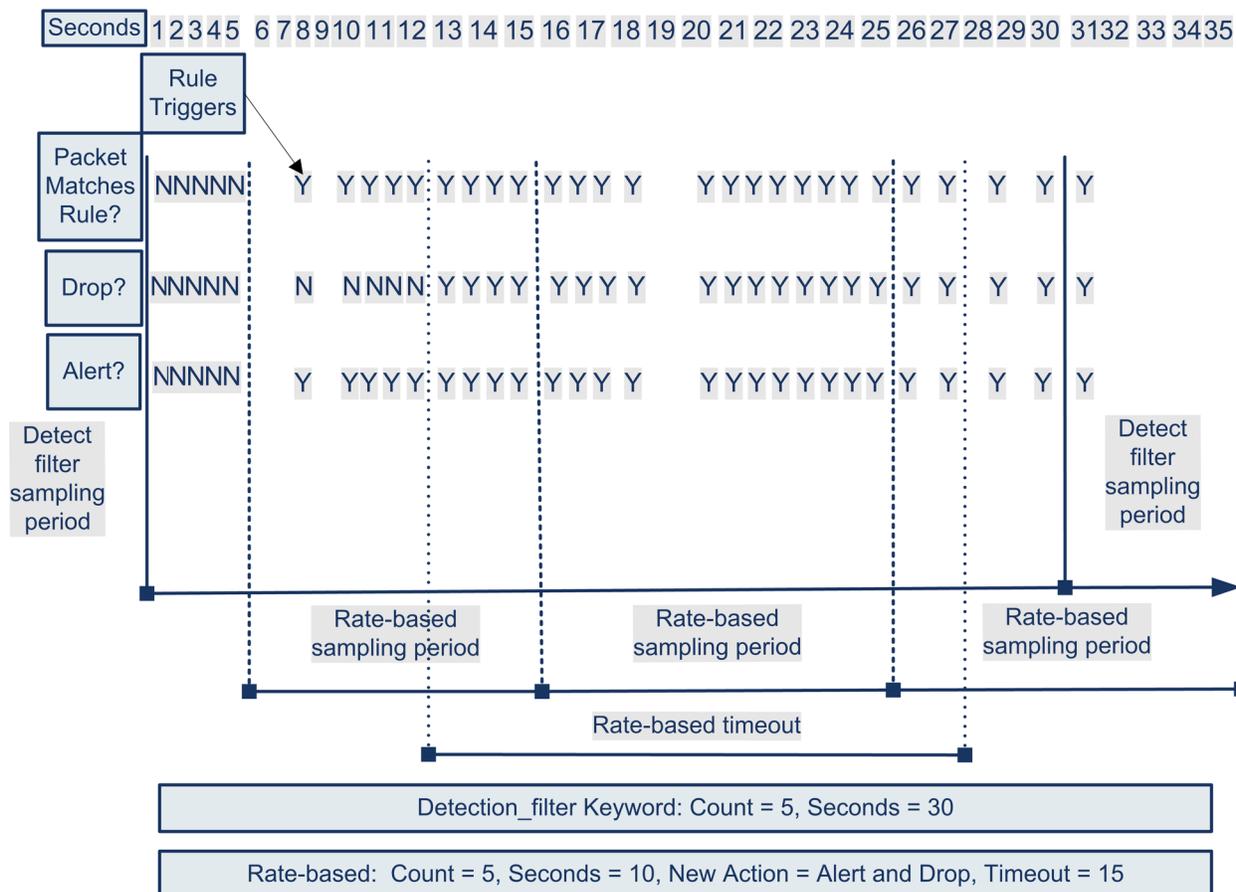
The `detection_filter` keyword prevents a rule from triggering until a threshold number of rule matches occur within a specified time. When a rule includes the `detection_filter` keyword, the system tracks the number of incoming packets matching the pattern in the rule per timeout period. The system 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](#) on page 743.

Dynamic Rule States and Thresholding or Suppression

LICENSE: Protection

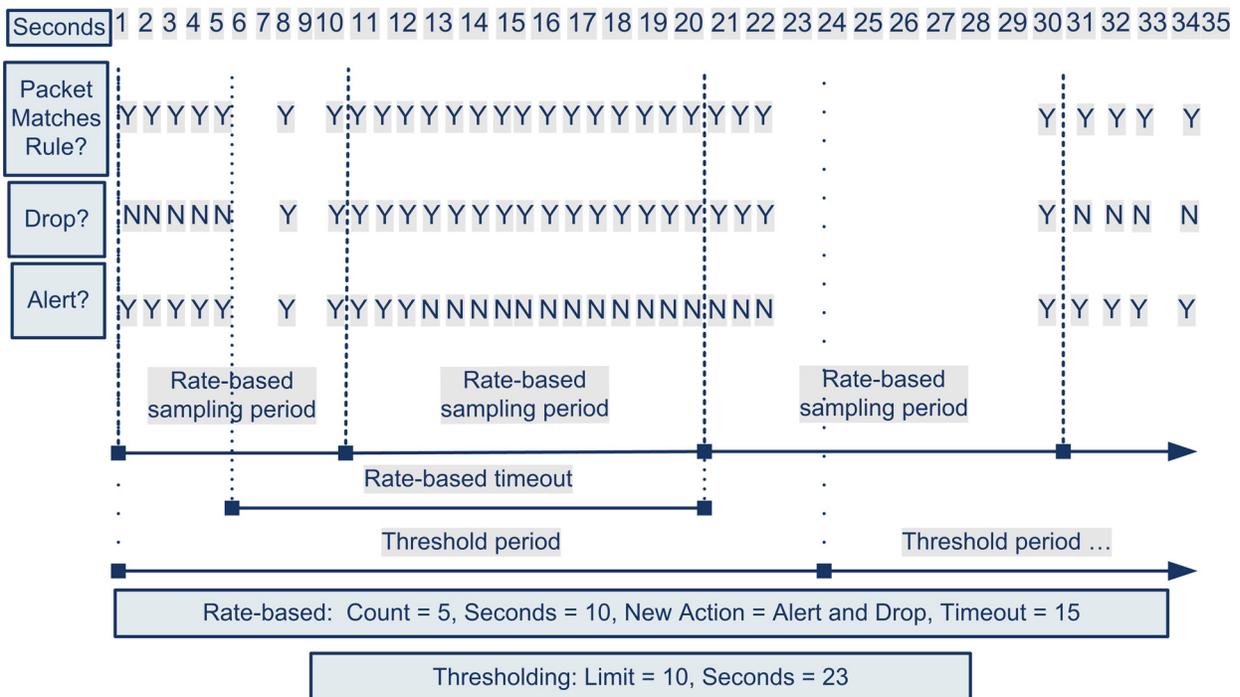
You can use thresholding and suppression to reduce excessive events by limiting the number of event notifications for a rule or by suppressing notifications altogether for that rule. For more information on the available options for thresholding and suppression, see [Configuring Event Thresholding](#) on page 747 and [Configuring Suppression Per Intrusion Policy](#) on page 753.

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](#) on page 1012, [Configuring Event Thresholding](#) on page 747, and [Configuring Suppression Per Intrusion Policy](#) on page 753.

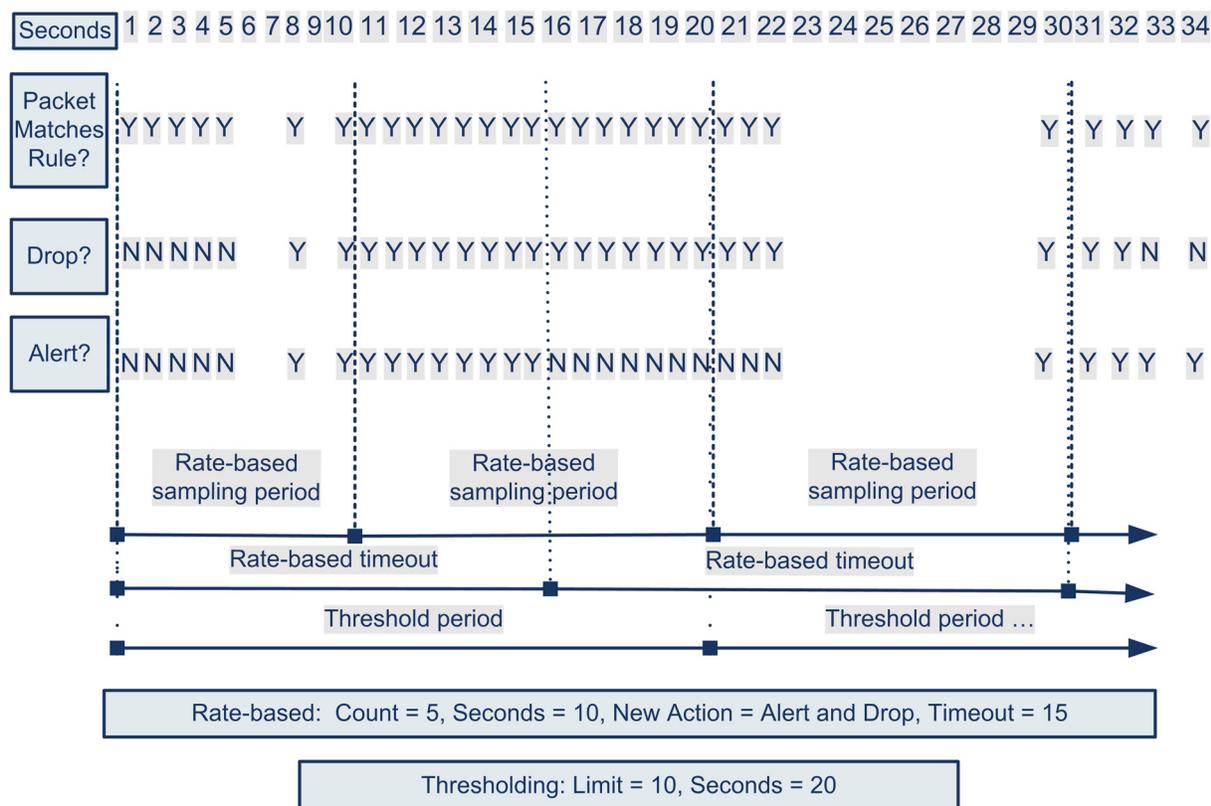
If suppression is applied to a rule, event notifications for that rule for all applicable IP addresses are suppressed even if a rate-based action change occurs because of a policy-wide or rule-specific rate-based setting. However, the interaction between thresholding and rate-based criteria is more complex.

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 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

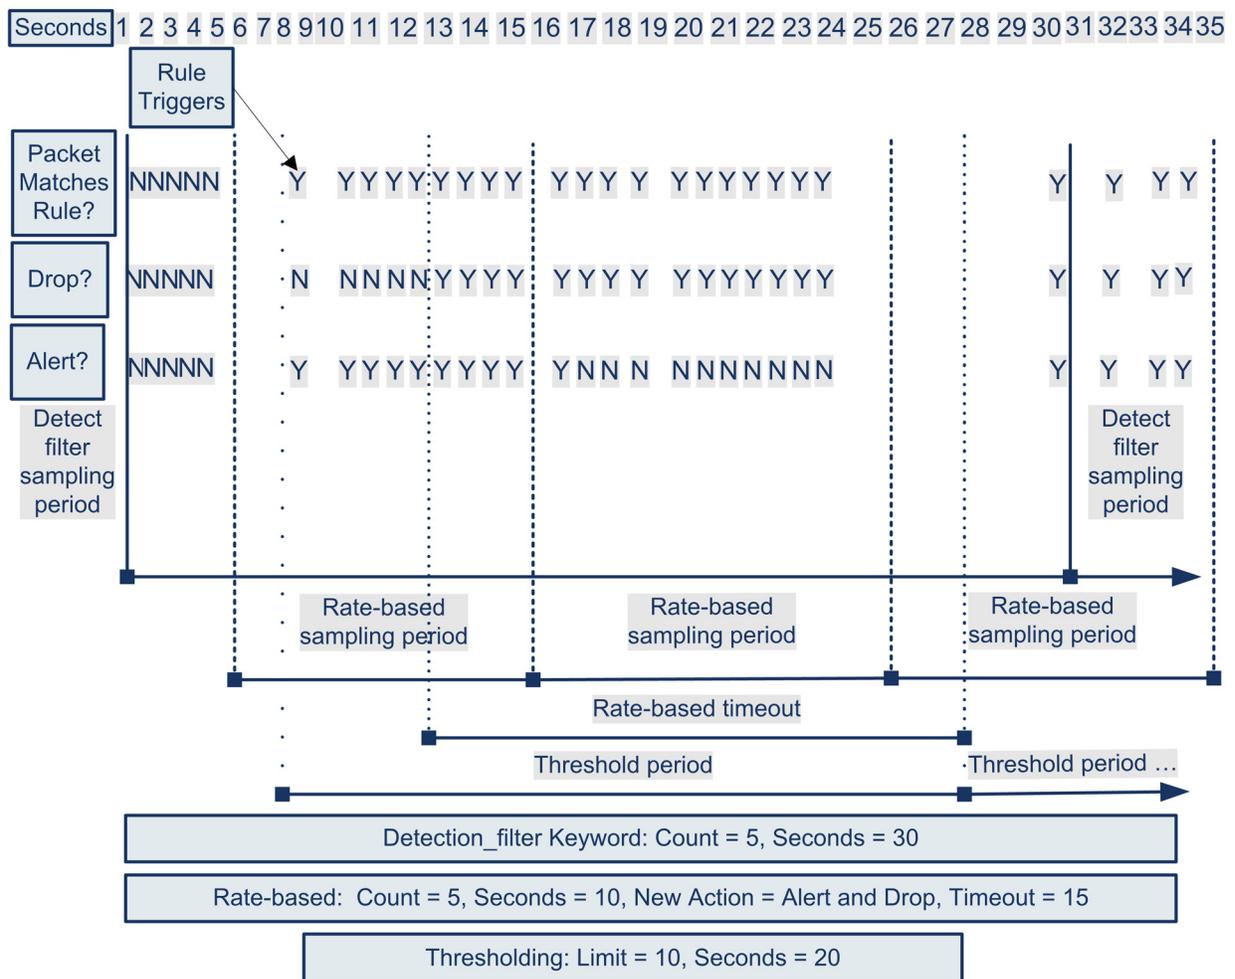
You may encounter situations where the `detection_filter` keyword, thresholding or suppression, and rate-based criteria all apply to the same traffic. When you enable suppression for a rule, events are suppressed for the specified IP addresses even if a rate-based change occurs.

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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. 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.

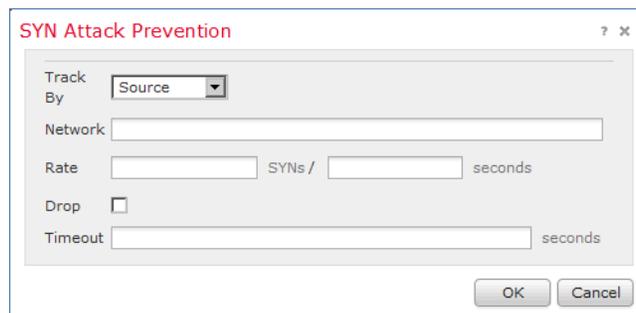
Rate-Based Attack Prevention < Back					
SYN Attack Prevention Add					
Track By	Network	Rate	Drop	Timeout	
Control Simultaneous Connections Add					
Track By	Network	Count	Drop	Timeout	

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You have two options:

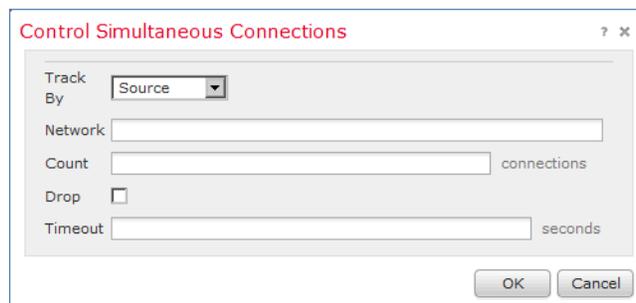
- To prevent incomplete connections intended to flood a host, click **Add** under **SYN Attack Prevention**.

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.



6. 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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

7. 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.
8. To drop packets matching the rate-based attack prevention settings, select **Drop**.
9. 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.

WARNING! 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.

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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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)-<i>123-4567</i>`. 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. Sourcefire 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 [Automatically Enabling Advanced Settings](#) on page 787 for more information.

Because the system uses TCP stream preprocessing to establish monitored sessions, TCP stream preprocessing must be enabled to use sensitive data detection in your policy. When you save changes to your policy, you are given the option to automatically enable TCP stream preprocessing if sensitive data detection is enabled and TCP stream preprocessing is disabled. See [Using TCP Stream Preprocessing](#) on page 940 for more information.

See the following sections for more information:

- [Deploying Sensitive Data Detection](#) on page 985
- [Selecting Global Sensitive Data Detection Options](#) on page 985
- [Selecting Individual Data Type Options](#) on page 987
- [Using Predefined Data Types](#) on page 988
- [Configuring Sensitive Data Detection](#) on page 989
- [Selecting Application Protocols to Monitor](#) on page 993
- [Special Case: Detecting Sensitive Data in FTP Traffic](#) on page 995
- [Using Custom Data Types](#) on page 996

Deploying Sensitive Data Detection

LICENSE: Protection

Because sensitive data detection can have a high impact on the performance of your Sourcefire 3D System, Sourcefire recommends that you adhere to the following guidelines when creating your intrusion policy and applying it as part of an access control policy:

- Select the No Rules Active default policy as your base policy; see [Selecting the Base Policy](#) on page 695 for more information.
- Ensure that the IP Defragmentation, FTP and Telnet Configuration, and TCP Stream Configuration advanced settings are enabled in your intrusion policy; see [Modifying Advanced Settings](#) on page 774 for more information.
- Apply the access control policy that includes the intrusion policy containing your sensitive data configuration to a separate device reserved for sensitive data detection; see [Applying an Access Control Policy](#) on page 466 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 within an intrusion policy. That is, you can configure different global sensitive data settings in different intrusion policies, but not for different data types within the same intrusion policy.

The [Global Sensitive Data Detection Options](#) table describes the global sensitive data detection options you can configure.

Global Sensitive Data Detection Options

OPTION	DESCRIPTION
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 web interface and in downloaded packets. See Using the Packet View on page 622 for more information.
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 in the Sourcefire 3D System, see IP Address Conventions on page 57.
Global Threshold	<p>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.</p> <p>Sourcefire 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 on page 987 for more information.</p> <p>Note the following points regarding global thresholds:</p> <ul style="list-style-type: none">• You must enable preprocessor rule 139:1 to detect and generate events on combined data type occurrences. See Setting Rule States on page 743 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 Sourcefire uses an otherwise inaccessible `sd_pattern` keyword to define a built-in data pattern to detect in traffic. See the [Sensitive Data Types table](#) on page 989 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](#) on page 996 for more information.

Note that data type names and patterns are system-wide; all other data type options are policy-specific.

The [Individual Data Type Options](#) table describes the data type options you can configure.

Individual Data Type Options

OPTION	DESCRIPTION
Data Type	Displays the unique name for the data type.
Threshold	<p>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.</p> <p>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.</p>
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.

Individual Data Type Options (Continued)

OPTION	DESCRIPTION
Application Protocols Note that this feature requires Protection and Control licenses.	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. At least one detector must be enabled (see Activating and Deactivating Detectors on page 1678) for each application protocol you select. By default, all Sourcefire-provided detectors are activated. If no detector is enabled for an application protocol, the system automatically enables all Sourcefire-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. See Selecting Application Protocols to Monitor on page 993 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 Sourcefire are predefined). See Using Custom Data Types on page 996 for more information. The web 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 to enable detection, and event generation, for each data type you want to use in your policy. See [Setting Rule States](#) on page 743 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](#) on page 729 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 [Sensitive Data Types](#) 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.

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 <code>(\d{3}) ?\d{3}-\d{4}</code> .	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 on page 996 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 will be prompted to enable sensitive data detection when you save changes to your policy. See [Automatically Enabling Advanced Settings](#) on page 787 for more information.

The [Sensitive Data Configuration Actions](#) table describes actions you can take on the Sensitive Data Detection page.

Sensitive Data Configuration Actions

To...	YOU CAN...
modify global settings	see the Global Sensitive Data Detection Options table on page 986 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 on page 987 for information on the options you can modify.

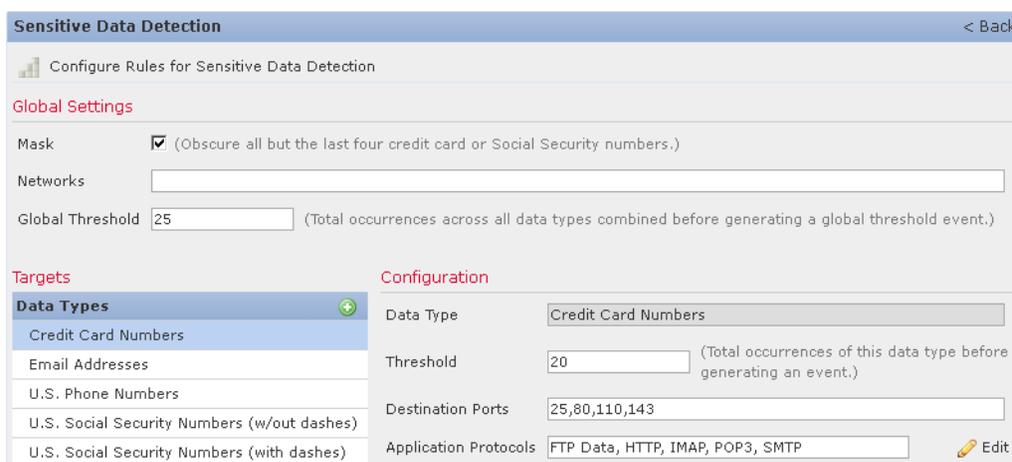
Sensitive Data Configuration Actions (Continued)

To...	YOU CAN...
<p>add or remove application protocols to monitor for a data type</p> <p>Note that this feature requires Protection and Control licenses.</p>	<p>click inside the Application Protocols field, or click Edit next to the field. The Application Protocols pop-up window appears:</p> <ul style="list-style-type: none">• 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. <p>Use Ctrl or Shift while clicking to select multiple application protocols. You can also click and drag to select multiple adjacent application protocols.</p> <p>At least one detector must be enabled (see Activating and Deactivating Detectors on page 1678) for each application protocol you select. By default, all Sourcefire-provided detectors are activated. If no detector is enabled for an application protocol, the system automatically enables all Sourcefire-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application.</p> <p>IMPORTANT! To detect sensitive data in FTP traffic, you must add the Ftp data application protocol and enable the FTP/Telnet preprocessor. See Special Case: Detecting Sensitive Data in FTP Traffic on page 995 for more information.</p>
<p>create a custom data type</p>	<p>click the + sign next to Data Types on the left side of the page. The Add Data Type pop-up window appears.</p> <p>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 on page 996 for more information.</p>
<p>display sensitive data preprocessor rules</p>	<p>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.</p> <p>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 on page 743 for more information.</p> <p>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 Managing Rules in an Intrusion Policy on page 717 for more information.</p> <p>Click Back to return to the Sensitive Data Detection page.</p>

To configure sensitive data detection:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
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 an Intrusion Policy](#) on page 792 for more information.

5. You can take any of the actions described in the [Sensitive Data Configuration Actions table](#) on page 990.
6. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Selecting Application Protocols to Monitor

LICENSE: Control

You can specify up to eight application protocols to monitor for each data type. See [Working with Servers](#) on page 1418 for more information on the application protocols the system can detect on your network.

At least one detector must be enabled (see [Activating and Deactivating Detectors](#) on page 1678) for each application protocol you select. By default, all Sourcefire-provided detectors are activated. If no detector is enabled for an application protocol, the system automatically enables all Sourcefire-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application.

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, Sourcefire 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 and enable the FTP/Telnet preprocessor, and there is no advantage in specifying a port number. See [Special Case: Detecting Sensitive Data in FTP Traffic](#) on page 995 and [Decoding FTP and Telnet Traffic](#) on page 832 for more information.

To modify application protocols to detect sensitive data:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

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.

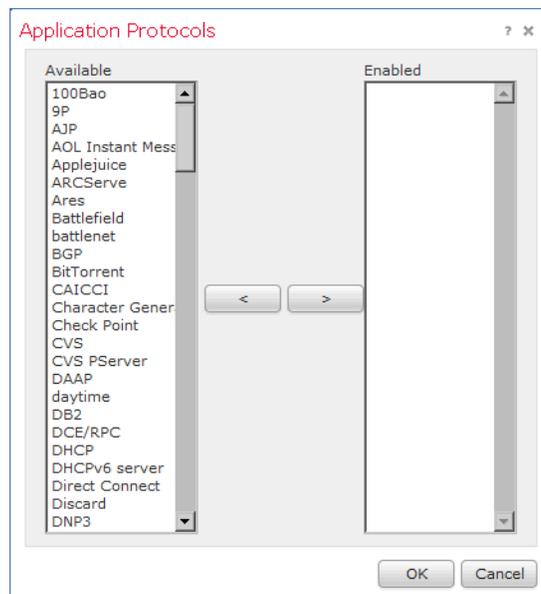
A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

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.

6. Click inside the **Application Protocols** field, or click **Edit** next to the field.

The Application Protocols pop-up window appears.



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.

IMPORTANT! To detect sensitive data in FTP traffic, you must add the **FTP data** application protocol and ensure that the FTP/Telnet preprocessor is enabled. See [Special Case: Detecting Sensitive Data in FTP Traffic](#) on page 995 for more information.

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](#) on page 993 for more information.
- Ensure that the FTP/Telnet preprocessor is enabled.

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](#) on page 832 for more information.

- Ensure that the FTP Data detector, which is enabled by default, is enabled. See [Activating and Deactivating Detectors](#) on page 1678.
- 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, Sourcefire recommends that you specify the FTP command port 23. See [Configuring Sensitive Data Detection](#) on page 989 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 data 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](#) on page 743 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 custom sensitive data rules by selecting the local rule filtering category on the Rules page. See [Filtering Rules in an Intrusion Policy](#) on page 729 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](#) on page 997
- [Configuring Custom Data Types](#) on page 999
- [Editing Custom Data Type Names and Detection Patterns](#) on page 1001

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 [Sensitive Data Pattern Metacharacters](#) table describes the metacharacters you can use when defining a custom data pattern.

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.	<code>colou?r</code> matches <code>color</code> or <code>colour</code>
{ <i>n</i> }	Matches the preceding character or escape sequence <i>n</i> times.	For example, <code>\d{2}</code> matches <code>55</code> , <code>12</code> , and so on; <code>\l{3}</code> matches <code>Abc</code> , <code>www</code> , and so on; <code>\w{3}</code> matches <code>a1B</code> , <code>25C</code> , and so on; <code>x{5}</code> matches <code>xxxxx</code>
\	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class. See the Sensitive Data Pattern Character Classes table on page 998 for a description of the character classes you can use in sensitive data patterns.	<code>\?</code> matches a question mark, <code>\\</code> matches a backslash, <code>\d</code> matches numeric characters, and so on

You must use a backslash to escape the characters in the [Escaped Sensitive Data Pattern Characters](#) table for the sensitive data preprocessor to interpret them correctly as literal characters.

Escaped Sensitive Data Pattern Characters

USE THIS ESCAPED CHARACTER...	TO REPRESENT THIS LITERAL CHARACTER...
\?	?
\{	{
\}	}
\\	\

The [Sensitive Data Pattern Character Classes](#) table describes the character classes you can use when defining a custom sensitive data pattern.

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
\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](#) on page 987 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](#) on page 743 for information on enabling rules in your intrusion policy.

To create or modify a custom data type:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:

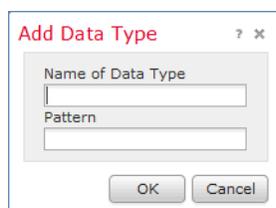
- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

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 an Intrusion Policy](#) on page 792 for more information.

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](#) on page 1001 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](#) on page 989 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](#) on page 1001.
- 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](#) on page 987 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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

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**.

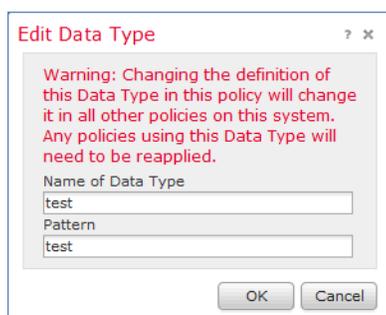
The Sensitive Data Detection page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

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.

6. Click the **Edit Data Type Name and Pattern** link.
The Edit Data Type pop-up window appears.



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](#) on page 997 for information on specifying the data pattern.

The Sensitive Data Detection page appears. If you clicked **OK**, the page displays your changes.

CHAPTER 26

USING ADAPTIVE PROFILES

Typically, the system uses the static settings you configure in an intrusion policy to process and analyze traffic. With the adaptive profiles feature, however, the system can adapt to network traffic by associating traffic with host information from the network map and then 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, Sourcefire recommends you configure adaptive profiles. In an inline deployment, Sourcefire recommends you configure the inline normalization preprocessor, with the Normalize TCP and Normalize TCP Payload options enabled. For more information, see [TCP Normalization](#) on page 921 and [Configuring Inline Normalization](#) on page 922.

For more information on using adaptive profiles to improve reassembly of packet fragments and TCP streams, see the following topics:

- [Understanding Adaptive Profiles](#) on page 1004
- [Configuring Adaptive Profiles](#) on page 1006

Understanding Adaptive Profiles

LICENSE: FireSIGHT + Protection

Adaptive profiles enable use of the most appropriate operating system profiles for IP defragmentation and for TCP stream preprocessing. For more information on the aspects of the intrusion policy affected by adaptive profiles, see [Defragmenting IP Packets](#) on page 927 and [Using TCP Stream Preprocessing](#) on page 940.

The system can use host information detected by network discovery, obtained through an Nmap scan, or added through the host input feature to adapt processing behavior.

IMPORTANT! When you input host information from a third-party application using the command line import utility or the host input API, you must first map the data to product definitions so the system can use it for adaptive profiles. For more information, see [Managing Third-Party Product Mappings](#) on page 1682.

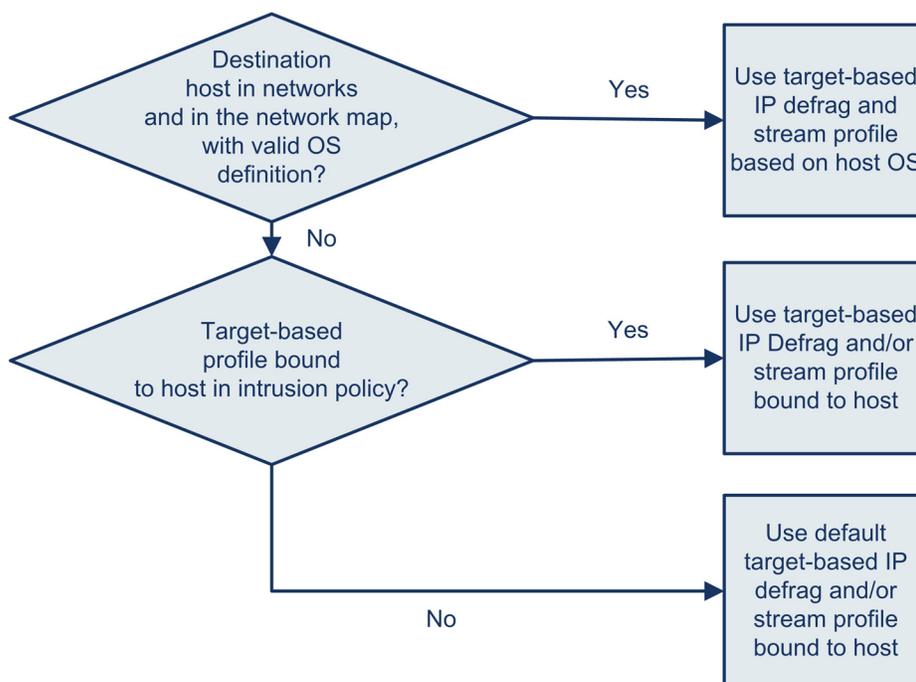
Using Adaptive Profiles with Preprocessors

LICENSE: FireSIGHT + Protection

Adaptive profiles, like the target-based profiles you can configure in an intrusion policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The rules engine then analyzes the data in the same format as that used by the destination host.

Manually configured target-based profiles only apply the default operating system profile you select or profiles you bind to specific hosts. Adaptive profiles, however, 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 an intrusion policy where adaptive profiles are enabled for the 10.6.0.0/16 subnet and where you have set the default IP Defragmentation target-based policy to Linux. The Defense Center where you configure the policy has a network map that 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 from the network map, where Host B is listed as 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](#) on page 927 for information on the IP Defragmentation preprocessor. See [Using TCP Stream Preprocessing](#) on page 940 for information on the stream preprocessor.

Adaptive Profiles and FireSIGHT Recommended Rules

LICENSE: FireSIGHT + Protection

Like FireSIGHT recommended rules, adaptive profiles compare metadata in a rule to host information to determine whether a rule should apply for a particular host. However, while FireSIGHT recommended rules provide recommendations for

enabling or disabling rules using that information, adaptive profiles use the information to apply specific rules to specific traffic.

FireSIGHT recommended rules require your interaction to implement suggested changes to rule states. Adaptive profiles, on the other hand, do not modify the intrusion policy. Adaptive treatment of rules happens on a packet-by-packet basis.

Additionally, FireSIGHT recommended rules can result in enabling disabled rules. Adaptive profiles, in contrast, only affect the application of rules that are already enabled in the intrusion policy. Adaptive profiles never change the rule state.

You can use adaptive profiles and FireSIGHT recommended rules in the same policy. Adaptive profiles use the rule state for a rule when the policy is applied to determine whether to include it as a candidate for applying, and your choices to accept or decline recommendations are reflected in that rule state. You can use both features to ensure that you have enabled or disabled the most appropriate rules for each network you monitor, and then to apply enabled rules most efficiently for specific traffic.

See [Managing FireSIGHT Rule State Recommendations](#) on page 764 for more information.

Configuring Adaptive Profiles

LICENSE: FireSIGHT + 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 exist in the network map and must be in the segment monitored by the device where you apply the access control policy that includes your intrusion policy.

You can indicate the hosts in the network map where adaptive profiles should be used to process traffic by specifying an IP address or block of addresses. You can also use a variable to indicate the network where adaptive profiles should be applied. See [Managing Variables](#) on page 698 for information on using variables.

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 in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

TIP! You can apply adaptive profiles to all hosts in the network map by using a variable with a value of `any` or by specifying `0.0.0.0/0` as the network value.

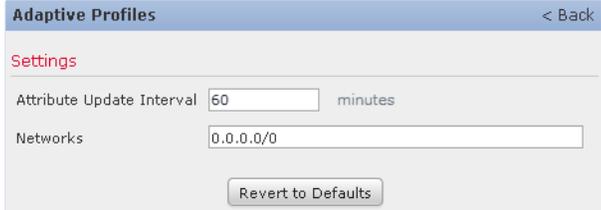
You can also control how frequently network map data is synced from the Defense Center to the managed device. The system uses the data to determine what profiles should be used when processing traffic.

To configure adaptive profiles:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. You have two choices, depending on whether **Adaptive Profiles** under **Detection Enhancement** is enabled:
 - If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Adaptive Profiles page appears.



Adaptive Profiles < Back

Settings

Attribute Update Interval 60 minutes

Networks 0.0.0.0/0

Revert to Defaults

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Optionally, in the **Attribute Update Interval** field, type the number of minutes that should elapse between synchronization of network map data from the Defense Center to the managed device.

IMPORTANT! Increasing the value for **Attribute Update Interval** could improve performance in a large network.

6. In the **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 map for which you want to use adaptive profiles.
See [Managing Variables](#) on page 698 for information on configuring variables in an intrusion policy. See [Creating a Network Discovery Policy](#) on page 1275 for information on configuring the network map.
7. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 27

USING GLOBAL RULE THRESHOLDING

You can use thresholds to limit the number of times the system logs and displays intrusion events. Thresholds 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.

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](#) on page 1009 and [Configuring Global Thresholds](#) on page 1012.
- 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](#) on page 747.

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

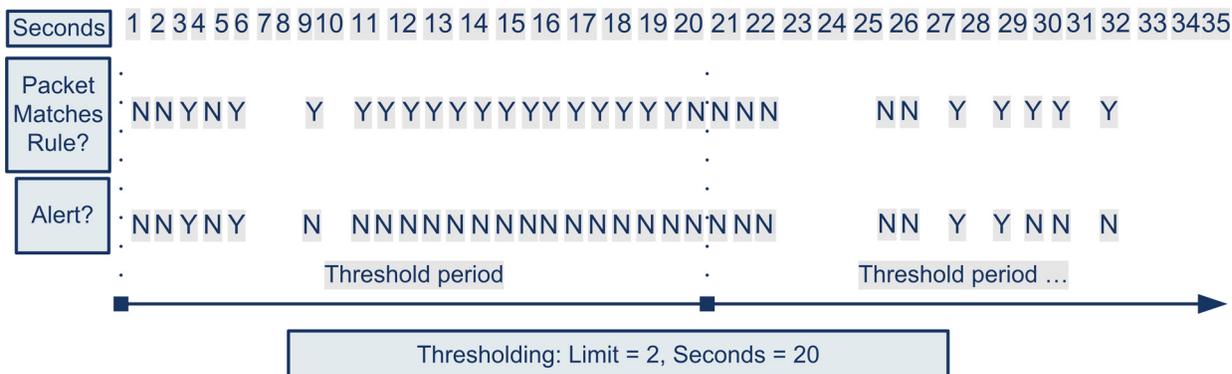
SID 1315. All other rules generate no more than five events in each 60 second period, but the system generates up to ten events for each 60 second period for SID 1315.

For more information on setting rule-based thresholds, see [Configuring Event Thresholding](#) on page 747.

TIP! A global or individual threshold on a managed device with multiple CPUs may result in a higher number of events than expected.

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, first, specify the thresholding type, as described in the following table.

Thresholding Options

OPTION	DESCRIPTION
Limit	<p>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.</p>
Threshold	<p>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.</p>
Both	<p>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:</p> <ul style="list-style-type: none"> • 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.

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](#) on page 1009.

A global threshold is configured by default. The default values are as follows:

- **Type** — Limit
- **Track By** — Destination
- **Count** — 1
- **Seconds** — 60

To configure global thresholding:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
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.



Global Rule Thresholding < Back

Settings

Type Limit Threshold Both

Track By Source Destination

Count

Seconds seconds

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Select the type of threshold from the **Type** drop-down list to do the following during the time specified by the seconds argument:
 - 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.
6. Select the tracking method from the **Track By** drop-down list:
 - 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.

7. You have the following options:
 - For a **Threshold** threshold, specify the number of rule matches you want to use as your threshold 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 in the **Count** field.
8. You have the following options:
 - For a **Limit** threshold, specify the number of seconds that make up the time period for which attacks are tracked in the **Seconds** field.
 - For a **Threshold** threshold specify the number of seconds that elapse before the count resets in the **Seconds** field. Note that the count resets if the number of rule matches indicated by the **Count** field occur before the number of seconds indicated elapse.
9. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. Disable **Global Rule Thresholding** under Intrusion Rule Thresholds.
5. Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 28

USING PERFORMANCE SETTINGS IN AN INTRUSION POLICY

Sourcefire provides several features for improving the performance of your system as it analyzes traffic for attempted intrusions. See the following sections for more information:

- [Event Queue Configuration](#) on page 1016 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.
- [Understanding Packet Latency Thresholding](#) on page 1017 describes how you can balance security with the need to maintain device latency at an acceptable level with packet latency thresholding.
- [Understanding Rule Latency Thresholding](#) on page 1021 describes how you can balance security with the need to maintain device latency at an acceptable level with rule latency thresholding.
- [Performance Statistics Configuration](#) on page 1026 describes how you can configure the basic parameters of how your managed devices monitor and report on their own performance.
- [Constraining Regular Expressions](#) on page 1028 describes how you can override default match and recursion limits on PCRE regular expressions.
- [Rule Processing Configuration](#) on page 1030 describes how you can configure rule processing event queue settings.

Event Queue Configuration

LICENSE: Protection

You can specify the number of packets to allow in the event queue, and enable or disable, before and after stream reassembly, inspection of packets that will be rebuilt into larger streams.

To configure event queue settings:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. You have two choices, depending on whether **Event Queue Configuration** under Performance Settings is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Event Queue Configuration page appears.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can modify the following options:
 - Type a value for the maximum number of events to allow in queue in the **Maximum Queued Events** field.
 - To inspect packets which will be rebuilt into larger streams of data before and after stream reassembly, select **Disable content checks that will be inserted through the stream reassembly process**. 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 that will be inserted through the stream reassembly process**. Disabling inspection decreases the processing overhead for inspection of stream inserts and may boost performance.
6. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Understanding Packet Latency Thresholding

LICENSE: Protection

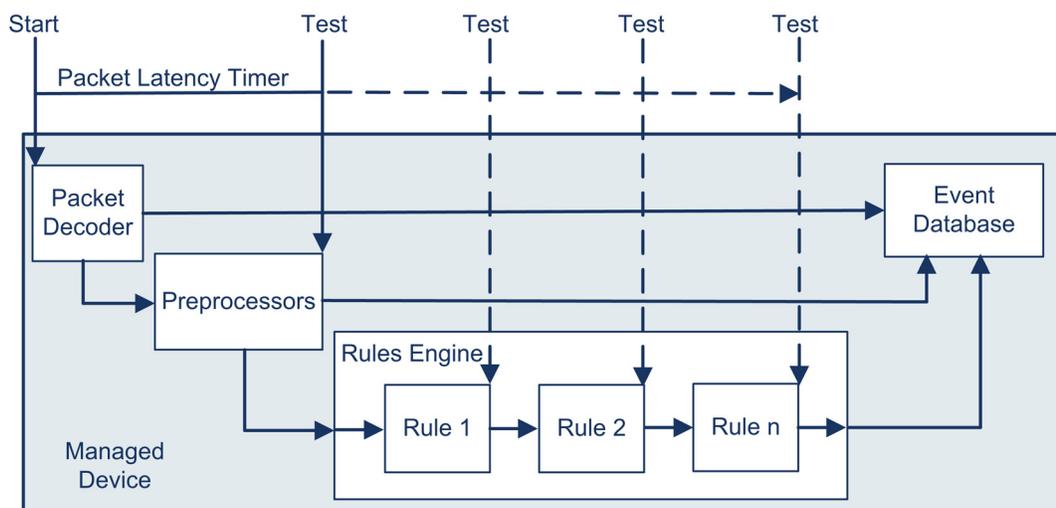
You can balance security with the need to maintain latency at an acceptable level by enabling packet latency thresholding. Packet latency thresholding measures the total elapsed time taken to process a packet by applicable decoders, preprocessors, and rules, and ceases inspection of the packet if the 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.

When you enable packet latency thresholding, 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.

IMPORTANT! 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](#) on page 743.

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.

See the following sections for more information:

- [Setting Packet Latency Thresholding Options](#) on page 1019.
- [Configuring Packet Latency Thresholding](#) on page 1020.

Setting Packet Latency Thresholding Options

LICENSE: Protection

The [Packet Latency Thresholding Options](#) table describes the options you can set to configure packet latency thresholding.

Packet Latency Thresholding Options

OPTION	DESCRIPTION
Threshold	Specifies the time in microseconds when inspection of a packet ceases. See Minimum Packet Latency Threshold Settings on page 1020 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 [Viewing Intrusion Events](#) on page 602 and [Setting Rule States](#) on page 743 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, Sourcefire recommends that, if you enable packet latency thresholding, you use the threshold settings in the [Minimum Packet Latency Threshold Settings](#) table

until your own calculations provide you with settings tailored to your particular network environment.

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, said differently, 4 microseconds per packet. Multiplying by a factor of twenty-five results in a recommended minimum threshold of 100 microseconds.

Configuring Packet Latency Thresholding

LICENSE: Protection

You can enable or disable packet latency thresholding and modify the latency threshold.

To configure packet latency thresholding:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. You have two choices, depending on whether **Latency-Based Packet Handling** under Performance Settings is enabled:
 - If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Latency-Based Packet Handling page appears.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. See the [Minimum Packet Latency Threshold Settings table](#) on page 1020 for recommended minimum **Threshold** settings.
6. Optionally, click **Configure Rules for Latency-Based Packet Handling** at the top of the page to display rules associated with individual options.
Click **Back** to return to the Latency-Based Packet Handling page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

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.

When you enable rule latency thresholding, 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.

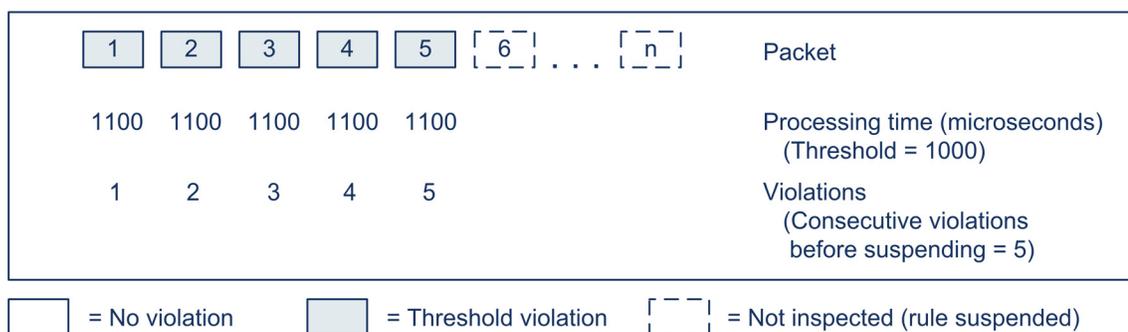
The following example shows five consecutive rule processing times that do not result in rule suspension.

1	2	3	4	5	Packet
1100	1100	1100	500	1100	Processing time (microseconds) (Threshold = 1000)
1	2	3	0	1	Violations (Consecutive violations before suspending = 5)

= No violation = Threshold violation

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](#) on page 743.

IMPORTANT! 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

See the following sections for more information:

- [Setting Rule Latency Thresholding Options](#) on page 1024.
- [Configuring Rule Latency Thresholding](#) on page 1025.

Setting Rule Latency Thresholding Options

LICENSE: Protection

When enabled, 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 [Viewing Intrusion Events](#) on page 602 and [Setting Rule States](#) on page 743 for more information.

The [Rule Latency Thresholding Options](#) table further describes the options you can set to configure rule latency thresholding.

Rule Latency Thresholding Options

OPTION	DESCRIPTION
Threshold	Specifies the time in microseconds that rules should not exceed when examining a packet. See Minimum Rule Latency Threshold Settings on page 1025 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, Sourcefire recommends that, if you enable rule latency thresholding, you use the threshold settings in the

[Minimum Rule Latency Threshold Settings](#) table until your own calculations provide you with settings tailored to your particular network environment.

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.

Configuring Rule Latency Thresholding

LICENSE: Protection

You can enable or disable rule latency thresholding, and modify the rule latency threshold, the suspension time for suspended rules, and the number of consecutive threshold violations that must occur before suspending rules.

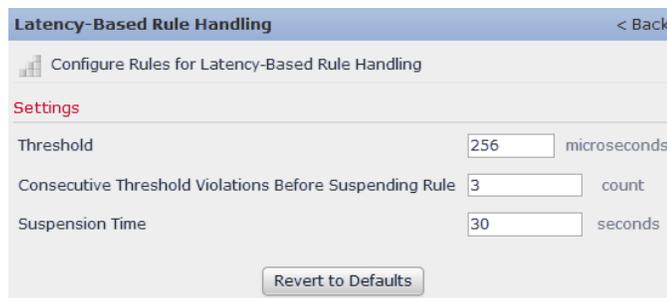
To configure rule latency thresholding:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. You have two choices, depending on whether **Latency-Based Rule Handling** under Performance Settings is enabled:
 - If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Latency-Based Rule Handling page appears.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. See the [Minimum Rule Latency Threshold Settings table](#) on page 1025 for recommended minimum **Threshold** settings.
6. Optionally, click **Configure Rules for Latency-Based Rule Handling** at the top of the page to display rules associated with individual options. Click **Back** to return to the Latency-Based Rule Handling page.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Performance Statistics Configuration

LICENSE: Protection

You can configure the basic parameters of how devices monitor and report on their own performance. This allows you to specify the intervals at which the system updates performance statistics on your devices by configuring the following:

- number of seconds
- number of packets analyzed

WARNING! Do not apply an access control policy that includes an intrusion policy with the Performance Statistics **Log Session/Protocol Distribution** check box selected unless directed to do so by Sourcefire Support.

When the number of seconds specified has elapsed since the last performance statistics update, the system verifies that the specified number of packets has been analyzed. If so, the system updates performance statistics. If not, the system waits until the specified number of packets has been analyzed.

To configure basic performance statistics parameters:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

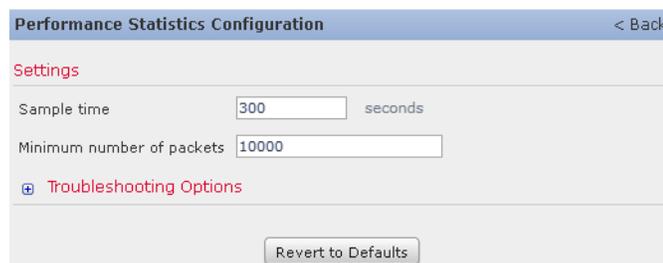
The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

4. Click **Edit** next to **Performance Statistics Configuration** under **Performance Settings**.

The Performance Statistics Configuration page appears.



TIP! You cannot disable the Performance Statistics Configuration advanced setting. This ensures that Sourcefire Support can troubleshoot your system.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Optionally, you can modify any of the performance statistics options:
 - To specify the number of seconds for the system to wait since the last performance statistics update before counting the number of packets that have been analyzed, modify the value for **Sample time**.
 - To specify the number of packets to analyze before updating performance statistics, modify the value for **Minimum number of packets**.

6. Optionally, modify the troubleshooting options only if asked to do so by Sourcefire Support; click the + sign next to **Troubleshooting Options**. See [Understanding Troubleshooting Options](#) on page 790 for more information.

WARNING! Do not apply an access control policy that includes an intrusion policy with the **Log Session/Protocol Distribution** troubleshooting option enabled unless directed to do so by Support.

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 the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Constraining Regular Expressions

LICENSE: Protection

You can override default match and recursion limits on PCRE regular expressions that are used in intrusion rules to examine packet payload content. See [Searching for Content Using PCRE](#) on page 1086 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.

WARNING! Do not override default PCRE limits unless you are an experienced intrusion rule writer with knowledge of the impact of degenerative patterns.

Note that when a rule that requires this feature is enabled in an intrusion policy where this feature is disabled, you must enable the feature or choose to allow the system to enable it automatically before you can save the policy. For more information, see [Automatically Enabling Advanced Settings](#) on page 787.

The [Regular Expression Constraint Options](#) table describes the options you can configure to override the default limits.

Regular Expression Constraint Options

OPTION	DESCRIPTION
Match Limit State	<p>Specifies whether to override Match Limit. You have the following options:</p> <ul style="list-style-type: none"> 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	<p>Specifies the number of times to attempt to match a pattern defined in a PCRE regular expression.</p>
Match Recursion Limit State	<p>Specifies whether to override Match Recursion Limit. You have the following options:</p> <ul style="list-style-type: none"> 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 <p>Note that for Match Recursion Limit to be meaningful, it must be smaller than Match Limit.</p>
Match Recursion Limit	<p>Specifies the number of recursions when evaluating a PCRE regular expression against the packet payload.</p>

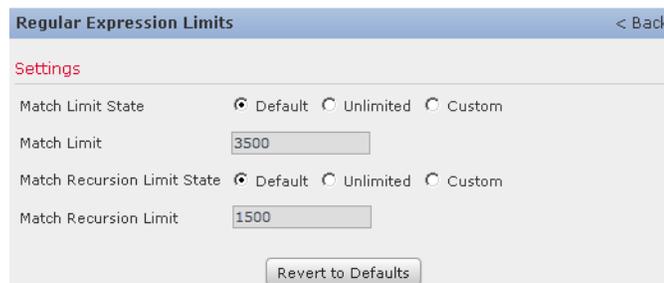
To configure PCRE overrides:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
 The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
 The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
4. You have two choices, depending on whether **Regular Expression Limits** under Performance Settings is enabled:
 - If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Regular Expression Limits page appears.



A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can modify any of the options in the [Regular Expression Constraint Options table](#) on page 1029.
6. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Rule Processing Configuration

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 [Rule Processing Configuration Options](#) table describes the options you can configure to determine how many events are logged per packet or stream.

Rule Processing Configuration Options

OPTION	DESCRIPTION
Maximum Queued Events	The maximum number of events that can be stored for a given packet or packet stream.
Logged Events	The number of events logged for a given packet or packet stream. This cannot exceed the Max Events value.
Order Events 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: <ul style="list-style-type: none">• 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.

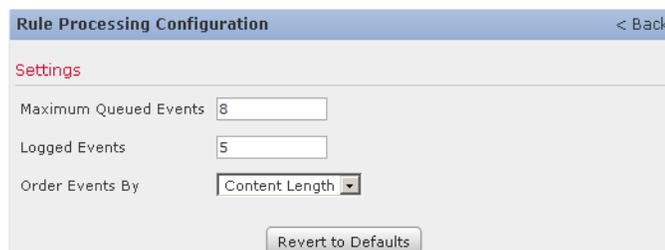
To configure how many events are logged per packet or stream:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.

4. You have two choices, depending on whether **Rule Processing Configuration** under Performance Settings is enabled:
 - If the configuration is enabled, click **Edit**.
 - If the configuration is disabled, click **Enabled**, then click **Edit**.

The Rule Processing Configuration page appears.



The screenshot shows a web interface for configuring rule processing. It has a title bar 'Rule Processing Configuration' with a back arrow. Below is a 'Settings' section with three rows: 'Maximum Queued Events' (input field with '8'), 'Logged Events' (input field with '5'), and 'Order Events By' (dropdown menu with 'Content Length'). A 'Revert to Defaults' button is at the bottom.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. You can modify any of the options on the Rule Processing Configuration page. See the [Rule Processing Configuration Options table](#) on page 1031 for a full description of each available option.
6. 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 [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

CHAPTER 29

CONFIGURING EXTERNAL RESPONSES TO INTRUSION EVENTS

While the Sourcefire 3D System provides various views of intrusion events within the web interface, some enterprises prefer to define external intrusion event notification to facilitate constant monitoring of critical systems. If you want to immediately notify a specific person of critical events, you can set up email alerts to do so. You can also 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](#) on page 746 for more information.

There is another type of alerting you can perform in the Sourcefire 3D System, outside of your intrusion policies. You can configure email, SNMP, and syslog alert responses for other types of events, including intrusion events with specific impact flags, or connection events logged by specific access control rules. For more information, see [Configuring External Alerting](#) on page 524.

See the following sections for more information on external intrusion event notification:

- [Using SNMP Responses](#) on page 1034 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](#) on page 1038 describes the options you can configure to send event data to an external syslog and provides the procedure for specifying the syslog alerting options.
- [Understanding Email Alerting](#) on page 1042 describes the options you can configure to send notifications of intrusion events by email.

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 [Modifying Advanced Settings](#) on page 774.

TIP! If your network management system requires a management information base file (MIB), you can obtain it from the Defense Center at `/etc/sf/DCEALERT.MIB`.

SNMP v2 Options

For SNMP v2, you can specify the options described in the [SNMP v2 Options](#) table.

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 [SNMP v3 Options](#) table.

IMPORTANT! 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.

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.
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](#) on page 1037.

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](#) on page 1034 and [Applying an Access Control Policy](#) on page 466.

To configure SNMP alerting options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.

The Intrusion Policy page appears.

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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.

The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

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.



SNMP Alerting < Back

Settings

Trap Type as Binary as String

SNMP Version Version2 Version3

SNMP v3

Trap Server

Authentication Password

Private Password (SNMP v3 passwords must be 8 or more characters)

Username

Revert to Defaults

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

5. Specify the trap type format that you want to use for IP addresses that appear in the alerts, **as Binary** or **as String**.

IMPORTANT! 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.

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](#) on page 1035.
 - 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](#) on page 1036 for more information.

IMPORTANT! You must select SNMP v2 **or** SNMP v3.

IMPORTANT! When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format.

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 the [Common Intrusion Policy Editing Actions table](#) on page 675 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 [Available Syslog Facilities](#) 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.

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.

Available Syslog Facilities (Continued)

FACILITY	DESCRIPTION
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:

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](#) on page 1038 and [Applying an Access Control Policy](#) on page 466.

To configure syslog alerting options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Intrusion Policy**.
The Intrusion Policy page appears.
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 [Committing Intrusion Policy Changes](#) on page 679 for information on saving unsaved changes in another policy.
The Policy Information page appears.
3. Click **Advanced Settings** in the navigation panel on the left.
The Advanced Settings page appears.
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.

The screenshot shows the 'Syslog Alerting' configuration interface. It includes a 'Logging Hosts' text field, 'Facility' and 'Priority' dropdown menus, and a 'Revert to Defaults' button.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See [Using Layers in an Intrusion Policy](#) on page 792 for more information.

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.
6. Select facility and priority levels from the drop-down lists.
See [Using Syslog Responses](#) on page 1038 for details on facility and priority options.
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 the [Common Intrusion Policy Editing Actions table](#) on page 675 for more information.

Understanding Email Alerting

LICENSE: Protection

Email alerts are notifications of intrusion events by email. Email alerts include the following information:

- total number of alerts in the database
- last email time (the time that the system generated the last email report)
- current time (the time that the system generated the current email report)
- total number of new alerts
- number of events that matched specified email filters (if events are configured for specific rules)
- timestamp, protocol, event message, and session information (source and destination IPs and ports with traffic direction) for each event (if Summary Output is off)

IMPORTANT! If multiple intrusion events originate from the same source IP, a note appears beneath the event that displays the number of additional events.

- number of events per destination port
- number of events per source IP

For each rule or rule group, you can enable or disable email alerting on intrusion events. Your email alert settings are used regardless of which intrusion policy you apply to the device as part of an access control policy.

The following list describes the parameters you can set for email alerting.

On/Off

Enables or disables email notification.

From Address

Specifies the email address or addresses from which the system sends intrusion events.

To Address

Specifies the email address where the system sends intrusion events. To send email to multiple recipients, separate email addresses with commas. For example:

`user1@example.com, user2@example.com`

Max Alerts

Specifies the maximum number of intrusion events the system sends via email in the time frame specified by Frequency (seconds).

Frequency (seconds)

Specifies how often the system mails intrusion events. The Frequency setting also specifies the frequency with which email settings are saved.

Minimum frequency: 300 seconds

Maximum frequency: 4 billion seconds

Coalesce Alerts

Enables or disables grouping of intrusion events by source IP and event so that multiple identical intrusion events generated against the same source IP only present one event on the page.

Note that alert coalescence (grouping) occurs after events are filtered. Therefore, if you configure email alerting on specific rules, you will only receive a list of events that match the rules you specified in the Mail Alerting Configuration.

Summary Output

Enables or disables brief email alerting, which is suitable for text-limited devices such as pagers. Brief email alerts contain:

- event timestamp
- for Defense Centers, the IP address for the device that generated the event
- event protocol
- source IP and port
- destination IP and port
- event message
- the number of intrusion events generated against the same source IP

For example:

```
2011-05-18 10:35:10 10.1.1.100 icmp 10.10.10.1:8 -> 10.2.1.3:0  
snort_decoder: Unknown Datagram decoding problem! (116:108)
```

Email Alerting on Specific Rules Configuration

Specifies the rules or rule groups whose events you want mailed to the specified email address or addresses.

For information about configuring email alerting, see [Configuring Email Alerting](#) on page 1044.

Configuring Email Alerting

LICENSE: Protection

You can configure email alerting so that your appliance notifies you whenever an intrusion event occurs for an specific rule or rule group.

Before you can receive email alerts, you **must**:

- configure your mail host to receive email alerts (see [Configuring a Mail Relay Host and Notification Address](#) on page 1983)
- make sure that both the managed device and the Defense Center can reverse resolve their own IP addresses

To configure email alerting options:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Email**.

The Email Alerting page appears.

Email Alerting

State on off

From Address

To Address

Max Alerts

Min Frequency Note: Frequency is set by number of seconds

Coalesce Alerts on off

Summary Output on off

Time Zone

Rules Email Alerting per Rule Configuration
 Select All

2. Next to **State**, select **on** to enable email alerting.
3. In the **From Address** field, type the address you want to display in the From field in the email alerts.
4. In the **To Address** field, type the address where you want to receive the email alerts.
5. In the **Max Alerts** field, type the maximum number of events you want included in a single email.
6. In the **Min Frequency** field, type the number of seconds for the minimum frequency with which you want to receive email alerts.
7. To group events by IP address, next to **Coalesce Alerts**, select **on**.

- To send brief email alerts, next to **Summary Output**, select **on**.

TIP! If you enable **Summary Output**, consider enabling **Coalesce Alerts** to reduce the number of alerts generated. Also consider setting **Max Alerts** to 1 to avoid overflowing your device's text message buffer.

- In the **Time Zone** field, select your time zone from the drop-down list.
- To enable email alerting per rule, click **Email Alerting per Rule Configuration**. The rule groups appear.

TIP! To receive email alerts for all rules in all categories, select **Select All**.

- Perform one or both of the following:
 - Click **All** next to rule categories for which you want to receive email alerts for all rules belonging to the category.
 - Click the category folder within which you want to specify email alerting on individual rules in that category, then enable the rules for which you want to receive email alerts.
- Click **Save**.

The system saves your email alerting configuration. When applicable intrusion events occur, you receive email alerts.

CHAPTER 30

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 Defense Center web interface.

WARNING! 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 your Sourcefire 3D 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](#) on page 743.
- Sourcefire provides two types of intrusion rules: shared object rules and standard text rules. The Sourcefire 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 [Working with Intrusion Events](#) on page 593.

See the following sections for more information:

- [Understanding Rule Anatomy](#) on page 1047 describes the components, including the rule header and rule options, that make up a valid standard text rule.
- [Understanding Rule Headers](#) on page 1049 provides a detailed description of the parts of a rule header.
- [Understanding Keywords and Arguments in Rules](#) on page 1055 explains the usage and syntax of the intrusion rule keywords available in the Sourcefire 3D System.
- [Constructing a Rule](#) on page 1181 explains how to build a new rule using the rule editor.
- [Searching for Rules](#) on page 1189 explains how to search for existing rules.
- [Filtering Rules on the Rule Editor Page](#) on page 1192 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:

Rule Header

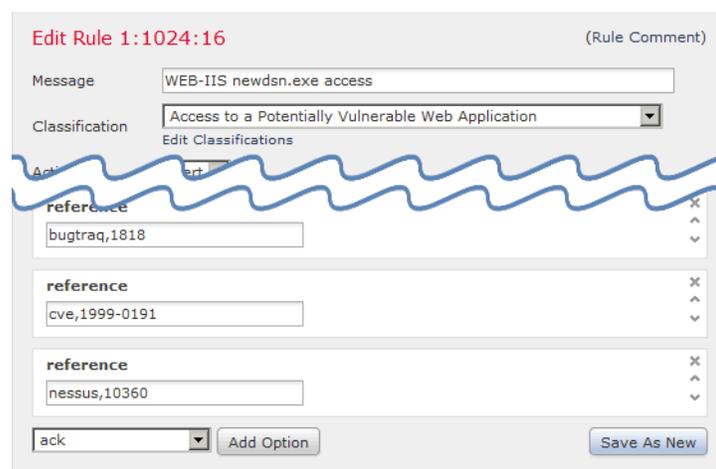
```
alert tcp $EXTERNAL_NET any -> $HTTP_SERVERS $HTTP_PORTS
```

Rule Keywords and Arguments

```
(msg:"WEB-IIS newdsn.exe access";  
flow:to_server,established; uricontent:"/scripts/  
tools/newdsn.exe"; nocase; metadata:service http;  
reference:bugtraq,1818; reference:cve,1999-0191;  
reference:nessus,10360; classtype:web-application-  
activity; sid:1024; rev:10; )
```

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.

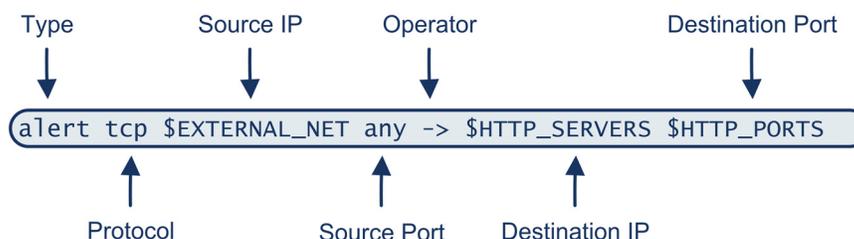
The following displays the same standard text rule within the rule editor:



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 [Rule Header Values](#) table describes each part of the rule header shown above.

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.

IMPORTANT! The previous example uses system variables, as do most intrusion rules. See [Managing Variables](#) on page 698 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](#) on page 1050 describes rule types and explains how to specify the action that occurs when the rule triggers.
- [Specifying Protocols](#) on page 1051 explains how to define the traffic protocol for traffic that the rule should test.
- [Specifying Source and Destination IP Addresses](#) on page 1051 explains how to define the individual IP addresses and IP ranges in the rule header.
- [Specifying Source and Destination Ports](#) on page 1053 explains how to define the individual ports and port ranges in the rule header.
- [Specifying Direction](#) on page 1055 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.

IMPORTANT! 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](#) on page 743.

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](#) on page 1181.

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)

IMPORTANT! The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`. For more information, see [Specifying Source and Destination Ports](#) on page 1053.

- 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.

IMPORTANT! 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](#) on page 1181 for more information about the procedures you use to build a rule header using the rule editor.

Specifying Source and Destination IP Addresses

LICENSE: Protection

Rule headers include the source and destination IP addresses of traffic that the rule inspects. You can restrict packet inspection to the packets originating from specific IP addresses or those destined to a specific IP address. This 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.

WARNING! The system recognizes only IP addresses and will not accept host names for source or destination IP addresses.

The [Source/Destination IP Address Syntax](#) table describes the various ways you can specify source and destination IP addresses.

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 range 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 range of IP addresses except one or more specific IP addresses	a range of addresses followed by a list of negated addresses or ranges	[10.0.0/8, !10.2.3.4, !10.1.0.0/16] [2001:db8::/32, !2001:db8::8329, !2001:db8::0202]
IP addresses defined by an IP address variable	the variable name, in uppercase letters, preceded by \$ See Managing Variables on page 698 for more information.	\$HOME_NET

Source/Destination IP Address Syntax (Continued)

To SPECIFY...	USE...	EXAMPLE
all IP addresses except addresses defined by an IP address variable	the variable name, in uppercase letters, preceded by !\$ See Managing Variables on page 698 for more information.	!\$HOME_NET
IP addresses defined by a network object or network object group	the object or group name using the format <code>!{object_name}</code> . See Working with Network Objects on page 161 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 on page 161 for more information.	!\${192.168sub16}

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, see [Defining IP Addresses in Variables and Rules](#) on page 713, [Managing Variables](#) on page 698, and [IP Address Conventions](#) on page 57.

Within the rule editor, you specify source and destination IP addresses in the **Source IP** and **Destination IP** fields. See [Constructing a Rule](#) on page 1181 for more information about the procedures you use to build a rule header using the rule editor.

Specifying Source and Destination Ports

LICENSE: Protection

Rule headers specify the source and destination ports of the traffic the rule inspects. You can restrict packet inspection to packets originating from or destined to specific ports. This reduces the amount of packet inspection the system must perform. This also reduces false positives, making the rule more specific and removing the possibility for the rule to trigger against packets whose source and destination ports do not indicate suspicious behavior.

IMPORTANT! The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`. For more information, see [Specifying Protocols](#) on page 1051.

The [Source/Destination Port Syntax](#) table describes the different ways you can specify source and destination ports.

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 <i>any</i> , which if negated would indicate <i>no port</i> .	!20
all ports defined by a port variable	the variable name, in uppercase letter, preceded by \$ See Managing Variables on page 698 for more information.	\$HTTP_PORTS
all ports except ports defined by a port variable	the variable name, in uppercase letter, preceded by !\$ See Managing Variables on page 698 for more information.	!\$HTTP_PORTS

You can currently substitute a colon (:) for the hyphen (-) in any port definition in an rule header, although eventually the colon will be deprecated.

You can create a list of ports up to 64 characters in length by surrounding the port list with brackets and separating the ports with commas, as shown in the following example:

[80, 8080, 8138, 8600-9000, !8650-8675]

For information about using variables to specify ports, see [Defining Ports in Variables and Rules](#) on page 715.

Within the rule editor, you specify source and destination ports in the **Source Port** and **Destination Port** fields. See [Constructing a Rule](#) on page 1181 for more information about the procedures you use to build a rule header using the rule editor.

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 [Directional Options in Rule Headers](#) table describes these options.

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](#) on page 1181 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 Sourcefire 3D System 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 [Using Adaptive Profiles](#) on page 1003.

See the following sections for more information:

- [Defining Intrusion Event Details](#) on page 1057 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](#) on page 1063 describes how to use the `content` keyword to test the content of the packet payload.
- [Constraining Content Matches](#) on page 1065 describes how to use modifying keywords for the `content` keyword.
- [Replacing Content in Inline Deployments](#) on page 1078 describes how to use the `replace` keyword in inline deployments to replace specified content of equal length.
- [Using Byte_Jump and Byte_Test](#) on page 1079 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](#) on page 1086 describes how to use the `pcr` keyword to use Perl-compatible regular expressions in rules.
- [Adding Metadata to a Rule](#) on page 1096 describes how to use the `metadata` keyword to add information to a rule.
- [Inspecting IP Header Values](#) on page 1101 describes the syntax and use of keywords that test values in the packet's IP header.
- [Inspecting ICMP Header Values](#) on page 1104 describes the syntax and use of keywords that test values in the packet's ICMP header.
- [Inspecting TCP Header Values and Stream Size](#) on page 1106 describes the syntax and use of keywords that test values in the packet's TCP header.
- [Enabling and Disabling TCP Stream Reassembly](#) on page 1112 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](#) on page 1113 describes the use and syntax of keywords that extract version and state information from encrypted traffic.
- [Reading Packet Data into Keyword Arguments](#) on page 1156 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](#) on page 1116 describes the use and syntax of keywords that test application layer protocol properties.
- [Inspecting Packet Characteristics](#) on page 1153 describes the use and syntax of the `dsize`, `sameIP`, `isdataat`, `fragoffset`, and `cvs` keywords.

- [Initiating Active Responses with Rule Keywords](#) on page 1160 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](#) on page 1165 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](#) on page 1166 describes how to log additional traffic for the host or session.
- [Detecting Attacks That Span Multiple Packets](#) on page 1168 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](#) on page 1175 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.
- [Pointing to a Specific Payload Type](#) on page 1177 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](#) on page 1178 describes how to point to the beginning of the packet payload.
- [Decoding and Inspecting Base64 Data](#) on page 1179 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](#) on page 1058
- [Defining the Event Priority](#) on page 1058
- [Defining the Intrusion Event Classification](#) on page 1058
- [Defining the Event Reference](#) on page 1062

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](#) on page 1181 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](#) on page 1181 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 [Rule Classifications](#) table lists the name and number for each classification.

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

Rule Classifications (Continued)

NUMBER	CLASSIFICATION NAME	DESCRIPTION
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

Rule Classifications (Continued)

NUMBER	CLASSIFICATION NAME	DESCRIPTION
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](#) on page 1182 for more information on the rule editor.

Adding Custom Classifications

LICENSE: Protection

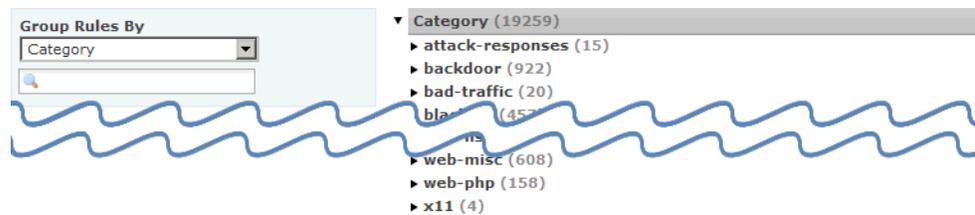
If you want more customized content for the packet display description of the events generated by a rule you define, create a custom classification.

To add classifications to the **Classification** list:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.



2. Click **Create Rule**.

The Create Rule page appears.

3. Under the **Classification** drop-down list, click **Edit Classifications**.

A pop-up window appears.

Default Classification	Description	Priority	Custom Classification	Description	Priority
attempted-admin	Attempted Administrator Privilege Gain	high	No Custom Classifications		
attempted-dos	Attempted Denial of Service	medium			
web-application-activity	Potentially Vulnerable Web Application	medium			
web-application-attack	Web Application Attack	high			

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.
5. Type a description of the classification in the **Classification Description** field.
 You can use up to 255 alphanumeric characters and spaces.
6. Select a priority from the **Priority** list.
 You can select **high**, **medium**, or **low**.
7. Click **Add**.
 The new classification is added to the list and becomes available for use in the rule editor.

New Classification

Classification Name

Classification Description

Priority high

Default Classification	Description	Priority	Custom Classification	Description	Priority
attempted-admin	Attempted Administrator Privilege Gain	high	TestClassification	test classification	high
attempted-dos	Attempted Denial of Service	medium			
web-application-activity	Potentially Vulnerable Web Application	medium			
web-application-attack	Web Application Attack	high			

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 [External Attack Identification Systems](#) table lists some of the external systems that can provide data on known exploits and attacks.

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](#) on page 1181 for more information about using the rule editor to build rules.

Searching for Content Matches

LICENSE: Protection

Use the **content** keyword to specify content that you want to detect in a packet. The rules engine searches the packet payload or stream for that string. For

example, if you enter `/bin/sh` as the value for the `content` keyword, 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.

You should almost always follow a `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` keyword.

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](#) on page 1078 for more information.

To enter content to be matched:

ACCESS: Admin/Intrusion Admin

1. In the `content` field, type the content you want to find (for example, `|90C8 C0FF FFFF|/bin/sh`).

If you want to specify that the rule search for any content that is **not** the specified content, check the **Not** box.

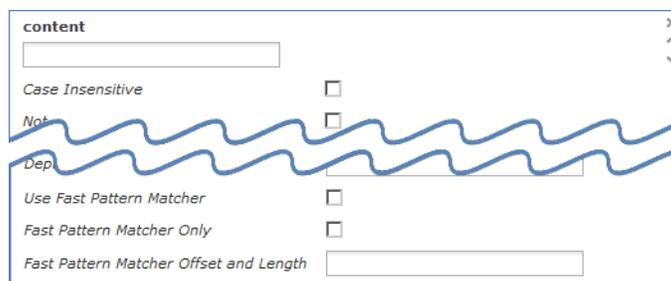
WARNING! You could invalidate your intrusion policy if you create a rule that includes only one `content` keyword and that keyword has the **Not** option selected. For more information, see [Not](#) on page 1066.

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](#) on page 1055. For more information on constraining the `content` keyword, see [Constraining Content Matches](#) on page 1065.
3. Continue with creating or editing the rule. See [Writing New Rules](#) on page 1182 or [Modifying Existing Rules](#) on page 1184 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** keyword. Configure options that modify the **content** keyword to specify the content for which you want to search.



For more information, see the following sections:

- [Case Insensitive](#) on page 1065
- [Raw Data](#) on page 1065
- [Not](#) on page 1066
- [Search Location Options](#) on page 1067
- [HTTP Content Options](#) on page 1069
- [Use Fast Pattern Matcher](#) on page 1074

Case Insensitive

LICENSE: Protection

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:

ACCESS: Admin/Intrusion Admin

1. Select **Case Insensitive** for the **content** keyword you are adding.
2. Continue with creating or editing the rule. See [Constraining Content Matches](#), [Searching for Content Matches](#) on page 1063, [Writing New Rules](#) on page 1182 or [Modifying Existing Rules](#) on page 1184 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 (data decoded by a Sourcefire 3D System preprocessor) 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** keyword with any HTTP content option. See [HTTP Content Options](#) on page 1069 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, when the HTTP Inspect preprocessor is enabled. For more information, see [Selecting Server-Level HTTP Normalization Options](#) on page 854.

To analyze raw data:

ACCESS: Admin/Intrusion Admin

1. Select the **Raw Data** check box for the **content** keyword you are adding.
2. Continue with creating or editing the rule. See [Constraining Content Matches](#), [Searching for Content Matches](#) on page 1063, [Writing New Rules](#) on page 1182, or [Modifying Existing Rules](#) on page 1184 for more information.

Not

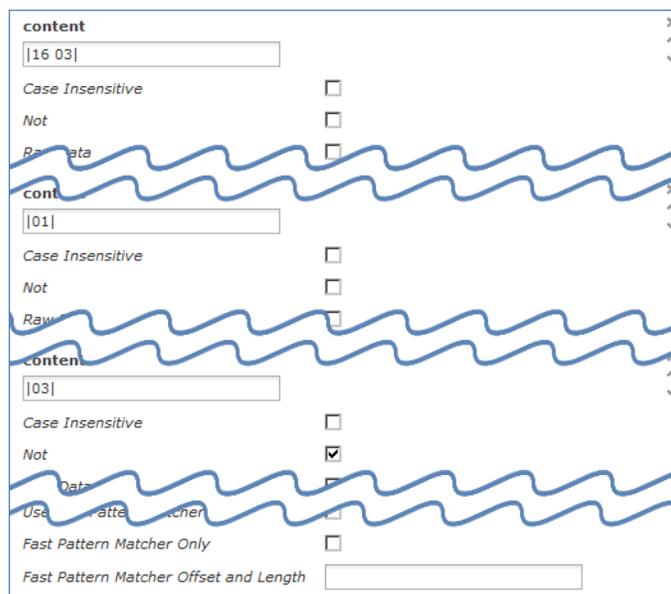
LICENSE: Protection

Select the **Not** option to search for content that does not match the specified content. If you create a rule that includes a **content** keyword with the **Not** option selected, you must also include in the rule at least one other **content** keyword without the **Not** option selected.

WARNING! Do not create a rule that includes only one **content** keyword if that keyword has the **Not** option selected. You could invalidate your intrusion policy. For more information, see [Configuring Intrusion Policies](#) on page 667.

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:

ACCESS: Admin/Intrusion Admin

1. Select the **Not** check box for the **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.

2. Include in the rule at least one other **content** keyword that does not have the **Not** option selected.
3. Continue with creating or editing the rule. See [Constraining Content Matches](#), [Searching for Content Matches](#) on page 1063, [Writing New Rules](#) on page 1182, or [Modifying Existing Rules](#) on page 1184 for more information.

Search Location Options

LICENSE: Protection

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](#) on page 1156 for more information.

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.

Depth

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.

Within

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.

To specify a search location value in the user interface:

ACCESS: Admin/Intrusion Admin

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.

2. Continue with creating or editing the rule. See [Constraining Content Matches](#) on page 1065, [Searching for Content Matches](#) on page 1063, [Writing New Rules](#) on page 1182 or [Modifying Existing Rules](#) on page 1184 for more information.

HTTP Content Options

LICENSE: Protection

HTTP **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](#) on page 1069 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](#) on page 1069 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 improve performance and reduce false positives, ensure that the HTTP Inspect preprocessor is enabled so HTTP message traffic can be normalized and evaluated against rules that include HTTP **content** options.
- 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** 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** 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** 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** 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 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](#) on page 1074 for more information.

The above restrictions are reflected in the description of each option in the following list describing the HTTP **content** keyword options.

Note that the HTTP preprocessor must be enabled to allow processing of rules using any of these **content** keyword options. When the HTTP preprocessor is disabled and you enable rules that use any of these keywords, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

The following list describes the HTTP **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](#) on page 1092 for more information.

IMPORTANT! 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](#) on page 1092 for more information.

IMPORTANT! 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](#) on page 1092 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 **Inspect 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](#) on page 1092 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](#) on page 854 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](#) on page 1092 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](#) on page 854 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](#) on page 1092 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](#) on page 854 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](#) on page 854 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](#) on page 854 for more information.

To specify an HTTP content option when doing a content search of TCP traffic:

ACCESS: Admin/Intrusion Admin

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** keyword you are adding; also, optionally, select the **HTTP Cookie** or **HTTP Raw Cookie** option.
2. Continue with creating or editing the rule. See [Constraining Content Matches](#) on page 1065, [Searching for Content Matches](#) on page 1063, [Writing New Rules](#) on page 1182, or [Modifying Existing Rules](#) on page 1184 for more information.

Use Fast Pattern Matcher

LICENSE: Protection

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](#) on page 1069 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](#) on page 1179 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:

ACCESS: Admin/Intrusion Admin

1. Select **Use Fast Pattern Matcher** for the content keyword you are adding.
2. Optionally, select **Fast Pattern Matcher Only** to determine without rules engine evaluation if the specified pattern exists in the packet.
Evaluation will proceed only if the fast pattern matcher detects the specified content.

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.

4. Continue with creating or editing the rule. See [Constraining Content Matches](#) on page 1065, [Searching for Content Using PCRE](#) on page 1086, [Writing New Rules](#) on page 1182, or [Modifying Existing Rules](#) on page 1184 for more information.

Replacing Content in Inline Deployments

LICENSE: Protection

You can use the `replace` keyword in an inline deployment to replace specified content.

IMPORTANT! You cannot use the `replace` keyword to replace content in SSL traffic detected by the Sourcefire SSL Appliance. The original encrypted data, not the replacement data, will be transmitted. See the *Sourcefire 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.

Optionally, you can enclose the replacement string in quotation marks for backward compatibility with previous Sourcefire 3D System 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.

IMPORTANT! 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.

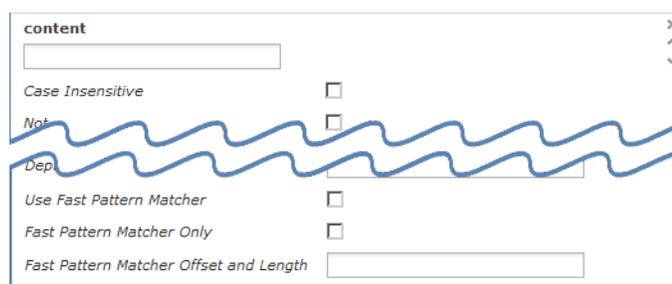
Note that you cannot use the `replace` keyword in combination with HTTP request message `content` keyword options. See [Searching for Content Matches](#) on page 1063 and [HTTP Content Options](#) on page 1069 for more information.

To replace content in an inline deployment:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **content** in the drop-down list and click **Add Option**.

The `content` keyword appears.



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.

3. Select **replace** in the drop-down list and click **Add Option**.

The `replace` keyword appears beneath the `content` keyword.



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](#) on page 810 and [DCE/RPC Keywords](#) on page 1119 for more information.

See the following sections for more information:

- [byte_jump](#) on page 1080
- [byte_test](#) on page 1083

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 [Required byte_jump Arguments](#) table describes the arguments required by the `byte_jump` keyword.

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 <code>offset</code> counter starts at byte 0, so calculate the <code>offset</code> 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 <code>byte_extract</code> variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 1156 for more information.

The [Additional Optional byte_jump Arguments](#) table describes options you can use to define how the system interprets the values you specified for the required arguments.

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.

Additional Optional byte_jump Arguments (Continued)

ARGUMENT	DESCRIPTION
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 [Endianness Arguments](#) table (if neither argument is specified, network byte order is used).

Endianness Arguments

ARGUMENT	DESCRIPTION
Big Endian	Processes data in big endian byte order, which is the default network byte order.

Endianness Arguments (Continued)

ARGUMENT	DESCRIPTION
Little Endian	Processes data in little endian byte order.
DCE/RPC	<p>Specifies a <code>byte_jump</code> keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 810 for more information.</p> <p>The DCE/RPC preprocessor determines big endian or little endian byte order, and the Number Type, Endian, and From Beginning arguments do not apply.</p> <p>When you enable this argument, you can also use <code>byte_jump</code> in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords on page 1119 for more information.</p> <p>The DCE/RPC preprocessor must be enabled to allow processing of rules that include this option. When the DCE/RPC preprocessor is disabled and you enable rules that use this option, you are prompted whether to enable the preprocessor when you save the policy. See Automatically Enabling Advanced Settings on page 787.</p>

Define how the system views string data in a packet by using one of the arguments in the [Number Type Arguments](#) table.

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`:

ACCESS: Admin/Intrusion Admin

- ▶ 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 [Required byte_test Arguments](#) table describes the required arguments for the `byte_test` keyword.

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 <code><</code> , <code>></code> , <code>=</code> , <code>!</code> , <code>&</code> , <code>^</code> , <code>!></code> , <code>!<</code> , <code>!=</code> , <code>!&</code> , or <code>!^</code> . For example, if you specify <code>!1024</code> , <code>byte_test</code> 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 <code>!</code> and <code>!=</code> are equivalent. You can also use an existing <code>byte_extract</code> variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 1156 for more information.
Offset	The number of bytes into the payload to start processing. The <code>offset</code> counter starts at byte 0, so calculate the <code>offset</code> 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 <code>byte_extract</code> variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 1156 for more information.

You can further define how the system uses `byte_test` arguments with the arguments described in the [Additional Optional byte_test Arguments](#) table.

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 [Endianness byte_test Arguments](#) table. If neither argument is specified, network byte order is used.

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	<p>Specifies a <code>byte_test</code> keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 810 for more information.</p> <p>The DCE/RPC preprocessor determines big endian or little endian byte order, and the Number Type and Endian argument do not apply.</p> <p>When you enable this argument, you can also use <code>byte_test</code> in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords on page 1119 for more information.</p> <p>The DCE/RPC preprocessor must be enabled to allow processing of rules that include this option. When the DCE/RPC preprocessor is disabled and you enable rules that use this option, you are prompted whether to enable the preprocessor when you save the policy. See Automatically Enabling Advanced Settings on page 787.</p>

You can define how the system views string data in a packet by using one of the arguments in the [Number Type byte-test Arguments](#) table.

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`:

ACCESS: Admin/Intrusion Admin

- ▶ 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.

The screenshot shows a configuration window for the `byte_test` keyword. It contains the following fields and controls:

- Bytes:** A text input field.
- Offset:** A text input field.
- Value:** A dropdown menu with a right-pointing arrow, followed by a text input field.
- Number Type:** A dropdown menu with a right-pointing arrow.
- Endian:** A dropdown menu with a right-pointing arrow.
- Relative:** A checkbox.
- DCE/RPC:** A checkbox.

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.
- `i smxAEGRBUIPHDMCKSY` 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.

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? i smxAEGRBUIPHDMCKSY` where `regex` is your Perl-compatible regular expression and `i smxAEGRBUIPHDMCKSY` is any combination of modifier options. See [Perl-Compatible Regular Expression Basics](#) on page 1088 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](#) on page 1088 describes the common syntax used in Perl-compatible regular expressions.
- [PCRE Modifier Options](#) on page 1090 describes the options you can use to modify your regular expression.
- [Example PCRE Keyword Values](#) on page 1094 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 [PCRE Metacharacters](#) table describes the metacharacters you can use with PCRE and gives examples of each.

PCRE Metacharacters

METACHARACTER	DESCRIPTION	EXAMPLE
.	Matches any character except newlines. If <code>s</code> is used as a modifying option, it also includes newline characters.	<code>abc.</code> matches <code>abcd</code> , <code>abc1</code> , <code>abc#</code> , and so on.
*	Matches zero or more occurrences of a character or expression.	<code>abc*</code> matches <code>abc</code> , <code>abcc</code> , <code>abccc</code> , <code>abccccc</code> , and so on.
?	Matches zero or one occurrence of a character or expression.	<code>abc?</code> matches <code>abc</code> .
+	Matches one or more occurrences of a character or expression.	<code>abc+</code> matches <code>abc</code> , <code>abcc</code> , <code>abccc</code> , <code>abccccc</code> , and so on.
()	Groups expressions.	<code>(abc)+</code> matches <code>abc</code> , <code>abcabc</code> , <code>abcabcabc</code> 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.	<code>a{4,6}</code> matches <code>aaaa</code> , <code>aaaaa</code> , or <code>aaaaaa</code> . <code>(ab){2}</code> matches <code>abab</code> .

PCRE Metacharacters (Continued)

METACHARACTER	DESCRIPTION	EXAMPLE
[]	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.
^	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 on page 1089 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](#) on page 1088), 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 [PCRE Character Classes](#) table describes and provides examples of the predefined character classes accepted by PCRE.

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]

PCRE Character Classes (Continued)

CHARACTER CLASS	DESCRIPTION	CHARACTER CLASS DEFINITION
\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\n\t\f]
\S	Matches anything that is not a white space character.	[^\r\n\t\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/i smxAEGRBUIPHDMCKSY`

where `i smxAEGRBUPHMC` 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/i smxAEGRBUIPHDMCKSY"`. 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 [Perl-Related Post Regular Expression Options](#) table describes options you can use to perform Perl processing functions.

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 [PCRE-Related Post Regular Expression Options](#) table describes the PCRE modifiers you can use after the regular expression.

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 [Snort-Specific Post Regular Expression Modifiers](#) table describes the Snort-specific modifiers that you can use after the regular expression.

The HTTP preprocessor must be enabled to allow processing of rules using the C, H, U, M, or P expression modifiers. When the HTTP preprocessor is disabled and you enable rules that use these modifiers, you are prompted whether to enable

the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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 keyword).
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 keyword HTTP URI option to search the same content. See HTTP Content Options on page 1069 for more information. IMPORTANT! 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 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 keyword HTTP Raw URI option to search the same content. See HTTP Content Options on page 1069 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 keyword HTTP Client Body option in HTTP Content Options on page 1069 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 keyword HTTP Header option to search the same content. See HTTP Content Options on page 1069 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 keyword HTTP Raw Header option to search the same content. See HTTP Content Options on page 1069 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 keyword HTTP Method option in HTTP Content Options on page 1069 for more information.

Snort-Specific Post Regular Expression Modifiers (Continued)

OPTION	DESCRIPTION
C	<p>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.</p> <p>Note the following:</p> <ul style="list-style-type: none">• Cookies included in the message body are treated as body content.• You cannot use this option in combination with the content keyword HTTP Cookie option to search the same content. See HTTP Content Options on page 1069 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.
K	<p>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.</p> <p>Note the following:</p> <ul style="list-style-type: none">• Cookies included in the message body are treated as body content.• You cannot use this option in combination with the content keyword HTTP Raw Cookie option to search the same content. See HTTP Content Options on page 1069 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	<p>Searches the 3-digit status code in an HTTP response. See the content keyword HTTP Status Code option in HTTP Content Options on page 1069 for more information.</p>
Y	<p>Searches the textual description that accompanies the status code in an HTTP response. See the content keyword HTTP Status Message option in HTTP Content Options on page 1069 for more information.</p>

IMPORTANT! 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 Sourcefire Vulnerability Research Team (VRT) for adding metadata to rules provided by Sourcefire.

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`

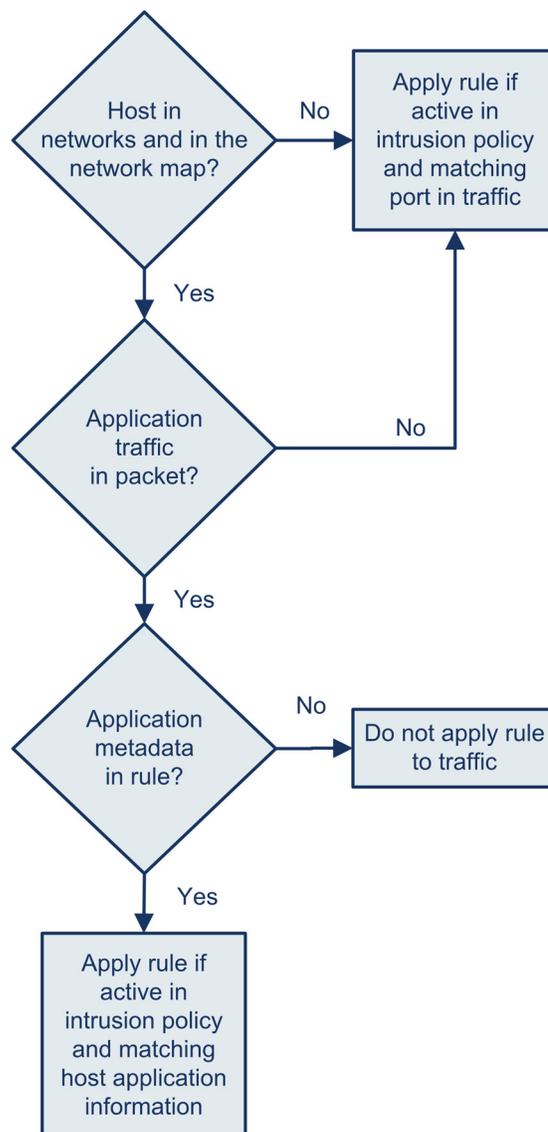
All other characters are permitted.

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 [service Values](#) table describes the most common application values.

IMPORTANT! Contact Sourcefire Support for assistance in defining applications not in the [service Values](#) table.

service Values

VALUE	DESCRIPTION
dcerpc	Distributed Computing Environment/Remote Procedure Calls System
dns	Domain Name System
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

service Values (Continued)

VALUE	DESCRIPTION
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

IMPORTANT! Contact Sourcefire Support for assistance in adding restricted metadata to local rules that might not otherwise function as expected. See [Importing Local Rule Files](#) on page 2080 for more information.

Searching for Rules with Metadata

LICENSE: Protection

To search for rules that use the `metadata` keyword, select the `metadata` keyword on the rules Search page and, optionally, type any portion of the metadata. For example, you can type:

- `author` to display all rules where you have used `author` for `key`.
- `author snortguru` to display all rules where you have used `author` for `key` and `SnortGuru` for `value`.
- `author s` to display all rules where you have used `author` for `key` and any terms such as `SnortGuru` or `SnortUser1` or `SnortUser2` for `value`.

TIP! When you search for both `key` and `value`, use the same connecting operator (equal to [=] or a space character) in searches that is used in the `key value` statement in the rule; searches return different results depending on whether you follow `key` with equal to (=) or a space character.

Note that regardless of the format you use to add metadata, the system interprets your metadata search term as all or part of a `key value` or `key=value` statement. For example, the following would be valid metadata that does not follow a `key value` or `key=value` format:

```
ab cd ef gh
```

However, the system would interpret each space in the example as a separator between a key and value. Thus, you could successfully locate a rule containing the example metadata using any of the following searches for juxtaposed and single terms:

```
cd ef
ef gh
ef
```

but you would not locate the rule using the following search, which the system would interpret as a single `key value` statement:

```
ab ef
```

For more information, see [Searching for Rules](#) on page 1189.

Setting a Red (Level 1) Impact Flag

LICENSE: Protection

You can use the following reserved `key value` statement in a `metadata` keyword:

```
impact_flag red
```

This `key value` statement sets the impact flag to red (level 1) for a local rule you import or a custom rule you create using the rule editor.

Note that when the Sourcefire Vulnerability Research Team (VRT) includes the `impact_flag red` statement in a rule provided by Sourcefire, VRT has

determined that a packet triggering the rule indicates that the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software. See [Using Impact Flags to Evaluate Events](#) on page 641 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](#) on page 1101
- [Inspecting the IP Header Identification Value](#) on page 1102
- [Identifying Specified IP Options](#) on page 1102
- [Identifying Specified IP Protocol Numbers](#) on page 1103
- [Inspecting a Packet's Type of Service](#) on page 1103
- [Inspecting a Packet's Time-To-Live Value](#) on page 1103

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.

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 [Fragbit Operators](#) table after the argument value in the rule.

Fragbit Operators

OPERATOR	DESCRIPTION
plus sign (+)	The packet must match against all specified bits.

Fragbit Operators (Continued)

OPERATOR	DESCRIPTION
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.

IMPORTANT! `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 [IPoption Arguments](#) table lists the available argument values.

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

IPoption Arguments (Continued)

ARGUMENT	DESCRIPTION
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.

IMPORTANT! 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 `tT1` 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 managed 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 Sourcefire 3D System 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](#) on page 1104
- [Inspecting the ICMP Message Type](#) on page 1105
- [Inspecting the ICMP Message Code](#) on page 1105

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 `i type` 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 `i type` 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 `i code` 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 `i code` 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 Sourcefire 3D System 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 Acknowledgement Value](#) on page 1106
- [Inspecting TCP Flag Combinations](#) on page 1107
- [Applying Rules to a TCP or UDP Client or Server Flow](#) on page 1108
- [Identifying Static TCP Sequence Numbers](#) on page 1110
- [Identifying TCP Windows of a Given Size](#) on page 1110
- [Identifying TCP Streams of a Given Size](#) on page 1111

Inspecting the TCP Acknowledgement Value

LICENSE: Protection

You can use the `ack` keyword to compare a value against a packet's TCP acknowledgement number. The rule triggers if a packet's TCP acknowledgement 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.

IMPORTANT! 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](#) on page 1108 for more information about the `flow` keyword.

You can specify the values described in the [flag Arguments](#) table for the `flag` keyword.

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.
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 [Operators Used with flags](#) table describes these operators.

Operators Used with flags

OPERATOR	DESCRIPTION	EXAMPLE
all	The packet must contain all specified flags.	Select <code>urg</code> and <code>all</code> 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 <code>ack</code> , <code>psh</code> , and <code>any</code> to specify that either or both the <code>ack</code> and <code>psh</code> 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 <code>urg</code> and <code>not</code> 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 `f1ow` 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 [State-Related flow Arguments](#) table describes the stream-related arguments you can specify for the `f1ow` keyword:

State-Related flow Arguments

ARGUMENT	DESCRIPTION
Established	Triggers on established connections.
Stateless	Triggers regardless of the state of the stream processor.

The [flow Directional Arguments](#) table describes the directional options you can specify for the `f1ow` keyword:

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 [Stream-Related flow Arguments](#) table describes the stream-related arguments you can specify for the `flow` keyword:

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.

To use the `Established` and `only stream traffic` arguments in TCP or UDP stream preprocessing rules, TCP or UDP stream preprocessing must be enabled as needed. When the required preprocessor is disabled and you enable rules that include these arguments, you are prompted whether to enable the required TCP or UDP preprocessor when you save the policy. See [Using TCP Stream Preprocessing](#) on page 940 and [Reassembling TCP Streams](#) on page 949 for information about using TCP stream preprocessing. See [Using UDP Stream Preprocessing](#) on page 956 for information about using UDP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 for more information on automatically enabling processors.

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.

Note that you must separate each option in the argument with a comma (,).

TCP stream preprocessing must be enabled to use the `stream_size` keyword in a rule. See [Using TCP Stream Preprocessing](#) on page 940 for more information. When TCP stream preprocessing is disabled and you enable rules that use this keyword, you are prompted whether to enable TCP stream preprocessing when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787 for more information.

The [stream_size Keyword Directional Arguments](#) table describes the case-insensitive directional options you can specify for the `stream_size` keyword:

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 <code>both, >, 200</code> 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 <code>either, >, 200</code> would trigger when traffic from the client is greater than 200 bytes OR traffic from the server is greater than 200 bytes.

The [stream_size Keyword Argument Operators](#) table describes the operators you can use with the `stream_size` keyword:

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 [stream_reassemble Optional Arguments](#) table describes the optional arguments you can use with the `stream_reassemble` keyword.

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
```

Note that the TCP stream preprocessor must be enabled to allow processing of rules using the `stream_reassemble` keyword. When the TCP stream preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

To use `stream_reassemble`:

ACCESS: Admin/Intrusion Admin

- ▶ 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](#) on page 1113
- [ssl_version](#) on page 1115

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.

Note that the SSL preprocessor must be enabled to allow processing of rules using the `ssl_state` keyword. When the SSL preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

The `ssl_state` keyword takes the following identifiers as arguments:

ssl_state Arguments

ARGUMENT	PURPOSE
<code>client_hello</code>	Matches against a handshake message with <code>clientHello</code> as the message type, where the client requests an encrypted session.
<code>server_hello</code>	Matches against a handshake message with <code>serverHello</code> as the message type, where the server responds to the client's request for an encrypted session.
<code>client_keyx</code>	Matches against a handshake message with <code>clientKeyExchange</code> as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
<code>server_keyx</code>	Matches against a handshake message with <code>serverKeyExchange</code> as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
<code>unknown</code>	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.

Note that the SSL preprocessor must be enabled to allow processing of rules using the `ssl_version` keyword. When the SSL preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

The `ssl_version` keyword takes the following SSL/TLS version identifiers as arguments:

ssl_version Arguments

ARGUMENT	PURPOSE
<code>sslv2</code>	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 2.
<code>sslv3</code>	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 3.
<code>tls1.0</code>	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.0.
<code>tls1.1</code>	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.1.
<code>tls1.2</code>	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](#) on page 1116
- [ASN.1](#) on page 1117
- [urilen](#) on page 1118
- [DCE/RPC Keywords](#) on page 1119
- [SIP Keywords](#) on page 1125
- [GTP Keywords](#) on page 1128
- [Modbus Keywords](#) on page 1145
- [DNP3 Keywords](#) on page 1148

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 [rpc Keyword Arguments](#) table lists the arguments that the `rpc` keyword accepts.

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 [asn.1 Keyword Arguments](#) table describes the arguments for the `asn1` keyword.

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, <code>pcre</code> , or <code>byte_jump</code> . 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.
- TCP stream preprocessing must be enabled. See [Using TCP Stream Preprocessing](#) on page 940 for more information.
- The HTTP preprocessor must be enabled to allow processing of rules using the `urilen` keyword. When the HTTP preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.
- The rule protocol is always TCP. See [Specifying Protocols](#) on page 1051 for more information.
- Target ports are always HTTP ports. See [Specifying Source and Destination Ports](#) on page 1053 and [Understanding Existing Variables](#) on page 699 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.

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 [DCE/RPC Keywords](#) 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](#) on page 810 for more information.

The DCE/RPC preprocessor must be enabled to allow processing of rules that include these keywords. When the DCE/RPC preprocessor is disabled and you enable rules that use these keywords, you are prompted whether to enable the

preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

DCE/RPC Keywords

USE THIS KEYWORD...	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](#) on page 1079 and [Reading Packet Data into Keyword Arguments](#) on page 1156.

Sourcefire 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 Use Fast Pattern Matcher` argument. See [Searching for Content Matches](#) on page 1063 and [Use Fast Pattern Matcher](#) on page 1074 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.

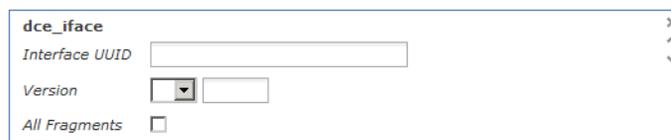
See the following sections for more information:

- [dce_iface](#) on page 1121
- [dce_opnum](#) on page 1123
- [dce_stub_data](#) on page 1124

dce_iface

LICENSE: Protection

You can use the **dce_iface** keyword to identify a specific DCE/RPC service.



The screenshot shows a configuration window for the **dce_iface** keyword. It contains the following fields:

- Interface UUID:** A text input field.
- Version:** A dropdown menu with a small arrow icon to its left.
- All Fragments:** A checkbox.

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](#) on page 1123 and [dce_stub_data](#) on page 1124 for more information.

Note that the DCE/RPC preprocessor must be enabled to allow processing of rules using the **dce_iface** keyword. When the DCE/RPC preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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 `svcsvc` 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 [dce_iface Arguments](#) table summarizes the `dce_iface` keyword arguments.

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.

dce_iface Arguments (Continued)

ARGUMENT	DESCRIPTION
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.

Note that the DCE/RPC preprocessor must be enabled to allow processing of rules using the **dce_opnum** keyword. When the DCE/RPC preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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, as shown in the following example. See [dce_iface](#) on page 1121 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

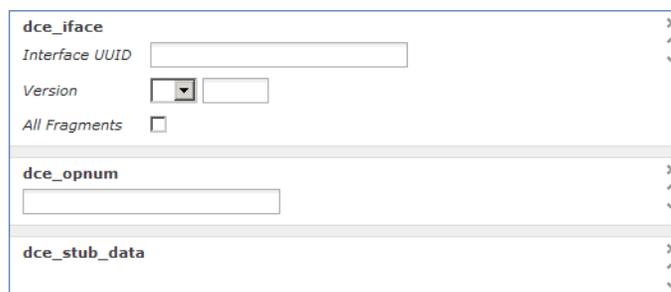
dce_stub_data

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.

Note that the DCE/RPC preprocessor must be enabled to allow processing of rules using the `dce_stub_data` keyword. When the DCE/RPC preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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, as shown in the following example.



The screenshot shows a configuration window with three sections. The first section is titled 'dce_iface' and contains a text input field for 'Interface UUID', a dropdown menu for 'Version', and a checkbox for 'All Fragments'. The second section is titled 'dce_opnum' and contains a text input field. The third section is titled 'dce_stub_data' and is currently empty. Each section has a close button (X) and a scroll indicator (up and down arrows) on the right side.

The `dce_stub_data` keyword has no arguments. See [dce_iface](#) on page 1121 and [dce_opnum](#) on page 1123 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](#) on page 756 and [Preventing Rate-Based Attacks](#) on page 971 for more information.

See the following sections for more information:

- [sip_header](#) on page 1125
- [sip_body](#) on page 1125
- [sip_method](#) on page 1126
- [sip_stat_code](#) on page 1127

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.

A screenshot of a text input field with a light blue border. The text 'sip_header' is entered in the field. On the right side of the field, there are three small icons: a close button (an 'x'), an up arrow, and a down arrow.

The `sip_header` keyword has no arguments. See [sip_method](#) on page 1126 and [sip_stat_code](#) on page 1127 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"; )
```

Note that the SIP preprocessor must be enabled to allow processing of rules using the `sip_header` keyword. When the SIP preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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.

A screenshot of a text input field with a light blue border. The text 'sip_body' is entered in the field. On the right side of the field, there are two small icons: a close button (an 'x') and an up arrow.

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.

Note also that the SIP preprocessor must be enabled to allow processing of rules using the `sip_body` keyword. When the SIP preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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.

A screenshot of a configuration field for the 'sip_method' keyword. The field is a text input box with a blue border. Above the input box, the text 'sip_method' is displayed. To the right of the input box, there are three small icons: a close button (x), an up arrow (^), and a down arrow (v).

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](#) on page 873 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
```

The SIP preprocessor must be enabled to allow processing of rules using the `sip_method` keyword. When the SIP preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

Sourcefire 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](#) on page 1063 and [Use Fast Pattern Matcher](#) on page 1074 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.

`sip_stat_code` Values

TO DETECT...	SPECIFY...	FOR EXAMPLE...	DETECTS...
a specific status code	the three-digit status code	189	189

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 that the SIP preprocessor must be enabled to allow processing of rules using the `sip_stat_code` keyword. When the SIP preprocessor is disabled and you enable rules that use this keyword, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

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.

The GTP preprocessor must be enabled to allow processing of rules using GTP keywords. When the GTP preprocessor is disabled and you enable rules that use these keywords, you are prompted whether to enable the preprocessor when you save the policy. See [Automatically Enabling Advanced Settings](#) on page 787.

See the following sections for more information:

- [gtp_version](#) on page 1128
- [gtp_type](#) on page 1129
- [gtp_info](#) on page 1137

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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **gtp_version** in the drop-down list and click **Add Option**.

The `gtp_version` keyword appears.



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.

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

GTP Message Types (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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
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

GTP Message Types (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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_context	N/A
59	N/A	forward_relocation_ complete_ack	N/A
60	N/A	forward_srns_context_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

GTP Message Types (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
97	N/A	mbms_notification_ response	update_bearer_request
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

GTP Message Types (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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

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

GTP Message Types (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **gtp_type** in the drop-down list and click **Add Option**.

The **gtp_type** keyword appears.

A screenshot of a configuration field for the 'gtp_type' keyword. The field is a rectangular box with a light blue border. Inside the box, the text 'gtp_type' is displayed in a small font. To the right of the text, there are three small icons: a close button (an 'x' in a square), a search button (a magnifying glass), and a dropdown arrow (a downward-pointing triangle). Below the text, there is a small, empty rectangular input area.

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](#) on page 1130 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.

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
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

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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
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

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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_ctxt
125	N/A	N/A	src_rnc_pdc
126	N/A	N/A	udp_src_port
127	charge_id	charge_id	apn_restriction
128	end_user_address	end_user_address	selection_mode
129	mm_ctxt	mm_ctxt	src_id
130	pdp_ctxt	pdp_ctxt	N/A
131	apn	apn	change_report_action
132	protocol_config	protocol_config	fq_cs

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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_pdcph	nمبر
162	N/A	addi_trace_info	N/A
163	N/A	hop_counter	N/A
164	N/A	plmn_id	N/A
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

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
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

GTP Information Elements (Continued)

VALUE	VERSION 0	VERSION 1	VERSION 2
201	N/A	apn_ambr_nsapi	N/A
202	N/A	ggsn_backoff_timer	N/A
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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **gtp_info** in the drop-down list and click **Add Option**.

The `gtp_info` keyword appears.



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](#) on page 1138 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](#) on page 1146
- [modbus_func](#) on page 1146
- [modbus_unit](#) on page 1147

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:

ACCESS: Admin/Intrusion Admin

- ▶ 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.

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

Modbus Function Codes (Continued)

VALUE	STRING
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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **modbus_func** in the drop-down list and click **Add Option**.

The `modbus_func` keyword appears.



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](#) on page 1146 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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **modbus_unit** in the drop-down list and click **Add Option**.

The `modbus_unit` keyword appears.



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](#) on page 1148
- [dnp3_func](#) on page 1149
- [dnp3_ind](#) on page 1151
- [dnp3_obj](#) on page 1152

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:

ACCESS: Admin/Intrusion Admin

- ▶ 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.

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

DNP3 Function Codes (Continued)

VALUE	STRING
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
130	unsolicited_response
131	authenticate_resp

To specify DNP3 function codes:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **dnp3_func** in the drop-down list and click **Add Option**.

The **dnp3_func** keyword appears.

A screenshot of a web form. At the top, the text 'dnp3_func' is displayed. Below it is a text input field. To the right of the input field is a small rectangular box containing a close button (an 'X') and a dropdown arrow.

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](#) on page 1149 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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **dnp3_ind** in the drop-down list and click **Add Option**.

The **dnp3_ind** keyword appears.

A screenshot of a web interface showing a drop-down menu. The text 'dnp3_ind' is displayed in the menu, with a small 'x' icon to its right and a downward-pointing arrow below it. Below the menu is an empty input field.

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.

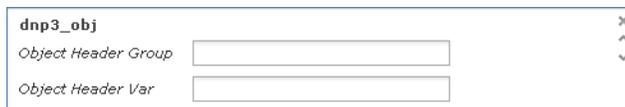
You identify object headers by specifying the decimal number for the type of object header group and the decimal number for the type of object variation. The combination of the two defines a specific type of DNP3 object.

To specify a DNP3 object:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **dnp3_obj** in the drop-down list and click **Add Option**.

The **dnp3_obj** keyword appears.

A screenshot of a web interface showing a drop-down menu. The text 'dnp3_obj' is displayed in the menu, with a small 'x' icon to its right and a downward-pointing arrow below it. Below the menu are two input fields. The first is labeled 'Object Header Group' and the second is labeled 'Object Header Var'.

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 Sourcefire 3D System provides the following keywords to evaluate packet characteristics:

- [dsize](#) on page 1153
- [isdataat](#) on page 1153
- [sameip](#) on page 1155
- [fragoffset](#) on page 1155
- [cvs](#) on page 1155

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, use `400<>500`.

WARNING! 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 [isdataat Arguments](#) table lists the arguments you can use with the `isdataat` keyword.

isdataat Arguments

ARGUMENT	TYPE	DESCRIPTION
Offset	Required	<p>The specific location in the payload. For example, to test that data appears at byte 50 in the packet payload, you would specify <code>50</code> as the offset value. A <code>!</code> modifier negates the results of the <code>isdataat</code> test; it alerts if a certain amount of data is not present within the payload.</p> <p>You can also use an existing <code>byte_extract</code> variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 1156 for more information.</p>
Relative	Optional	<p>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 <code>8</code>.</p>
Raw Data	Optional	<p>Specifies that the data is located in the original packet payload before decoding or application layer normalization by any Sourcefire 3D System preprocessor. You can use this argument with Relative if the previous content match was in the raw packet data.</p>

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`:

ACCESS: Admin/Intrusion Admin

- ▶ 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.

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 `cvsv` 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](#) on page 950.

To detect malformed CVS entries:

ACCESS: Admin/Intrusion Admin

- ▶ 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 [Required byte_extract Arguments](#) table describes the arguments required by the `byte_extract` keyword.

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 on page 1157 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 [Additional Optional byte_extract Arguments](#) table.

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 on page 1157 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 [Endianness `byte_extract` Arguments](#) table. The rules engine uses big endian byte order if you do not select either argument.

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	<p>Specifies a <code>byte_extract</code> keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 810 for more information.</p> <p>The DCE/RPC preprocessor determines big endian or little endian byte order, and the Number Type and Endian arguments do not apply.</p> <p>When you enable this argument, you can also use <code>byte_extract</code> in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords on page 1119 for more information.</p> <p>The DCE/RPC preprocessor must be enabled to allow processing of rules that include this option. When the DCE/RPC preprocessor is disabled and you enable rules that use this option, you are prompted whether to enable the preprocessor when you save the policy. See Automatically Enabling Advanced Settings on page 787.</p>

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 [Number Type `byte_extract` arguments](#) table.

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 [Arguments Accepting a byte_extract Variable](#) table lists the keyword arguments where you can specify a variable defined in the `byte_extract` keyword.

Arguments Accepting a byte_extract Variable

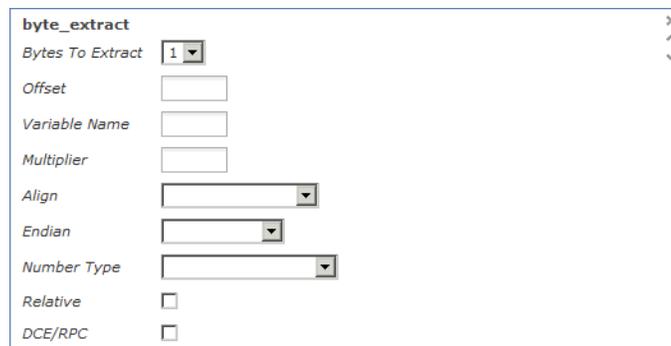
KEYWORD	ARGUMENT
content	Depth, Offset, Distance, Within See Constraining Content Matches on page 1065 for more information.
byte_jump	Offset See byte_jump on page 1080 for more information.
byte_test	Offset, Value See byte_test on page 1083 for more information.
isdataat	Offset See isdataat on page 1153 for more information.

To use `byte_extract`:

ACCESS: Admin/Intrusion Admin

- ▶ On the Create Rule page, select `byte_extract` in the drop-down list and click **Add Option**.

The `byte_extract` section appears beneath the last keyword you selected.



The screenshot shows a configuration form for the `byte_extract` keyword. The form includes the following fields and options:

- Bytes To Extract:** A dropdown menu with the value '1' selected.
- Offset:** An empty text input field.
- Variable Name:** An empty text input field.
- Multiplier:** An empty text input field.
- Align:** A dropdown menu.
- Endian:** A dropdown menu.
- Number Type:** A dropdown menu.
- Relative:** A checkbox, currently unchecked.
- DCE/RPC:** A checkbox, currently unchecked.

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 Drop Rules](#) on page 941 for more

information.

Note that to initiate additional TCP resets you must ensure that TCP Stream Configuration is enabled, and to initiate additional ICMP unreachable packets you must ensure that UDP Stream Configuration is enabled. See [Modifying Advanced Settings](#) on page 774 for more information. Note also that initial active responses do not require that you enable either TCP or UDP Stream Configuration.

See the following sections for information specific to the keywords you can use to initiate active responses:

- [Initiating Active Responses by Type and Direction](#) on page 1161
- [Sending an HTML Page Before a TCP Reset](#) on page 1162
- [Setting the Active Response Reset Attempts and Interface](#) on page 1164

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](#) on page 1051 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 [resp Arguments](#) table lists the arguments you can use with the **resp** keyword to specify exactly what you want the Sourcefire 3D System to do when the rule triggers.

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.

resp Arguments (Continued)

ARGUMENT	DESCRIPTION
reset_both	Directs a TCP reset packet to both the sending and receiving endpoints. Alternatively, you can specify <code>rst_all</code> , 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: <ul style="list-style-type: none"> • 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](#) on page 1164 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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select `resp` in the drop-down list and click **Add Option**.

The `resp` keyword appears.



2. Specify any of the arguments in the [resp Arguments table](#) on page 1161 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](#) on page 1047 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.*

IMPORTANT! 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. Sourcefire recommends that you test `react` rules extensively before activating them in a production environment.

See [Setting the Active Response Reset Attempts and Interface](#) on page 1164 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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select `react` in the drop-down list and click **Add Option**.

The `react` keyword appears.

A screenshot of a web interface showing a dropdown menu. The word "react" is selected and displayed in the dropdown box. To the right of the box are three small icons: a close button (x), an up arrow, and a down arrow.

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 Drop Rules](#) on page 941 for more information.

You use the **config response** command by inserting it on a separate line in a custom variable with the reserved name `USER_CONF`. See [Understanding Custom Variables](#) on page 710 and [Configuring Custom Variables](#) on page 711 for information on using a `USER_CONF` variable.

WARNING! Do not use the reserved variable `USER_CONF` to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Sourcefire Support. Conflicting or duplicate configurations will halt the system.

To specify active response reset attempts, the active response interface, or both:

ACCESS: Admin/Intrusion Admin

- ▶ 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 a variable with the reserved name `USER_CONF`. 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 [detection_filter Track Arguments](#) table to specify how the system tracks event instances.

detection_filter Track Arguments

ARGUMENT	DESCRIPTION
<code>by_src</code>	Detection criteria count by source IP address.
<code>by_dst</code>	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](#) on page 747, [Configuring Suppression Per Intrusion Policy](#) on page 753, [Setting a Dynamic Rule State](#) on page 758, and [Importing Local Rule Files](#) on page 2080 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 [Tag Arguments](#) table, [Count Argument](#) table and [Logging Metrics Arguments](#) table describe the other available arguments.

You can choose from two types of tagging. The [Tag Arguments](#) table describes the two types of tagging.

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:

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 [Logging Metrics Arguments](#) table.

WARNING! 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.

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 per managed device. 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](#) on page 1171 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](#) on page 1172 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 [flowbits Options](#) 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 (-).

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.

flowbits Options (Continued)

OPERATOR	STATE OPTION	GROUP	DESCRIPTION
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.
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, all rules that affect `flowbits` assignment (`set`, `setx`, `unset`, `toggle`) and define a corresponding state name are also enabled. 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.
- Stream preprocessing must be enabled for the TCP or UDP protocol specified in rules using the `flowbits` keyword. If you enable a rule using the `flowbits` keyword in an intrusion policy where the necessary TCP or UDP stream preprocessing is disabled, you are prompted when you try to save your changes whether to allow the system to automatically enable the required TCP or UDP stream preprocessing. See [Automatically Enabling Advanced Settings](#) on page 787 for more information.

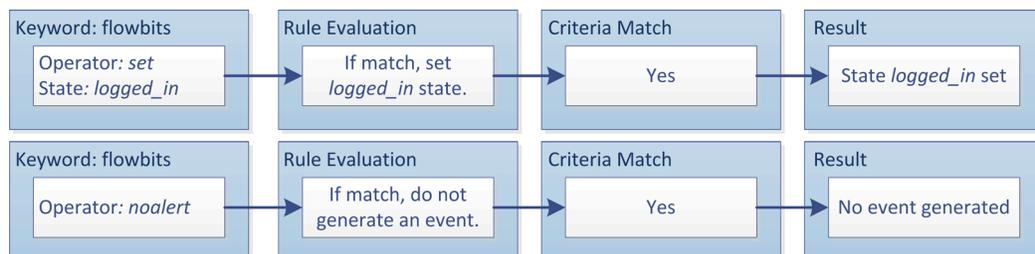
flowbits Example Using state_name

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:

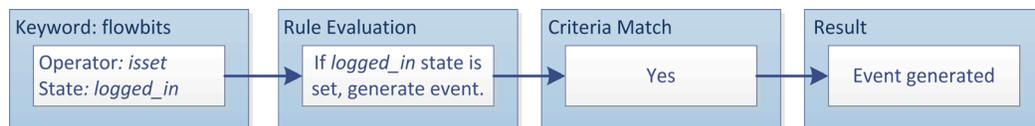


Note that `flowbits:set` sets a state of `logged_in`, while `flowbits:noalert` suppresses the alert because you are likely to see many innocuous login sessions on an IMAP server.

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.

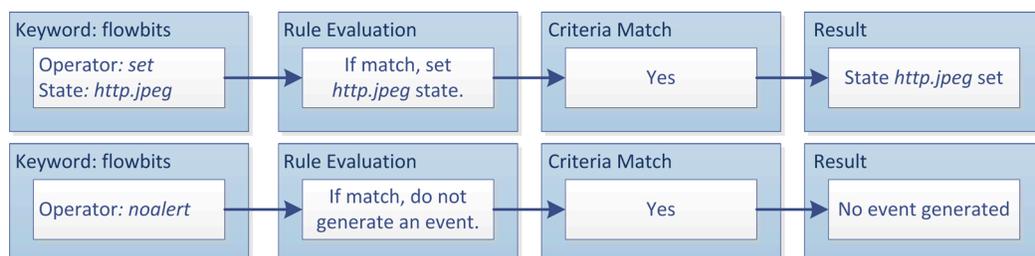
flowbits Example Resulting in a False Positive

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\x2f?jpe?g/smi";
flowbits:set,http.jpeg; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

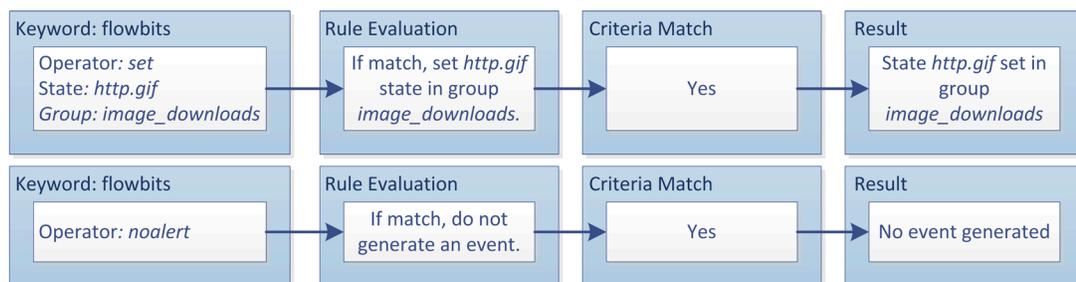


The `content` and `pcrc` keywords in the first rule fragment match a JPEG file download, `flowbits:set,http.jpeg` sets the `http.jpeg` flowbits state, and `flowbits:noalert` stops the rule from generating events. No event is generated because the rule's purpose is to detect the file download and set 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/"; pcrc:"/^Content-Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
flowbits:set,http.gif,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

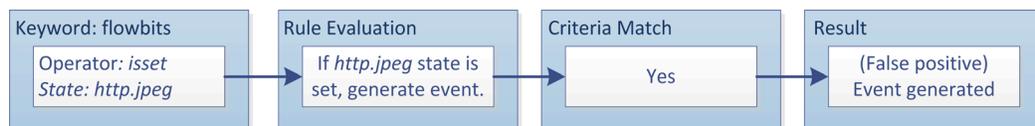


The `content` and `pcrc` keywords in the second rule match the GIF file download, `flowbits:set,http.gif` sets the `http.gif` flowbit state, and `flowbits:noalert` stops the rule from generating an event. Note that the `http.jpeg` state set by the first rule fragment is still set even though it is no 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|"; pcrc:"
/\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 `pcr` 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

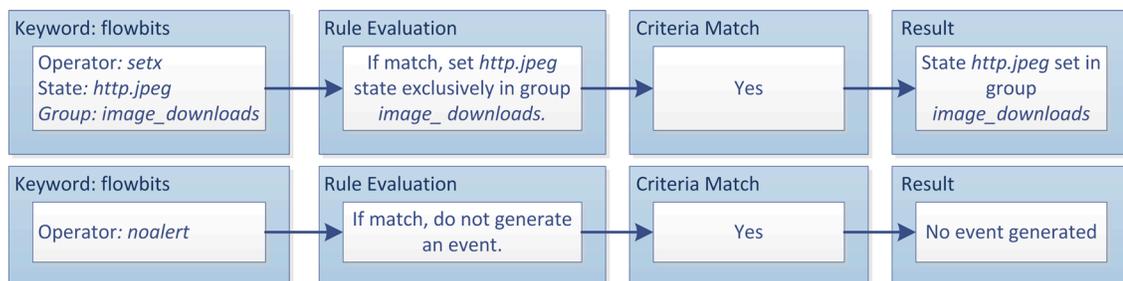
The following example illustrates how including state names in a group and using the `setx` operator can prevent false positives.

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/";pcr:"/^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:



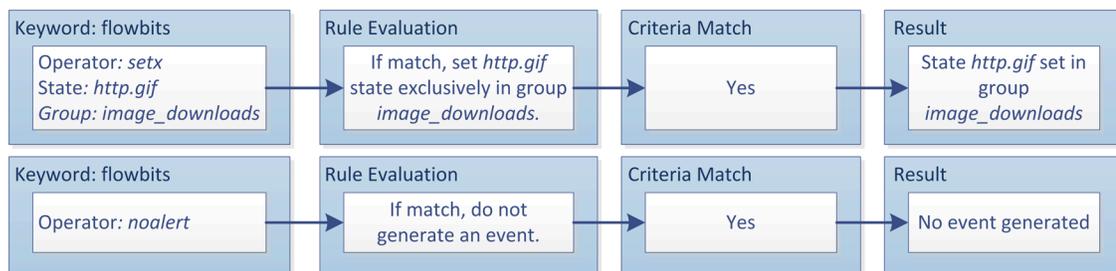
When the first rule fragment detects a JPEG file download, the `flowbits:setx,http.jpeg,image_downloads` keyword sets the `flowbits` state to `http.jpeg` and includes the state in the `image_downloads` group.

The next rule then detects a subsequent GIF file download:

```

(msg:"GIF transfer"; content:"image/"; pcr:"/^Content-Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
flowbits:setx,http.gif,image_downloads; flowbits:noalert;)
    
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

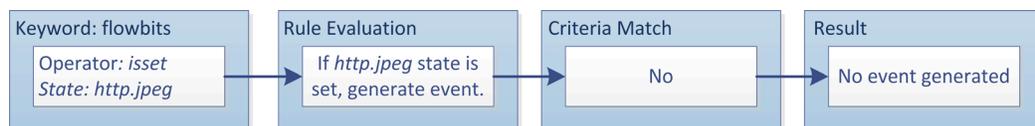


When the second rule fragment matches the GIF download, the `flowbits:setx,http.gif,image_downloads` keyword sets the `http.gif` `flowbits` state and unsets `http.jpeg`, the other state in the group.

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.

The HTTP Inspect preprocessor must be enabled for rules using the `http_encode` keyword to return matches. If you enable those rules in an intrusion policy where the HTTP preprocessor is disabled and try to save the policy, you are prompted whether to allow the system to automatically enable the HTTP preprocessor. For more information on automatically enabling processors and other advanced intrusion policy features, see [Automatically Enabling Advanced Settings](#) on page 787.

You must also configure the preprocessor to inspect HTTP responses and HTTP cookies to return matches for these. See [Decoding HTTP Traffic](#) on page 850 and

[Selecting Server-Level HTTP Normalization Options](#) on page 854 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](#) on page 862 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 [http_encode Encoding Types](#) table describes the encoding types this option can generate events for in HTTP URIs, headers, cookies, and set-cookies:

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.
unicode	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:

ACCESS: Admin/Intrusion Admin

1. Add the **http_encode** keyword to a rule.
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.

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.

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.

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](#) on page 850 and **Inspect HTTP Responses** in [Selecting Server-Level HTTP Normalization Options](#) on page 854 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](#) on page 850, and the **Inspect HTTP Responses** and **Inspect Compressed Data** options in [Selecting Server-Level HTTP Normalization Options](#) on page 854. 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](#) on page 850 and **Inspect HTTP Responses** in [Selecting Server-Level HTTP Normalization Options](#) on page 854 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](#) on page 895 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](#) on page 880, [Decoding POP Traffic](#) on page 884, and [Decoding SMTP Traffic](#) on page 889 for more information. You can use multiple `file_data` keywords in a rule.

To point to the beginning of a specific payload type:

ACCESS: Admin/Intrusion Admin

- ▶ 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 [Understanding Server-Level FTP Options](#) on page 838.
- To normalize telnet traffic for inspection, you must enable the FTP & Telnet preprocessor **Normalize** telnet option; see [Understanding Telnet Options](#) on page 835.
- To normalize SMTP traffic for inspection, you must enable the SMTP preprocessor **Normalize** option; see [Understanding SMTP Decoding](#) on page 890.

You can use multiple `pkt_data` keywords in a rule.

To point to the beginning of the packet payload:

ACCESS: Admin/Intrusion Admin

- ▶ 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](#) on page 1180
- [base64_data](#) on page 1181

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](#) on page 1181 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 [Optional base64_decode Arguments](#) table describes the arguments you can use with the `base64_decode` keyword.

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:

ACCESS: Admin/Intrusion Admin

1. On the Create Rule page, select **base64_decode** from the drop-down list and click **Add Option**.

The `base64_decode` keyword appears.



2. Optionally, select any of the arguments described in the [Optional base64_decode Arguments](#) table on page 1180.

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](#) on page 1074 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](#) on page 1069 for more information.

To inspect decoded Base64 data:

ACCESS: Admin/Intrusion Admin

- ▶ 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 Sourcefire and save your changes as a new rule. Note that for shared object rules provided by Sourcefire, 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](#) on page 1182
- [Modifying Existing Rules](#) on page 1184
- [Adding Comments to Rules](#) on page 1187
- [Deleting Custom Rules](#) on page 1188

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 Sourcefire. 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 (although intrusion rules created on the secondary Defense Center in a high availability pair begin with the number 1,000,000,000) 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.

IMPORTANT! The system assigns a new SID to any custom rule in an intrusion policy that you import. For more information, see [Importing and Exporting Objects](#) on page 2225.

To write a custom standard text rule using the rule editor:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.

2. Click **Create Rule**.

The Create Rule page appears.

The screenshot shows the 'Create New Rule' interface. It includes a 'Message' field, a 'Classification' dropdown menu (currently showing 'A Client was Using an Unusual Port'), an 'Action' dropdown menu (showing 'alert'), a 'Protocol' dropdown menu (showing 'icmp'), a 'Direction' dropdown menu (showing 'Directional'), and input fields for 'Source IPs', 'Source Port', 'Destination IPs', and 'Destination Port'. At the bottom, there is a 'Detection Options' section with a dropdown menu (showing 'ack'), an 'Add Option' button, and a 'Save As New' button.

3. In the **Message** field, enter the message you want displayed with the event. For details on event messages, see [Defining the Event Message](#) on page 1058.

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.

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](#) on page 1058.
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.
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](#) on page 1051.
7. In the **Source IPs** field, enter the originating IP address or address range for traffic that should trigger the rule. In the **Destination IPs** field, enter the destination IP address or address range for traffic that should trigger the rule. For more detailed information about the IP address syntax that the rule editor accepts, see [Defining IP Addresses in Variables and Rules](#) on page 713.

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.

IMPORTANT! 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 Variables and Rules](#) on page 715.

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
10. From the **Detection Options** list, select the keyword that you want to use.
11. Click **Add Option**.
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](#) on page 1055.

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.

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 [Applying an Access Control Policy](#) on page 466 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 Sourcefire and create one or more new instances of the rule by saving it.

Creating a rule or modifying a Sourcefire 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 Sourcefire Vulnerability Research Team (VRT). You can delete instances of a shared object rule that you create, but you cannot delete shared object rules provided by Sourcefire. See [Understanding Rule Headers](#) on page 1049 and [Deleting Custom Rules](#) on page 1188 for more information.

IMPORTANT! Do not modify the protocol for a shared object rule; doing so would render the rule ineffective.

To modify a rule:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.

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 rules by searching for them, enter the search criteria (most simply, the SID) for the rule or rules you want and click **Search**. Click a rule returned by the search as appropriate. See [Searching for Rules](#) on page 1189 for more information.
 - 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](#) on page 1192 for more information.

The rule editor opens, displaying the rule you selected.

The screenshot shows the 'Edit Rule 1:356:8' interface. At the top right, there is a '(Rule Comment)' field. Below it, the 'Message' field contains 'FTP passwd retrieval attempt'. The 'Classification' dropdown is set to 'A Suspicious Filename was Detected'. Under 'Edit Classifications', there is an 'Action' dropdown set to 'alert'. Below that, there is a 'Fast Pattern Matcher Offset and Length' field. A 'reference' section contains a list with 'arachnids,213'. At the bottom, there is an 'ack' dropdown, an 'Add Option' button, and a 'Save As New' button.

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.

3. Make any modifications to the rule (see [Writing New Rules](#) on page 1182 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](#) on page 743 and activate the local rule.

4. Activate the intrusion policy by applying it as part of an access control policy as described in [Applying an Access Control Policy](#) on page 466 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:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.

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 searching for it, enter the search criteria (most simply, the SID) for the rule you want and click **Search**. Click the rule returned by the search as appropriate. See [Searching for Rules](#) on page 1189 for more information.
- 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](#) on page 1192 for more information.

The rule editor appears.

3. Click **Rule Comment**.

The Rule Comment page appears.



4. Enter your comment in the text box and click **Add Comment**.

The comment is saved in the comment text box.

TIP! You can also add and view rule comments in an intrusion event's packet view. For more information, see [Viewing Event Information](#) on page 625.

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 provided by Sourcefire.

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](#) on page 1184 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](#) on page 2074.

See the following sections for more information:

- For information on creating custom rules, see [Writing New Rules](#) on page 1182.
- For information on importing local rules, see [Importing Rule Updates and Local Rule Files](#) on page 2072.
- For information on setting rule states, see [Setting Rule States](#) on page 743.

To delete custom rules:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.

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.

Searching for Rules

LICENSE: Protection

The Sourcefire 3D System provides thousands of standard text rules, and the Sourcefire Vulnerability Research Team continues to add rules as new vulnerabilities and exploits are discovered. You can easily search for specific rules so that you can activate, deactivate, or edit them.

The [Rule Search Criteria](#) table describes the available search options:

Rule Search Criteria

OPTION	DESCRIPTION
Signature ID	To search for a single rule based on Snort ID (also called the Signature ID), enter a Snort ID number. To search for multiple rules, enter a comma-separated list of Snort ID numbers. This field has an 80-character limit.
Generator ID	To search for standard text rules, select 1 . To search for shared object rules, select 3 .
Message	To search for a rule with a specific message, enter a single word from the rule message in the Message field. For example, to search for DNS exploits, you would enter DNS , or to search for buffer overflow exploits, enter overf1ow .

Rule Search Criteria (Continued)

OPTION	DESCRIPTION
Protocol	To search rules that evaluate traffic of a specific protocol, select the protocol. If you do not select a protocol, search results contain rules for all protocols.
Source Port	To search for rules that inspect packets originating from a specified port, enter a source port number or a port-related variable.
Destination Port	To search for rules that inspect packets destined for a specific port, enter a destination port number or a port-related variable.
Source IP	To search for rules that inspect packets originating from a specified IP address, enter a source IP address or an IP address-related variable.
Destination IP	To search for rules that inspect packets destined for a specified IP address, enter a destination IP address or an IP address-related variable.
Keyword	To search for specific keywords, you can use the keyword search options. You select a keyword and a keyword value for which to search. You can also precede the keyword value with an exclamation point (!) to match any value other than the specified value.
Category	To search for rules in a specific category, select the category from the Category list.
Classification	To search for rules that have a specific classification, select the classification name from the Classification list.
Rule State	To search for rules within a specific policy and a specific rule state, select the policy from the first Rule State list, and choose a state from the second list to search for rules set to Generate Events , Drop and Generate Events , or Disabled .

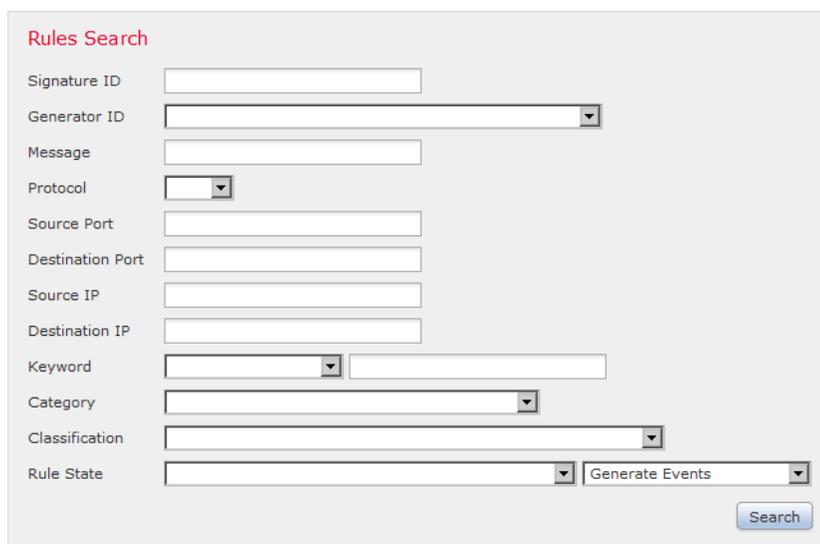
To search for specific rules:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.

2. Click **Search** on the toolbar.
The Search page appears.



The screenshot shows a web interface titled "Rules Search". It contains several input fields and dropdown menus for filtering rules. The fields are: Signature ID (text input), Generator ID (dropdown), Message (text input), Protocol (dropdown), Source Port (text input), Destination Port (text input), Source IP (text input), Destination IP (text input), Keyword (dropdown and text input), Category (dropdown), Classification (dropdown), Rule State (dropdown), and a "Generate Events" checkbox. A "Search" button is located at the bottom right of the form.

3. Add search criteria using any of the fields described in the [Rule Search Criteria table](#) on page 1189.

IMPORTANT! You must specify at least one search criterion to search for rules.

4. Perform the following steps to search for rules that contain specific keywords:
 - From the drop-down list in the **Keyword** section, select the keyword for which to search.
For a list of each available keyword, see [Understanding Keywords and Arguments in Rules](#) on page 1055.
 - In the **Keyword** field, enter the arguments for which you want to search.
5. Click **Search**.
The page reloads, showing a list of the rules that match your search criteria.
6. To view or edit a rule (or a copy of the rule, if it is a system rule), click the hyperlinked rule message. See [Modifying Existing Rules](#) on page 1184 for detailed information about editing rules.

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. You can also use any of the options in the context menu for the page.

See the following sections for more information:

- [Using Keywords in a Rule Filter](#) on page 1192
- [Using Character Strings in a Rule Filter](#) on page 1194
- [Combining Keywords and Character Strings in a Rule Filter](#) on page 1195
- [Filtering Rules](#) on page 1195

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](#) on page 1193 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](#) on page 1194 for more information.

The [Rule Filter Keywords](#) table describes the specific filtering keywords and arguments you can use to filter rules.

Rule Filter Keywords

KEYWORD	DESCRIPTION	EXAMPLE
<code>arachnids</code>	Returns one or more rules based on all or part of the Arachnids ID in a rule reference. See Defining the Event Reference on page 1062 for more information.	<code>arachnids:181</code>
<code>bugtraq</code>	Returns one or more rules based on all or part of the Bugtraq ID in a rule reference. See Defining the Event Reference on page 1062 for more information.	<code>bugtraq:2120</code>
<code>cve</code>	Returns one or more rules based on all or part of the CVE number in a rule reference. See Defining the Event Reference on page 1062 for more information.	<code>cve:2003-0109</code>
<code>gid</code>	The argument <code>1</code> returns standard text rules. The argument <code>3</code> returns shared object rules. See Reading Preprocessor Generator IDs on page 784 and the Rule Types table on page 719 for more information.	<code>gid:3</code>
<code>mcafee</code>	Returns one or more rules based on all or part of the McAfee ID in a rule reference. See Defining the Event Reference on page 1062 for more information.	<code>mcafee:10566</code>
<code>msg</code>	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 on page 1058 for more information.	<code>msg:chat</code>

Rule Filter Keywords (Continued)

KEYWORD	DESCRIPTION	EXAMPLE
nessus	Returns one or more rules based on all or part of the Nessus ID in a rule reference. See Defining the Event Reference on page 1062 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 on page 1062 and Defining the Event Message on page 1058 for more information.	ref:MS03-039
sid	Returns the rule with the exact Signature ID. See Reading Preprocessor Generator IDs on page 784 for more information.	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 on page 1062 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](#) on page 1058. For information on rule SIDs and GIDs, see [Reading Preprocessor Generator IDs](#) on page 784.

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, including selecting any of the features available in the context menu.

To filter for specific rules:

ACCESS: Admin/Intrusion Admin

1. Select **Policies > Intrusion > 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](#) on page 1184 for more information.

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.

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 (✕).

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.
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](#) on page 1192 for more information.
The page refreshes to display any group that contains at least one matching rule.
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, pre-filtered page, click the filter clearing icon ().
7. Optionally, make any changes to the rule that you would normally make on the page. See [Modifying Existing Rules](#) on page 1184.
To put any changes you make into effect, apply the intrusion policy part of an access control policy as described in [Applying an Access Control Policy](#) on page 466.

CHAPTER 31

WORKING WITH MALWARE PROTECTION AND FILE CONTROL

Malicious software, or *malware*, can enter your organization's network via multiple routes. To help you identify and mitigate the effects of malware, the Sourcefire 3D System's file control, network file trajectory, and advanced malware protection components can detect, track, 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.

If your organization has a FireAMP subscription, the Defense Center can also receive endpoint-based malware detection data from Sourcefire's malware awareness network (called the *Sourcefire cloud*). The Defense Center presents this data alongside the network-based file and malware data generated by the system.

The Context Explorer and the dashboard provide you with different types of high-level views of the files (including malware files) detected in your organization. To further target your analysis, you can use a malware file's *network file trajectory* page to track the spread of an individual threat across hosts over time, allowing you to concentrate outbreak control and prevention efforts where most useful.

IMPORTANT! Managed devices **must** be running at least Version 5.1.1 of the Sourcefire 3D System to perform network-based file control and malware detection, and Version 5.2 to view file trajectories and block malware.

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 target devices, as described in the following table.

License Requirements for File and Malware Detection

FEATURE	DESCRIPTION	LICENSE
file control	detect and optionally block the transmission of certain file types in network traffic	Protection
advanced malware protection	detect, track, and optionally block the transmission of malware files in network traffic	Malware
FireAMP integration	receive endpoint-based malware information from the Sourcefire cloud, using your organization's FireAMP subscription; track the transmission of malware files using that information	Any
geolocation	detect source and destination countries and other geographical information associated with file and malware events	FireSIGHT (with GeoDB update for detailed information)

Because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances to perform network-based malware protection or view file trajectories for files for which you conduct a malware cloud lookup.

For more information, see:

- [Understanding Malware Protection and File Control](#) on page 1199
- [Understanding and Creating File Policies](#) on page 1206
- [Working with Sourcefire Cloud Connections for FireAMP](#) on page 1219
- [Working with File Events](#) on page 1222
- [Working with Malware Events](#) on page 1229
- [Working with Network File Trajectory](#) on page 1241
- [Using Geolocation](#) on page 1816

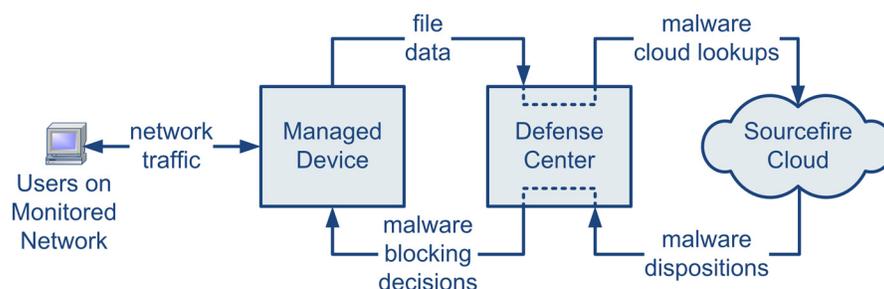
Understanding Malware Protection and File Control

LICENSE: Protection, Malware, or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

Using the *advanced malware protection* feature, you can configure the Sourcefire 3D System to detect, track, and optionally block malware files being transmitted on your network, as shown in the following diagram.



The system can detect and optionally block malware in many types of files, including PDFs, Microsoft Office documents, and others. Managed devices monitor specific application protocol-based network traffic for transmissions of those file types. When a device detects an eligible file, it sends the file's SHA-256 hash value to the Defense Center, which then performs a *malware cloud lookup* using that information. The Sourcefire cloud returns a disposition to the Defense Center.

To inspect or block files, you must enable a Protection license on the managed devices where you apply policies. To perform malware cloud lookups on and optionally block malware files, you must also enable a Malware license for those devices.

Understanding File Dispositions

A file can have one of the following dispositions:

- **malware** indicates that the cloud categorized the file as malware.
- **clean** indicates that the cloud categorized the file as clean.
- **neutral** indicates that a malware cloud lookup occurred before the cloud assigned a disposition. The file is uncategorized.
- **unavailable** indicates that the Defense Center could not perform a malware cloud lookup.

TIP! If several recent malware events have the disposition **unavailable**, check the cloud connection and port configuration. For more information, see [Security, Internet Access, and Communication Ports](#) on page 50.

Based on the file's disposition, the Defense Center instructs the managed device either to block the file or to allow its upload or download. To improve performance, if the system already knows the disposition for a file based on its SHA-256 value, the Defense Center uses the cached disposition rather than querying the Sourcefire 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 Defense Center so the system can take appropriate action the next time it detects that file being transmitted. A changed disposition is called a *retrospective* disposition.

File dispositions 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 disposition. Dispositions have the following TTL values:

- Clean - 4 hours
- Neutral - 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 against the cloud to determine disposition.

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 *global malware whitelist*. When the system detects a file with a SHA-256 on the whitelist, it does not perform a malware lookup or block the file as malware, even if the disposition for the file in the cloud is Malware. You can enable use of the global malware whitelist on a per-file-policy basis. For more information on managing the global malware whitelist, see [Working with the Global Malware Whitelist](#) on page 180.

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, managed devices monitor network traffic for transmissions of specific file types, then either block or allow 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 Sourcefire cloud.

Using File and Malware Events for Analysis

The system generates malware and file events when files are transferred or blocked; you can view these events using the Defense Center's web interface. Additionally, the Context Explorer and the dashboard provide you with different types of high-level views of the files (including malware files) detected by your organization.

To further target your analysis, the *network file trajectory* feature allows you to track individual files' paths of transmission. A file's trajectory page displays summary information about the file, a graphical map of the file's transmission from host to host (including blocked transmissions), and a list of the malware or file events associated with the detection or blocking of those files.

Note that because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances to perform network-based malware protection or view file trajectories for files for which you conduct a malware cloud lookup.

For more information, see the following sections:

- [Configuring Malware Protection and File Control](#) on page 1201
- [Logging Events Based on Malware Protection and File Control](#) on page 1202
- [Integrating FireAMP with the Sourcefire 3D System](#) on page 1202
- [Network-Based AMP vs Endpoint-Based FireAMP](#) on page 1204
- [Working with Network File Trajectory](#) on page 1241

Configuring Malware Protection and File Control

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

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. File rules can not only allow or block files based on simple file type matching, but can also block files based on malware disposition. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer.

As a simple example, you could implement a file policy that blocks your users from downloading executable files. As another example, you could examine downloaded PDFs for malware and block any instances you find. For detailed information on file policies and associating them with access control rules, see [Understanding and Creating File Policies](#) on page 1206 and [Performing File and Intrusion Inspection on Allowed Traffic](#) on page 511.

Because you cannot use a Malware license with a DC500, you cannot use that appliance to apply file policies that perform network-based malware protection. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy that performs network-based malware protection to those appliances.

Logging Events Based on Malware Protection and File Control

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Defense Center logs records of the system's file inspection and handling as file 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 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. Note that endpoint-based malware events generated by FireAMP Connectors (see [Integrating FireAMP with the Sourcefire 3D System](#) on page 1202) do not have corresponding file events.

You can use the Defense Center to view, manipulate, and analyze file and malware events, then communicate your analysis to others. The Context Explorer, dashboards, event viewer, network file trajectory map, and reporting features can give you a deeper understanding of the files and malware detected and blocked in network traffic and on endpoints. You can also use events to trigger correlation policy violations, or alert you via email, SMTP, or syslog. For detailed information on file and malware events, see [Working with File Events](#) on page 1222 and [Working with Malware Events](#) on page 1229.

Because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances to generate or analyze file and malware events associated with malware cloud lookups.

Integrating FireAMP with the Sourcefire 3D System

LICENSE: Any

FireAMP is Sourcefire's enterprise-class advanced malware analysis and protection solution that discovers, understands, and blocks advanced malware outbreaks, advanced persistent threats, and targeted attacks.

If your organization has a FireAMP subscription, individual users install *FireAMP Connectors* on *endpoints*: computers and mobile devices. A FireAMP Connector is a lightweight agent that, among other capabilities, can inspect files upon upload, download, execution, open, copy, move, and so on. These connectors communicate with the Sourcefire cloud to determine if inspected files contain malware.

When a file is positively identified as malware, the cloud sends the threat identification to the Defense Center. The cloud can also send other kinds of information to the Defense Center, including data on scans, quarantines, blocked executions, and cloud recalls. The Defense Center logs this information as malware events.

With a FireAMP deployment, you can not only configure Defense Center-initiated remediations and alerts based on malware events, but you can also use the FireAMP portal (<http://amp.sourcefire.com/>) to help you mitigate the effect of malware. The portal provides a robust, flexible web interface where you control all aspects of your FireAMP deployment and manage all phases of an outbreak. You can:

- configure custom malware detection policies and profiles for your entire organization, as well as perform flash and full scans on all your users' files
- perform malware analysis, including view heat maps, detailed file information, network file trajectory, and threat root causes
- configure multiple aspects of outbreak control, including automatic quarantines, application blocking to stop non-quarantined executables from running, and exclusion lists
- create custom protections, block execution of certain applications based on group policy, and create custom whitelists

For more information, see the following sections:

- [Network-Based AMP vs Endpoint-Based FireAMP](#) on page 1204 compares the malware protection strategies available in the Sourcefire family of products.
- [Working with Sourcefire Cloud Connections for FireAMP](#) on page 1219 explains how to establish communications between the Defense Center and the Sourcefire cloud.

TIP! For detailed information on FireAMP, refer to the online help on the FireAMP portal.

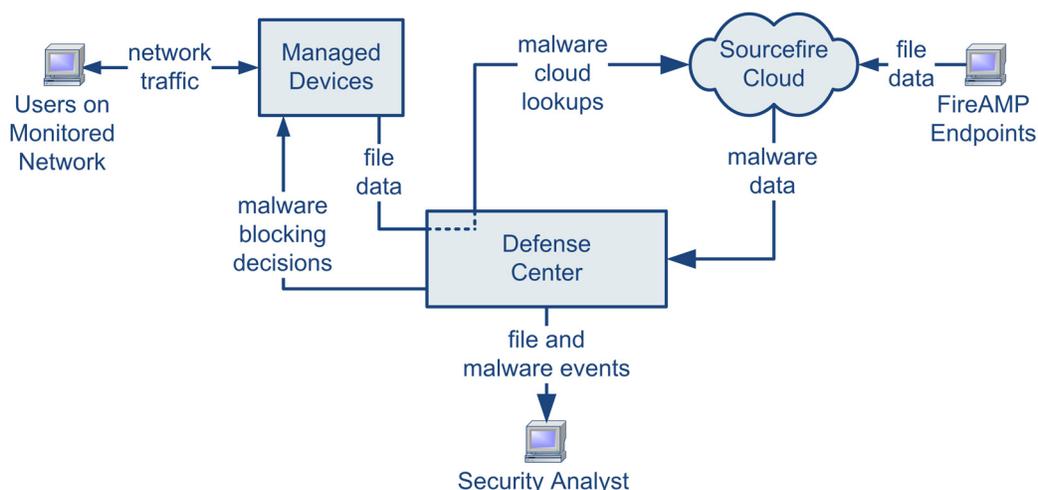
Network-Based AMP vs Endpoint-Based FireAMP

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The following diagram shows how you can use the Defense Center to work with data from both a network-based advanced malware protection strategy and an endpoint-based FireAMP strategy.



Note that because FireAMP malware detection is performed at the endpoint at download or execution time, while managed devices detect malware in network traffic, the information in the two types of malware events is different. For example, endpoint-based malware events contain information on file path, invoking client application, and so on, while malware detections in network traffic contain port, application protocol, and originating IP address information about the connection used to transmit the file.

As another example, for network-based malware events, user information represents the user most recently logged into the host where the malware was destined, as determined by network discovery. On the other hand, FireAMP-reported users represent the user currently logged into the endpoint where the malware was detected, as determined by the local connector.

IMPORTANT! The IP addresses reported in endpoint-based malware events may not be in your network map—and may not even be in your monitored network. Depending on your deployment, network architecture, level of compliance, and other factors, the endpoints where connectors are installed may not be the same hosts as those monitored by your managed devices.

Note that because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances

to perform network-based malware protection or view file trajectories for files for which you conduct a malware cloud lookup.

The following table summarizes the differences between the two strategies.

Network vs Endpoint-Based Malware Protection Strategies

FEATURE	NETWORK-BASED	ENDPOINT-BASED (FIREAMP)
file type detection and blocking method (file control)	in network traffic, using access control and file policies	not supported
malware detection and blocking method	in network traffic, using access control and file policies	on individual endpoints, using an installed connector that communicates with the Sourcefire cloud
network traffic inspected	traffic passing through a managed device	none; connectors installed on endpoints directly inspect files
malware detection robustness	limited file types	all file types
malware analysis choices	Defense Center-based	Defense Center-based, plus additional options on the FireAMP portal
malware mitigation	malware blocking in network traffic, Defense Center-initiated remediations	FireAMP-based quarantine and outbreak control options, Defense Center-initiated remediations
events generated	file events, malware events, and retrospective malware events	malware events
information in malware events	basic malware event information, plus connection data (IP address, port, and application protocol)	in-depth malware event information; no connection data
network file trajectory	Defense Center-based	Defense Center-based, plus additional options on the FireAMP portal
required licenses or subscriptions	Protection license to perform file control; Malware license to perform malware protection	FireAMP subscription (not license-based)

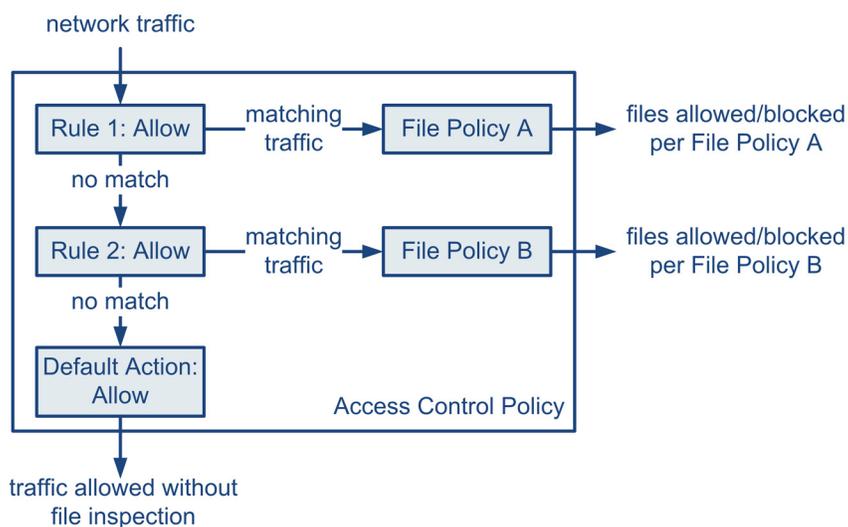
Understanding and Creating File Policies

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

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. Consider the following diagram of a simple access control policy in an inline deployment.



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. File rules can not only allow or block files based on simple file type matching, but can also block files based on malware disposition. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer. You can also enable use of a global malware whitelist to exempt specific files from inspection and blocking.

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 file policies with individual access control rules, you

have granular control over how you identify and block files transmitted on your network. In other words, this association tells the system that before it passes traffic that matches an access control rule's conditions, you first want to inspect the traffic with a file policy.

Keep in mind that the system can perform file inspection on interactively blocked traffic only if the user bypasses the warning and clicks through to the originally requested site. Otherwise, the connection is denied without either file or intrusion inspection; see [Understanding Rule Actions](#) on page 478 and [Adding an HTTP Response Page](#) on page 436.

You can associate different file policies with individual access control rules in the same access control policy. This allows you to match various file and malware detection profiles against different types of traffic on your network. Note, however, that you **cannot** use a file policy to inspect traffic handled by the access control default action.

Also note that because you cannot use a Malware license with a DC500, you cannot use that appliance to apply file policies that perform network-based malware protection. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy that performs network-based malware protection to those appliances.

File and Intrusion Policy Interaction

You can associate both a file policy and an intrusion policy with an access control rule. When you do so, note that the two policies interact in ways that may change how traffic is inspected.

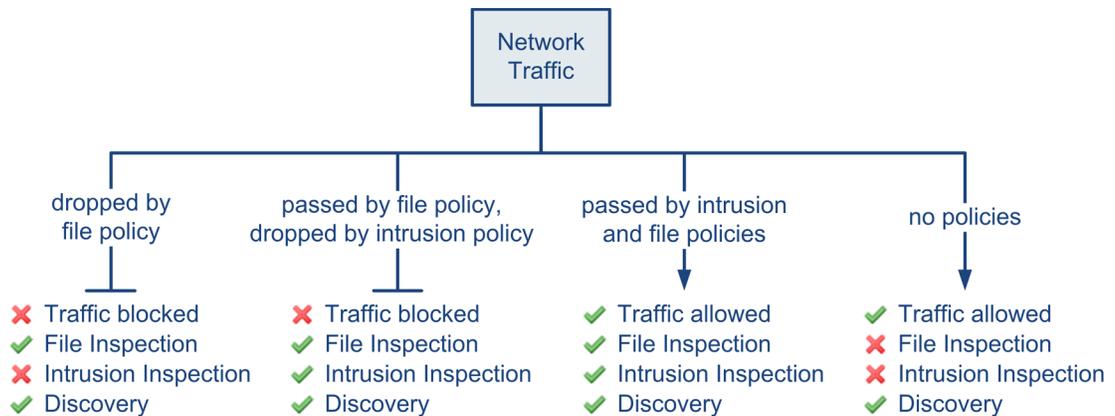
File inspection occurs before any intrusion policy inspection; that is, the system does not inspect files blocked by a file policy for intrusions. Within a file policy, simple blocking by type takes precedence over malware inspection and blocking.

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 any traffic and that is associated with both an intrusion policy and a file policy. The file policy has a rule that matches PDFs and has a Block Files action and a rule that matches executable files and has a Block Malware action. After you apply the policy:

- First, the system blocks the download of all PDF files, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware lookup nor intrusion inspection.
- Next, the system performs malware cloud lookups for downloaded executables. Any executables with a Malware 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.

The diagram below illustrates the types of inspection performed on traffic that meets the conditions of either an Allow or user-bypassed Interactive Block access control rule. For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with a single 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.

Regardless of whether the traffic is inspected or dropped by an intrusion or file policy, the system can inspect it using network discovery; see [Introduction to Network Discovery](#) on page 1250.



Special Considerations for Intrusion Policies when Blocking Malware over FTP

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 file 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

You populate a file policy with file rules. The following table describes the components of a file rule.

File Rule Components

FILE RULE COMPONENT	DESCRIPTION
application protocol	The system can detect and inspect files transmitted via FTP, HTTP, SMTP, IMAP, and POP3. 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, and POP3 traffic for downloaded files; you can inspect outgoing FTP, HTTP, and SMTP traffic for uploaded files.
file categories and types	<p>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.</p> <p>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.</p> <p>WARNING! 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.</p>
file rule action	<p>A file rule's action determines how the system handles traffic that matches the conditions of the rule.</p> <p>IMPORTANT! File rules are evaluated in rule-action, not numerical, order. For more information, see the next section, File Rule Actions and Evaluation Order.</p>

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:

1. *Block Files* rules allow you to block specific file types.
2. *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.
3. *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.
4. *Detect Files* rules allow you to log the detection of specific file types to the database, while still allowing their transmission.

Note that if you select either Block Files or Block Malware, you can choose to reset the connection where the blocked file transfer occurred. Sourcefire recommends that you enable **Reset Connection** to prevent blocked application sessions from remaining open until the TCP connection resets.

File rules with Block Files and Block Malware actions also 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.

Unlike in access control policies, where rules are evaluated in numerical order, file policies handle files in the rule-action order listed above. 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 web interface, these rules are evaluated in the following order:

File Rule Evaluation Order Example

APP. PROTOCOL	DIRECTION	ACTION	RESULT
SMTP	Upload	Block Files	Blocks users from emailing PDF files.
FTP	Download	Block Malware	Blocks the download of malware PDF files via file transfer.
POP3, IMAP	Download	Malware Cloud Lookup	Inspects PDF files received via email for malware.
Any	Any	Detect Files	Detects and logs, but allows the traffic, when users view PDF files on the web (that is, via HTTP).

The Defense Center uses warning icons () to designate conflicting file rules. For details, hover your pointer over a warning icon.

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.

Note that because you cannot use a Malware license with a DC500, you cannot create file rules that use the Block Malware or Malware Cloud Lookup action or use that appliance to apply file policies that contain rules with those actions. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy containing rules with those actions to those appliances.

Logging File and 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, and detected malware files
- *malware events*, which represent detected malware files
- *retrospective malware events*, which are generated when the malware disposition for a previously detected file changes

When a file policy generates a file or malware event, the system automatically logs the end of the associated connection to the Defense Center database, regardless of the logging configuration of the invoking access control rule. 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 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 global malware whitelist
 - **File Block** for Block Files or Block Malware file rules
 - **File Resume Allow** where file transmission was originally blocked by a Block Files or Block Malware rule, then the HTTP session automatically resumed after a new access control policy was applied that allowed the file
- For connections where a file or malware was blocked, the **Action** is **Block**.

As with any kind of event generated by the Sourcefire 3D System, you can view, manipulate, and analyze file and malware events using the Defense Center's web

interface. You can also use malware events to trigger correlation policy violations, or alert you via email, SMTP, or syslog.

IMPORTANT! The Defense Center can also receive malware events using your organization's FireAMP subscription. Because these malware events are generated on endpoints at download or execution time, their information is different from that in network-based malware events.

For more information on connection, file, and malware events, as well as additional details on how they are logged, see:

- [Logging Connection, File, and Malware Information](#) on page 515
- [Working with File Events](#) on page 1222
- [Working with Malware Events](#) on page 1229
- [Understanding Connection Data](#) on page 540

Internet Access and High Availability

The system uses port 32137 to perform malware cloud lookups for network-based AMP. You must open that port outbound on the Defense Center. Additionally, the Defense Center must have direct access to the internet, without a proxy.

Although they share file policies and related configurations, Defense Centers in a high availability pair share neither cloud connections nor file and malware events. To ensure continuity of operations, and to ensure that detected files' malware dispositions are the same on both Defense Centers, both primary and secondary Defense Centers must have access to the cloud.

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.



Name	Last Modified	
Detect All Files Detects all possible malware and file types.	2012-09-25 14:12:52 Modified by "jkpolk"	   
Block Executables	2012-09-25 14:12:32 Modified by "jkpolk"	   

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 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.

Also 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](#) on page 1213
- [Working with File Rules](#) on page 1214
- [Comparing Two File Policies](#) on page 1217

Creating a File Policy

LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

After you create a file policy and populate it with rules, you can use it in an access control policy.

You can enable a global malware whitelist for the policy to exempt files on the list from inspection or blocking by any of the rules in the policy.

Note that because you cannot use a Malware license with a DC500, you cannot create file rules that use the Block Malware or Malware Cloud Lookup action or use that appliance to apply file policies that contain rules with those actions. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy containing rules with those actions to those appliances.

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:

ACCESS: Admin/Access Admin

1. Select **Policies > Files**.

The File Policies page appears.

2. Click **New File Policy**.

The New File Policy dialog box appears.



For a new policy, the web interface indicates that the policy is not in use. If you are editing an in-use file policy, the web 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 [Managing Access Control Policies](#) on page 456.

3. Enter a **Name** and optional **Description** for your new policy, then click **Save**.

The File Policy Rules page appears.

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](#) on page 1214.

Because you cannot use a Malware license with a DC500, you cannot create file rules that use the Block Malware or Malware Cloud Lookup action or use that appliance to apply file policies that contain rules with those actions. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy containing rules with those actions to those appliances.

5. If you have added files to the global malware whitelist and want to exempt them from malware detection and blocking by rules in the file policy, select the **Enable Global Malware Whitelist** check box.

Note that enabling this checkbox on a DC500 or in a policy applied to a Series 2 device has no effect.

6. Click **Save**.

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

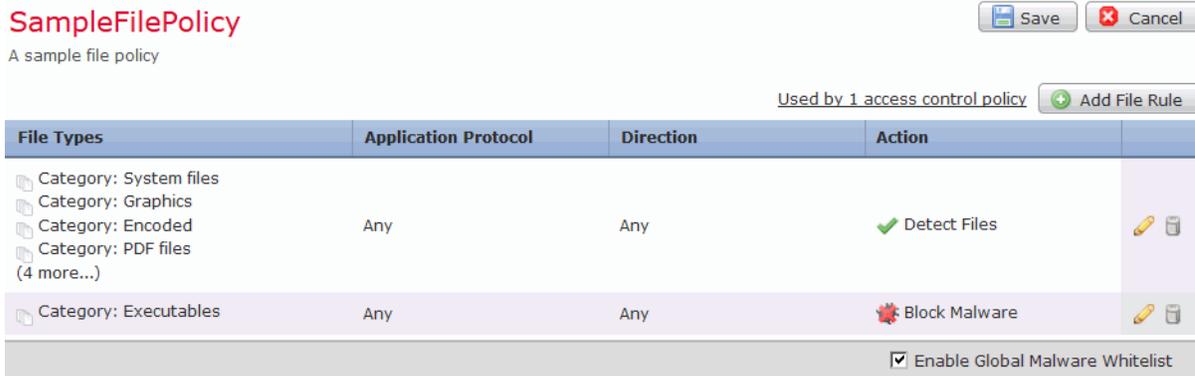
LICENSE: Protection or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

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:

ACCESS: Admin/Access Admin

1. Select Policies > Files.

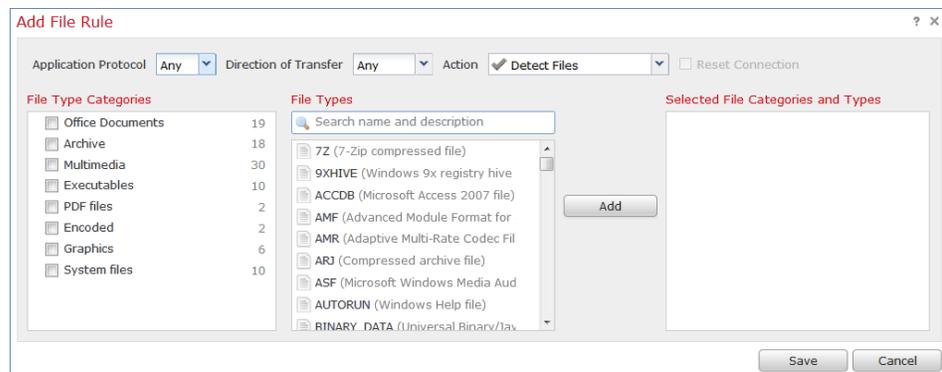
The File Policies page appears.

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](#) on page 1213.
- To add rules to an existing policy, click the edit icon (📎) next to the policy.

3. On the File Policy Rules page that appears, click Add File Rule.

The Add File Rule dialog box appears.



4. Select an **Application Protocol**.

Any, the default, detects files in HTTP, SMTP, IMAP, POP3, and FTP traffic.

5. Select a **Direction of Transfer**.

You can inspect incoming HTTP, IMAP, POP3, and FTP traffic for downloaded files and outgoing HTTP, FTP, and SMTP traffic for uploaded files. Use **Any** to detect files over multiple application protocols, regardless of whether users are sending or receiving.

6. Select a file rule **Action**.

For detailed information on file rule actions, see [File Rule Actions and Evaluation Order](#) on page 1209.

Note that because you cannot use a Malware license with a DC500, you cannot create file rules that use the Block Malware or Malware Cloud Lookup action or use that appliance to apply file policies that contain rules with those actions. Similarly, because you cannot enable a Malware license on a Series 2 device, you cannot apply a file policy containing rules with those actions to those appliances.

7. If you selected either Block Files or Block Malware, enable **Reset Connection** if you want to reset the connection where the blocked file transfer occurs.

IMPORTANT! Sourcefire recommends that you enable **Reset Connection** to prevent blocked application sessions from remaining open until the TCP connection resets.

8. 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.

TIP! Hover your pointer over a file type to view its description.

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.

9. 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.
10. Click **Save**.

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.

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.

▲ Previous ▼ Next (Difference 1 of 24) 📄 Comparison Report 🆕 New Comparison

Detect All Files (2012-09-28 17:01:35 by jkpolk)		Block Executables (2012-09-28 16:58:37 by jkpolk)	
Policy Information			
Name	Detect All Files	Name	Block Executables
Modified	2012-09-28 17:01:35 by jkpolk	Modified	2012-09-28 16:58:37 by jkpolk
File Rules			
Rule 1			
Application Protocol	Any	Application Protocol	Any
Direction	Any	Direction	Any
Action	Detect	Action	Block
File Categories			
Archive			
Encoded			
Graphics			
Multimedia			
Office Documents			
PDF files			
System files			

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:

ACCESS: Admin/Access Admin

1. Select Policies > Files.

The File Policies page appears.

2. Click Compare Policies.

The Select Comparison dialog box appears.



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.

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.

5. Click OK.

The comparison view appears.

6. Optionally, click **Comparison Report** to generate the access control policy comparison report.

The 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.

Working with Sourcefire Cloud Connections for FireAMP

LICENSE: Any

FireAMP is Sourcefire's enterprise-class advanced malware analysis and protection solution. If your organization has a FireAMP subscription, individual users install FireAMP Connectors on their computers and mobile devices. These lightweight agents communicate with the Sourcefire cloud, which in turn communicate with the Defense Center. After you configure the Defense Center to connect to the cloud, you can receive records of scans, malware detections, and quarantines. The records are stored in the Defense Center database as malware events. For more information, see [Understanding Malware Protection and File Control](#) on page 1199.

Each Defense Center in your deployment can connect to the Sourcefire cloud. By default, the cloud sends malware events for all groups within your organization, but you can restrict by group when you configure the connection.

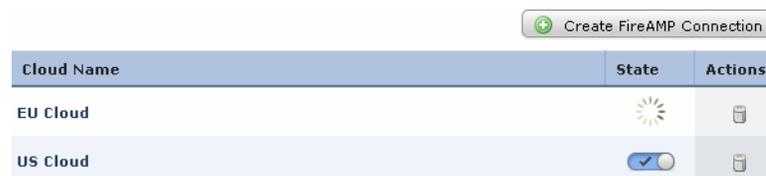
Internet Access and High Availability

The system uses port 443/HTTPS to connect to the Sourcefire cloud to receive endpoint-based malware events. You must open that port, both inbound and outbound, on the Defense Center. Additionally, the Defense Center must have direct access to the internet, without a proxy. The default health policy includes the FireAMP Status Monitor, which warns you if the Defense Center cannot connect to the cloud after an initial successful connection, or if the connection is deregistered using the FireAMP portal.

Cloud connections to receive endpoint-based malware events are **not** shared between members of a high availability pair. To ensure continuity of operations, connect both the primary and secondary Defense Centers to the cloud.

Managing Sourcefire Cloud Connections

Use the Defense Center's FireAMP Management page (**FireAMP > FireAMP Management**) to view and create connections to the Sourcefire cloud, as well as disable and delete those connections.



Cloud Name	State	Actions
EU Cloud		
US Cloud	<input checked="" type="checkbox"/>	

A spinning state icon indicates that the connection is pending, for example, if you configured the connection on the Defense Center, but now must authorize the connection using the FireAMP portal. A failed or denied icon (🚫) indicates that the cloud denied the connection or the connection failed for another reason.

TIP! Click any cloud name to open the FireAMP portal in a new browser window.

For more information, see:

- [Creating a Sourcefire Cloud Connection](#) on page 1220
- [Deleting or Disabling a Sourcefire Cloud Connection](#) on page 1221

Creating a Sourcefire Cloud Connection

LICENSE: Any

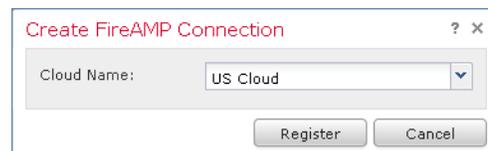
Creating a connection between the Defense Center and the Sourcefire cloud is a two-step process. First, configure the Defense Center to connect to the cloud. Then, log into the FireAMP portal to authorize the connection. If you do not have a FireAMP subscription, you cannot complete the registration process.

To re-register a Defense Center that was restored to factory defaults or reverted while registered to the cloud, you must connect to FireAMP and remove the Defense Center before re-registering it.

To create a Sourcefire cloud connection for FireAMP:

ACCESS: Admin

1. Select **FireAMP > FireAMP Management**.
The FireAMP Management page appears.
2. Click **Create FireAMP Connection**.
The Create FireAMP Connection dialog box appears.



3. Select the **Cloud Name** for the cloud you want to use, then click **Register**.
4. Confirm that you want to continue to the FireAMP portal, then log into the portal.
The Applications page on the portal appears. Use this page to authorize the Sourcefire cloud to send malware events to the Defense Center.

5. Optionally, select specific groups within your organization for which you want to receive malware events.

Select groups only if you want to restrict the events you receive. By default, the Defense Center receives malware events for all groups.

TIP! To manage groups, select **Management > Groups** on the FireAMP portal. For detailed information, refer to the online help on the portal.

6. Click **Allow**.

You are returned to the FireAMP Management page on the Defense Center. Your connection is enabled and the Defense Center begins receiving malware events from the cloud.

Clicking **Deny** also returns you to the Defense Center, where the cloud connection is marked as denied. Similarly, if you navigate away from the Applications page on the FireAMP portal, and neither deny nor allow the connection, the connection is marked as pending on the Defense Center's web interface. The health monitor does **not** alert in either of these situations. If you want to connect to the cloud later, you must delete the failed or pending connection, then recreate it.

Deleting or Disabling a Sourcefire Cloud Connection

LICENSE: Any

Delete a Sourcefire cloud connection if you no longer want to receive malware events from the cloud. To temporarily stop malware events from being sent for a particular connection, you can disable the connection rather than deleting it. In this situation, the cloud stores the events until you re-enable the connection; the cloud then sends the stored events.

WARNING! In rare cases — for example, with a very high event rate or a long-term disabled connection — the cloud may not be able to store all events generated while the connection is disabled.

Note that deregistering a connection using the FireAMP portal (instead of the Defense Center's web interface) stops events from being sent, but does not remove the connection from the Defense Center. Deregistered connections show a failed state on the FireAMP Management page and you must delete them.

To enable or disable a Sourcefire cloud connection using the Defense Center:

ACCESS: Admin

- ▶ On the FireAMP Management page, next to the connection you want to delete, click the slider, then confirm that you want to either enable or disable the connection.

When you enable a connection, the cloud begins sending events to the Defense Center, including any events that occurred while the connection was disabled. The cloud does not send events for disabled connections.

To delete a Sourcefire cloud connection using the Defense Center:

ACCESS: Admin

- ▶ On the FireAMP Management page, next to the connection you want to delete, click the delete icon (), then confirm you want to remove the connection.

The connection is removed and the cloud stops sending events to the Defense Center.

Working with File Events

LICENSE: Protection

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. Note that when the system generates a file event, the system also logs the end of the associated connection to the Defense Center database, regardless of the logging configuration of the invoking access control rule. For more information, see [Understanding and Creating File Policies](#) on page 1206.

IMPORTANT! Files detected in network traffic and identified as malware by the Sourcefire 3D System 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. Endpoint-based malware events do not have corresponding file events. For more information, see [Working with Malware Events](#) on page 1229.

You can use the Defense Center's event viewer to view, search, and delete file events. Additionally, the Files Dashboard provides an at-a-glance view of detailed information about the files (including malware files) detected on your network, using charts and graphs. Network file trajectory offers a more in-depth view of individual files, providing summary information about the file and how it has moved through the network over time. Using file identification data, you can trigger correlation rules and create reports, the latter using either the predefined Files Report template or a custom report template.

For more information, see:

- [Viewing File Events](#) on page 1223
- [Understanding the File Events Table](#) on page 1224
- [Using Geolocation](#) on page 1816
- [Searching for File Events](#) on page 1227

Viewing File Events

LICENSE: Protection

The Sourcefire 3D System's event viewer allows you to view file events in a table, as well as manipulate the event view depending on the information relevant to your analysis.

The page you see when you access file events differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. The system is delivered with the following predefined workflows for file events:

- *File Summary*, the default, provides a quick breakdown of the different file event categories and types, along with any associated malware dispositions.
- *Hosts Receiving Files* and *Hosts Sending Files* provide a list of hosts that have received or sent files, grouped by the associated malware dispositions for those files.

IMPORTANT! Dispositions appear only for files for which the system performed a malware cloud lookup; see [File Rule Actions and Evaluation Order](#) on page 1209.

You can also create a custom workflow that displays only the information that matches your specific needs. For information on specifying a different default workflow, including a custom workflow, see [Configuring Event View Settings](#) on page 2218.

Using the event viewer, you can:

- search for, sort, and constrain events, as well as change the time range for displayed events
- specify the columns that appear (table view only)
- view the host profile associated with an IP address, or the user details and host history associated with a user identity
- view the connections where specific files were detected
- view events using different workflow pages within the same workflow
- view events using a different workflow altogether
- drill down page-to-page within a workflow, constraining on specific values

- bookmark the current page and constraints so you can return to the same data (assuming the data still exists) at a later time
- view the sending and receiving countries and continents for routable IP addresses associated with a file
- view a file's trajectory
- use the SHA-256 value context menu to add a file to the global malware whitelist, with the filename as a description, or to view the full text of a file's SHA-256 value
- create a report template using the current constraints
- delete events from the database
- use the IP address context menu to whitelist, blacklist, or obtain additional available information about a host or IP address associated with a file event

For detailed information on using the event viewer, including creating custom workflows, see [Understanding and Using Workflows](#) on page 1791.

To view file events:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Files**.

The first page of your default file events workflow appears. For information on the columns that appear, see [Understanding the File Events Table](#) on page 1224.

TIP! To quickly view the connections where specific files were detected, select the files using the check boxes in the event viewer, then click the **Connections** tab above the column headings. For more information, see [Navigating Between Workflows](#) on page 1836.

Understanding the File Events Table

LICENSE: Protection

The Defense Center logs a file event when a managed device detects or blocks a file being transmitted in monitored network traffic, according to the settings in an applied file policy.

The table view of file events, which is the final page in predefined file event workflows, and which you can add to custom workflows, includes a column for each field in the files table. Some fields in the table view of file events are disabled by default. To enable a field for the duration of your session, click the expand arrow (▶) to expand the search constraints, then click the column name under **Disabled Columns**. The [File Event Fields](#) table below describes the file event fields.

Keep in mind that although you can perform file control with only a Protection license, a Malware license allows you to perform advanced malware protection for certain file types and track files transferred on your network.

File Event Fields

FIELD	DESCRIPTION
Time	The date and time the event was generated.
Action	The action associated with the file policy rule that detected the file.
Sending IP	The IP address of the host sending the detected file.
Sending Country	The country of the host sending the detected file. Note that the DC500 Defense Center does not support this feature.
Receiving IP	The IP address of the host receiving the detected file.
Receiving Country	The country of the host receiving the detected file. Note that the DC500 Defense Center does not support this feature.
Sending Port	The source port used by the traffic where the file was detected.
Receiving Port	The destination port used by the traffic where the file was detected.
User	The user logged into the host (Receiving IP) where the file was destined. Note that because the user is associated with the destination host, users are not associated with file events where the user uploaded a file.
File Name	The name of the file.
Disposition	If the system performed a malware cloud lookup for this file, the cloud-determined disposition of the file: Malware , Clean , Neutral , or Unavailable . If a Detect Files or Block Files rule handled the file and the Defense Center did not perform a malware cloud lookup, the event has the disposition N/A .

File Event Fields (Continued)

FIELD	DESCRIPTION
SHA256	If this file was detected as the result of a Malware Cloud Lookup or Block Malware rule, the SHA-256 hash value of the file.
Type	The type of file, for example, HTML or MSEXE .
Category	The general category of file type, for example: Office Documents , Archive , Multimedia , Executables , PDF files , Encoded , Graphics , or System Files .
Size (KB)	The size of the file, in kilobytes. Note that if the system determines the file type of a file before the file is fully received, the file size may not be calculated and this field is blank.
URI	The originating URI of the file, for example, the URL where a user downloaded it.
Application Protocol	The application protocol used by the traffic in which a managed device detected the file.
Application Protocol, Client, or Web Application Category or Tag	Criteria that characterize the application to help you understand the application's function. For more information, see the Application Characteristics table on page 1264.
Client	The client application used in the connection to transmit a file.
Web Application	For files transmitted using HTTP, the web application (content or requested URL) detected in the connection and used to transmit the file.
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. For more information, see the Application Characteristics table on page 1264.
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. For more information, see the Application Characteristics table on page 1264.

File Event Fields (Continued)

FIELD	DESCRIPTION
Message	For files where a malware disposition has changed, that is, for files associated with retrospective malware events, information about when and how the disposition changed.
File Policy	The file policy that detected the file.
Device	The name of the device that detected the file.
Count	The number of events that match the information in each row. This field appears after you apply a constraint that creates two or more identical rows.

Searching for File Events

LICENSE: Protection

Using the Defense Center's Search page, you can search for specific file events, display the results in the event viewer, and save your search criteria to reuse later. Custom Analysis dashboard widgets, report templates, and custom user roles can also use saved searches.

The screenshot shows the 'Search Information' form. At the top, there is a note: 'Note: If a search name is not specified, an automatically generated name will be used.' Below this, the form includes several fields: 'Table' is a dropdown menu set to 'File Events'; 'Name' is an empty text box with the placeholder 'Search 1, My Search'; 'Save As Private' is a checked checkbox; 'Constraint' is a section header; 'Action' is a text box with 'Detect, Block' as the value; 'Sending IP' is a text box with a green plus icon and the value '192.168.1.0/24, 192.168.1.3, 2001:0db8:85a3::1370'; 'Message' is a text box; 'File Policy' is a text box with 'My File Policy' as the value; 'Device' is a text box with 'device1.example.com, *.example.com, 192.168.1.3' as the value. At the bottom, there are two buttons: 'Search' and 'Save As New Search'.

Keep in mind that your search results depend on the available data in the events you are searching. In other words, depending on the available data, your search constraints may not apply. For example, the **Disposition** and **SHA256** fields are populated only for files for which the Defense Center performed a malware cloud lookup.

Note that because the DC500 does not support geolocation, searches using these fields from a DC500 return no results, regardless of whether geolocation information was detected.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Special Search Syntax for File Events

To supplement the general search syntax listed above, the following table describes some special search syntax for file events.

File Event Special Search Syntax

SEARCH CRITERION	SPECIAL SYNTAX
Sending/Receiving Continent	The system returns all events where either the Sending Continent or the Receiving Continent matches the continent you specify.
Sending/Receiving Country	The system returns all events where either the Sending Country or the Receiving Country matches the country you specify.
Sending/Receiving IP	The system returns all events where either the Sending IP or the Receiving IP matches the IP address you specify.
URI or Message	The system performs a partial match, that is, you can search for all or part of the field contents without using asterisks.

To search for file events:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.

2. From the **Table** drop-down list, select **File Events**.
The page reloads with the appropriate constraints.
3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields.
See the [File Event Fields table](#) on page 1225 for information on the fields in the file events table.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.
6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in your default malware events workflow, constrained by the current time range.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**).

Working with Malware Events

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The system logs malware events to the Defense Center database when:

- a managed device detects a file in network traffic that is then identified as malware by a malware cloud lookup
- the system learns that a file's malware disposition has changed; these are called retrospective malware events
- a FireAMP Connector installed on an endpoint in your organization detects a threat and communicates that threat to the Sourcefire cloud

Because FireAMP malware detection is performed at the endpoint at download or execution time, while managed devices detect files in network traffic, the information in these malware events is different. Retrospective malware events also contain slightly different data than other network-based malware events, or endpoint-based malware events.

The following sections briefly describe the different kinds of malware events. For information on the overall malware detection process, see [Understanding Malware Protection and File Control](#) on page 1199.

Endpoint-Based (FireAMP) Malware Events

If your organization has a FireAMP subscription, individual users install FireAMP Connectors on their computers and mobile devices. These lightweight agents communicate with the Sourcefire cloud, which in turn communicates with your Defense Center; see [Working with Sourcefire Cloud Connections for FireAMP](#) on page 1219. The cloud can send notification of threats, as well other kinds of information including data on scans, quarantines, blocked executions, and cloud recalls. The Defense Center logs this information to its database as malware events.

IMPORTANT! The IP addresses reported in endpoint-based malware events may not be in your network map — and may not even be in your monitored network at all. Depending on your deployment, level of compliance, and other factors, endpoints in your organization where FireAMP Connectors are installed may not be the same hosts as those monitored by your managed devices.

Malware Events Based on Network Traffic

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

With a Malware license, your managed devices can detect malware in network traffic as part of your overall access control configuration; see [Understanding and Creating File Policies](#) on page 1206.

When a managed device detects one of a set of specific file types, the Defense Center performs a malware cloud lookup, which returns a disposition to the Defense Center of **Malware**, **Clean**, **Neutral**, or **Unavailable**. The Defense Center logs records of these files' detection and dispositions, along with other contextual data, as malware events.

IMPORTANT! Files detected in network traffic and identified as malware by the Sourcefire 3D System generate both a file event and a malware event. This occurs because to detect malware in a file, the system must first detect the file itself. For more information, see [Working with File Events](#) on page 1222.

Retrospective Malware Events

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

For malware files detected in network traffic, file dispositions can change. For example, the Sourcefire 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.

The cloud notifies the Defense Center if the disposition changes for a file for which you performed a malware lookup in the last week. Then, two things happen:

- The Defense Center generates a new retrospective malware event.
This new retrospective malware event represents a disposition change for all files detected in the last week that have the same SHA-256 hash value. For that reason, these events contain limited information: the date and time the Defense Center was notified of the disposition change, the new disposition, the SHA-256 hash value of the file, and the threat name. They do not contain IP addresses or other contextual information.

- The Defense Center changes the disposition for previously detected files with the retrospective event's associated SHA-256 hash value.

If a file's disposition changes to malware, the Defense Center logs a new malware event to its database. Except for the new disposition, the information in this new malware event is identical to that in the file event generated when the file was initially detected.

If a file's disposition changes to clean, the Defense Center does not remove the malware event from the malware table. Instead, the event simply reflects the change in disposition. This means that files with clean dispositions can appear in the malware table, but only if they were originally thought to be malware. Files that were never identified as malware appear only in the files table.

In either case, the malware event's **Message** indicates how and when the disposition changed, for example:

```
Retrospective Event, Mon Oct 1 20:44:00 2012 (UTC), old  
Disp: Neutral, New Disp: Malware
```

Using Malware Events

You can use the Defense Center's event viewer to view, search, and delete malware events. Additionally, the Files Dashboard and Context Explorer provide an at-a-glance view of detailed information about the files (including malware files) detected on your network, using charts and graphs. Network file trajectory offers a more in-depth view of individual malware files, providing summary information about the file and how it has moved through the network over time. Using malware detection data, you can trigger correlation rules and create reports, the latter using either the predefined Malware Report template or a custom report template.

For more information, see:

- [Viewing Malware Events](#) on page 1232
- [Understanding the Malware Events Table](#) on page 1233
- [Searching for Malware Events](#) on page 1239

Viewing Malware Events

LICENSE: Malware or Any

The Sourcefire 3D System's event viewer allows you to view malware events in a table, as well as manipulate the event view depending on the information relevant to your analysis.

The page you see when you access malware events differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. The system is delivered with the following predefined workflows for malware events:

- *Malware Summary*, the default, provides a list of detected malware, grouped by individual threat.
- *Malware Event Summary* provides a quick breakdown of the different malware event types and subtypes.
- *Hosts Receiving Malware* and *Hosts Sending Malware* provide a list of hosts that have received or sent malware, grouped by the associated malware dispositions for those files. Note that dispositions appear only for files detected as the result of Malware Cloud Lookup or Block Malware file rules.
- *Applications Introducing Malware* provides a list of the client applications that accessed or executed the malware detected on endpoints in your organization. From this list, you can drill down into the individual malware files accessed by each parent client.

You can also create a custom workflow that displays only the information that matches your specific needs. For information on specifying a different default workflow, including a custom workflow, see [Configuring Event View Settings](#) on page 2218.

Using the event viewer, you can:

- search for, sort, and constrain events, as well as change the time range for displayed events
- specify the columns that appear (table view only)
- view the host profile associated with an IP address, or the user details and host history associated with a user identity
- view the connections where specific malware was detected (for network-based malware events only)
- view events using different workflow pages within the same workflow
- view events using a different workflow altogether
- drill down page-to-page within a workflow, constraining on specific values
- bookmark the current page and constraints so you can return to the same data (assuming the data still exists) at a later time
- view geolocation information for routable IP addresses associated with a file

- view a file's trajectory
- create a report template using the current constraints
- delete events from the database
- use the SHA-256 value context menu to add a file to the global malware whitelist, with the filename as the description, or to view the full text of a file's SHA-256 value
- use the IP address context menu to whitelist, blacklist, or obtain additional available information about a host or IP address associated with a malware event

Note that neither Series 2 devices nor the DC500 Defense Center support network-based malware protection, which can affect the data displayed. For example, a Series 3 Defense Center managing only Series 2 devices can display only endpoint-based malware events.

For detailed information on using the event viewer, including creating custom workflows, see [Understanding and Using Workflows](#) on page 1791.

To view malware events:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Malware**.

The first page of your default malware events workflow appears. For information on the columns that appear, see [Understanding the Malware Events Table](#) on page 1233.

Understanding the Malware Events Table

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The system logs malware events to the Defense Center database when a FireAMP Connector installed on an endpoint in your organization detects a threat, or a managed device detects a file in network traffic that is then identified as malware by a malware cloud lookup. The system also logs retrospective malware events when it learns that a file's malware disposition has changed. Note that neither Series 2 devices nor the DC500 Defense Center support network-based malware protection, which can affect the data displayed. For example, a Series 3 Defense Center managing only Series 2 devices can display only endpoint-based malware events. For more information, see [Understanding Malware Protection and File Control](#) on page 1199 and [Working with Malware Events](#) on page 1229.

The table view of malware events, which is the final page in predefined malware event workflows, and which you can add to custom workflows, includes a column for each field in the files table. Some fields in the table view of malware events are disabled by default. To enable a field for the duration of your session, click the

expand arrow (▶) to expand the search constraints, then click the column name under **Disabled Columns**.

Keep in mind that not every field is populated for every event; the different types of malware event can contain different information. For example, because FireAMP malware detection is performed at the endpoint at download or execution time, endpoint-based malware events contain information on file path, invoking client application, and so on. In contrast, because managed devices detect malware files in network traffic, their associated malware events contain port, application protocol, and originating IP address information about the connection used to transmit the file.

The following table lists each malware event field, and indicates whether the system displays information in that field, depending on the malware event type. Note that the DC500 Defense Center does not support sending or receiving continent or country geolocation information.

Malware Event Fields

FIELD	DESCRIPTION	NETWORK	ENDPOINT	RETROSPECTIVE FROM CLOUD
Time	The date and time the event was generated.	X	X	X
Action	The file rule action associated with the rule action for the rule the file matched.	X		X
Sending IP	The IP address of the host sending detected malware.	X		
Sending Continent	The continent of the host sending detected malware.	X		X
Sending Country	The country of the host sending detected malware.	X		
Receiving IP	For network-based malware events, the IP address of the host receiving detected malware. For endpoint-based malware events, the IP address of the endpoint where the FireAMP Connector is installed and where the malware event occurred.	X	X	
Receiving Continent	The continent of the host receiving detected malware.	X		X
Receiving Country	The country of the host receiving detected malware.	X		

Malware Event Fields (Continued)

FIELD	DESCRIPTION	NETWORK	ENDPOINT	RETROSPECTIVE FROM CLOUD
Sending Port	The source port used by the traffic in which a managed device detected malware.	X		
Receiving Port	The destination port used by the traffic in which a managed device detected malware.	X		
User	<p>The user of the host (Receiving IP) where the malware event occurred.</p> <p>For network-based malware events, this user is determined by network discovery. Because the user is associated with the destination host, users are not associated with malware events where the user uploaded a malware file.</p> <p>For endpoint-based malware events, FireAMP Connectors determine user names. FireAMP users cannot be tied to user discovery or control. They do not appear in the Users table, nor can you view details for these users.</p>	X	X	
Event Type	The type of malware event. For a full list of event types, see Malware Event Types on page 1238.	X	X	X
Event Subtype	The FireAMP action that led to malware detection, for example, Create , Execute , Move , or Scan .		X	
Threat Name	The name of the detected malware.	X	X	X
File Name	The name of the malware file.	X	X	
File Disposition	The results of the system's malware cloud lookup for a device-detected file: Malware , Clean , Neutral , or Unavailable . Note that clean files appear in the malware table only if they were changed to clean; see Retrospective Malware Events on page 1230.	X		X
File SHA256	The SHA-256 hash value of the file.	X	X	X
File Path	The file path of the malware file, not including the file name.		X	

Malware Event Fields (Continued)

FIELD	DESCRIPTION	NETWORK	ENDPOINT	RETROSPECTIVE FROM CLOUD
File Type	The file type of the malware file, for example, HTML or MSEXEC.	X	X	
File Type Category	The general category of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.	X	X	
File Timestamp	The time and date the malware file was created.		X	
File Size (KB)	The size of the malware file, in kilobytes.	X	X	
File URI	The originating URI of the malware file, for example, the URL where a user downloaded it.	X		
Application File Name	The client application accessing the malware file when detection occurred. These applications are not tied to network discovery or application control.		X	
Application File SHA256	The SHA-256 hash value of the parent file accessing the FireAMP-detected or quarantined file when detection occurred.		X	
Application Protocol	The application protocol used by the traffic in which a managed device detected a malware file.	X		
Application Protocol, Client, or Web Application Category or Tag	Criteria that characterize the application to help you understand the application's function. For more information, see the Application Characteristics table on page 1264.	X		X
Client	The client application that runs on one host and relies on a server to send a file.	X		X
Web Application	The application that represents the content or requested URL for HTTP traffic detected in the connection.	X		X

Malware Event Fields (Continued)

FIELD	DESCRIPTION	NETWORK	ENDPOINT	RETROSPECTIVE FROM CLOUD
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. For more information, see the Application Characteristics table on page 1264.	X		X
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. For more information, see the Application Characteristics table on page 1264.	X		X
Detector	The FireAMP detector that identified the malware, such as ClamAV, SPERO, or SHA.		X	
Message	Any additional information associated with the malware event. For network-based malware events, this field is populated only for files whose disposition has changed; see Retrospective Malware Events on page 1230.	X	X	
FireAMP Cloud	The name of the Sourcefire cloud where the event originated.		X	
Device	For network-based malware events, the name of the device that detected the malware file. For endpoint-based malware events and retrospective malware events generated by the cloud, the name of the Defense Center.	X	X	X
Count	The number of events that match the information in each row. This field appears after you apply a constraint that creates two or more identical rows.	n/a	n/a	n/a

Malware Event Types

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

For network-based malware events, the event type can be one of:

- Threat Detected in Network File Transfer
- Threat Detected in Network File Transfer (retrospective)

An endpoint-based malware event can have any of the following types:

- Blocked Execution
- Cloud Recall Quarantine
- Cloud Recall Quarantine Attempt Failed
- Cloud Recall Quarantine Started
- Cloud Recall Restore from Quarantine
- Cloud Recall Restore from Quarantine Failed
- Cloud Recall Restore from Quarantine Started
- Quarantine Failure
- Quarantined Item Restored
- Quarantine Restore Failed
- Quarantine Restore Started
- Scan Completed, No Detections
- Scan Completed With Detections
- Scan Failed
- Scan Started
- Threat Detected
- Threat Detected in Exclusion
- Threat Quarantined

If a file's trajectory map contains malware events, the events are one of the following types: Threat Detected in Network File Transfer, Threat Detected in Network File Transfer (retrospective), Threat Detected, Threat Detected in Exclusion, and Threat Quarantined. See [Working with Network File Trajectory](#) on page 1241 for more information.

Note that neither Series 2 devices nor the DC500 Defense Center support network-based malware protection, which can affect the data displayed. For example, a Series 3 Defense Center managing only Series 2 devices can display only endpoint-based malware events.

Searching for Malware Events

LICENSE: Malware or Any

Using the Defense Center's Search page, you can search for specific malware events, display the results in the event viewer, and save your search criteria to reuse later. Custom Analysis dashboard widgets, report templates, and custom user roles can also use saved searches.

Saved Searches

--- New Search ---
Threats Detected (Sourcefire)

Load Delete

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: Malware Events

Name: Search 1, My Search

Save As Private:

Constraint

Action: Malware Block, Malware Cloud Lookup

Sending IP: 192.168.1.0/24, 192.168.1.3, 2001:db8:85a3::1370

Receiving IP: 192.168.1.0/24, 192.168.1.3, 2001:db8:85a3::1370

Sending / Receiving IP: 192.168.1.0/24, 192.168.1.3, 2001:db8:85a3::1370

Sending Country: USA, United States, United*

Receiving Country: USA, United States, United*

Business Relevance: Very Low, High

Message

FireAMP Cloud: US Cloud

Device: device1.example.com, *.example.com, 192.168.1.3

Search Save As New Search

Searches delivered with the system, labeled with (sourcefire) in the Saved Searches list, serve as examples.

Keep in mind that your search results depend on the available data in the events you are searching. In other words, depending on the available data, your search constraints may not apply. For example, because endpoint-based malware events are not generated as a result of managed devices inspecting network traffic, they do not contain connection information (port, application protocol, and so on).

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
- Click the add object icon (🟢+) that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Special Search Syntax for Malware Events

To supplement the general search syntax listed above, the following table describes some special search syntax for malware events.

Malware Event Special Search Syntax

SEARCH CRITERION	SPECIAL SYNTAX
Sending/Receiving IP	The system returns all events where either the Sending IP or the Receiving IP matches the IP address you specify.
Event Type	When searching for events with a specific malware event type (see Malware Event Types on page 1238), enclose the event type in quotation marks, for example, " Scan Completed with Detection ". Otherwise, the system performs a partial match. That is, if you search using the same string but do not use quotation marks, the system returns events with the following types: <ul style="list-style-type: none">• Scan Completed, No Detections• Scan Completed With Detection
Initiator/Responder Continent	The system returns all events where either the Initiator Continent or the Responder Continent matches the continent you specify.
Initiator/Responder Country	The system returns all events where either the Initiator Country or the Responder Country matches the country you specify.
URI or Message	The system performs a partial match, that is, you can search for all or part of the field contents without using asterisks.

To search for malware events:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Malware Events**.
The page reloads with the appropriate constraints.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields.
See the [Malware Event Fields table](#) on page 1234 for information on the fields in the malware events table.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.
6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in your default malware events workflow, constrained by the current time range.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**).

Working with Network File Trajectory

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The network file trajectory feature maps how hosts transferred files, including malware files, across your network. You can use the map to determine which hosts may have transferred malware, which hosts are at risk, and observe file transfer trends.

The trajectory map charts file transfer data, the disposition of the file, and if a file transfer was blocked or the file was quarantined. The data used to build the map can come from network-based malware events (any file event for which the system performed a malware cloud lookup and returned a malware disposition) and certain endpoint-based malware events related to detecting and blocking malware (any Threat Detected or Threat Quarantined event type). Vertical lines between data points represent file transfers between hosts. Horizontal lines connecting the data points show a host's file activity over time.

You can track the transmission of any file type for which the system can perform a malware cloud lookup. To directly access a file's trajectory, you can use the Network File Trajectory List page (**Analysis > Files > Network File Trajectory**) and locate specific files. Additionally, if you are analyzing an intrusion and want to

review the trajectory for a related file, you can access the file's trajectory from the Context Explorer, dashboard, or event views of connection, file, or malware events.

The data a single trajectory map displays depends on the licenses applied to your appliance. The following table lists the licenses necessary to track different types of file trajectory.

License Requirements for Network File Trajectory

To view...	You need the following license...
network-based file and malware trajectories	Malware
endpoint-based threat and quarantine tracking	Any (you must have a FireAMP subscription)

See [Understanding Malware Protection and File Control](#) on page 1199 for more information.

Note that because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances to perform network-based malware protection or view file trajectories for files for which you conduct a malware cloud lookup. You can, however, still view file trajectories for endpoint-based threat and quarantine tracking.

For more information, see the following sections:

- [Reviewing Network File Trajectory](#) on page 1242
- [Analyzing Network File Trajectory](#) on page 1244

Reviewing Network File Trajectory

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

As you review file and malware events, you can view a file's trajectory map from the Context Explorer, properly configured dashboard widgets, and various event views. You can also review the most recently viewed network file trajectories and the most recently detected malware from the Network File Trajectory List page.

For more information, see the following sections:

- [Viewing the Top File Names Graph](#) on page 137
- [Drilling Down on Context Explorer Data](#) on page 148
- [Understanding the Custom Analysis Widget](#) on page 79
- [Understanding the File Events Table](#) on page 1224
- [Understanding the Malware Events Table](#) on page 1233

- [Information Available in Connection Events](#) on page 550
- [Accessing Network File Trajectory](#) on page 1243

Accessing Network File Trajectory

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Network File Trajectory List page allows you to locate files that have a SHA-256 hash value, whether to analyze the most recently detected malware, or to track a specific threat.

Recently Viewed Files

Time	File SHA256	File Names	File Type	Disposition	Events
2013-01-02 14:...	a0fe58d8...115f4f3c	Worm.Sober	MSEXE	Malware	6
2013-01-02 13:...	5988b530...578031c9	showcase.swf	SWF	Neutral	1
2013-01-02 13:...	5d846ea6...69d8d35a	wtsclock001.swf	SWF	Clean	2
2013-01-02 13:...	ff72d314...a93bef2c	Nw2tpVaQsRg	SWF	Neutral	2
2013-01-02 13:...	a0b13fb3...e5d4192e	tRhkE-AzBsE	SWF		2

Recent Malware

Time	File SHA256	File Names	File Type	Disposition	Events
2013-01-02 14:...	a0fe58d8...115f4f3c	Worm.Sober	MSEXE	Malware	6

The page displays the malware most recently detected on your network, as well as the files whose trajectory maps you have most recently viewed. From these lists, you can view when the file was most recently seen on the network, the file's SHA-256 hash value, name, type, current disposition, and the number of events associated with the file. For more information on the fields, see [Understanding the File Events Table](#) on page 1224.

The page also contains a search box that lets you locate files, either based on SHA-256 hash value or file name, or by the IP address of the host that transferred or received a file. After you locate a file, you can click the **File SHA256** value to view the detailed trajectory map. See [Analyzing Network File Trajectory](#) on page 1244 for more information.

Note that because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances to view file trajectories for files for which you conduct a malware cloud lookup.

To locate a file from the Network File Trajectory List page:

ACCESS: Any

1. Select **Analysis > Files > Network File Trajectory.**

The Network File Trajectory List page appears, displaying the lists of recently viewed files and recent malware.

2. Optionally, you can type a complete SHA-256 hash value, host IP address, or file name of a file you want to track into the search field and press Enter.

The Query Results page appears listing all files that match the search. If only one result matches, the Network File Trajectory page for that file appears.

Analyzing Network File Trajectory

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

You can trace a file through the network by viewing the detailed network file trajectory. The file's trajectory presents summary information about a file, displays the map charting data points over time, and also lists the event data tied to the data points in a table. Using the table and the map, you can pinpoint specific file events, hosts on the network that transferred or received this file, related events in the map, and other related events in a table constrained on selected values.

Note that because you cannot use a Malware license with a DC500, nor can you enable a Malware license on a Series 2 device, you cannot use those appliances to view file trajectories for files for which you conduct a malware cloud lookup.

For more information, see the following sections:

- [Summary Information](#) on page 1244
- [Trajectory Map](#) on page 1246
- [Events Table](#) on page 1249

Summary Information

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

A file's trajectory page displays basic information about the file, including file identification information, when the file was first seen and most recently seen on the network, the number of related events and hosts associated with the file, and the file's current disposition.

TIP! To view related file events, click a field value link. The first page in the File Events default workflow opens in a new window, displaying all file events that also contain the selected value.

The following table describes the summary information fields.

Network File Trajectory Summary Information Fields

NAME	DESCRIPTION
File SHA256	The SHA-256 hash value of the file. The hash is displayed by default in a condensed format. To view the full hash value, hover your pointer over it. If multiple SHA-256 hash values are associated with a file name, hover your pointer over the link to view all of the hash values.
File Names	The names of the file associated with the event, as seen on the network. If multiple file names are associated with a SHA-256 hash value, the most recent detected file name is listed. You can expand this to view the remaining file names by clicking more .
File Type	The file type of the file, for example, HTML or MSEXE .
File Category	The general category of file type, for example, Office Documents or System Files .
Parent Application	The client application accessing the malware file when detection occurred. These applications are not tied to network discovery or application control. This field only appears for endpoint-based malware events.
First Seen	The first time a managed device or FireAMP Connector detected the file, and the IP address of the host that first uploaded the file.
Last Seen	The most recent time a managed device or FireAMP Connector detected the file, and the IP address of the host that last downloaded the file.
Event Count	The number of events seen on the network associated with the file, and the number of events displayed in the map if there are more than 250 detected events.
Seen On	The number of hosts that either sent or received the file. Because one host can upload and download a file at different times, the total number of hosts may not match the total number of senders plus the total number of receivers in the Seen On Breakdown field.

Network File Trajectory Summary Information Fields (Continued)

NAME	DESCRIPTION
Seen On Breakdown	The number of hosts that sent the file, followed by the number of hosts that received the file.
Current Disposition	The results of the system's malware cloud lookup for a device-detected file: Malware , Clean , Neutral , or Unavailable if the cloud lookup failed. This field only appears for network-based malware events.
Threat Name	Name of the malware threat associated with the file. This field only appears for endpoint-based malware events.

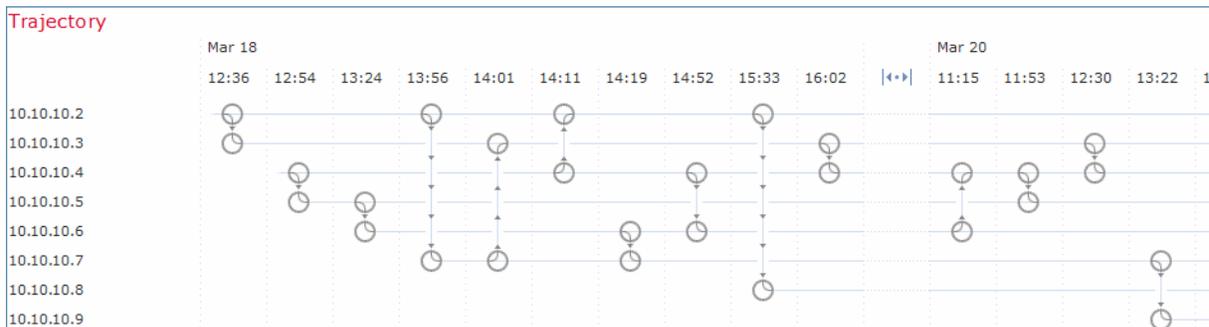
Trajectory Map

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

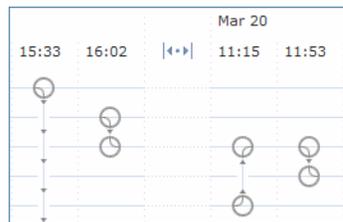
A file's trajectory map visually tracks a file from the first detection on your network to the most recent. The map shows when hosts transferred or received the file, how often they transferred the file, and when the file was blocked or quarantined. The map also shows how often file events occurred for the file and when the system assigned the file a disposition or retrospective disposition. You can select a data point in the map and highlight a path that traces back to the first instance the host transferred that file; this path also intersects with every occurrence involving the host as either sender or receiver of the file. The following screenshot shows an example trajectory map:



The map's y-axis contains a list of all host IP addresses that have interacted with the file. The IP addresses are listed in descending order based on when the system first detected the file on that host. Each row contains all events associated with that IP address, whether a single file event, file transfer, or retrospective event. The x-axis contains the date and time the system detected

each event. The timestamps are listed in chronological order. If multiple events occurred within a minute, all are listed within the same column. You can scroll the map horizontally and vertically to view additional events and IP addresses.

The map displays up to 250 events associated with the file SHA-256 hash. If there are more than 250 events, the map displays the first 10, then truncates extra events with an arrow icon (|<+>|). The map then displays the remaining 240 events. The following screenshot shows events truncated with the arrow icon:

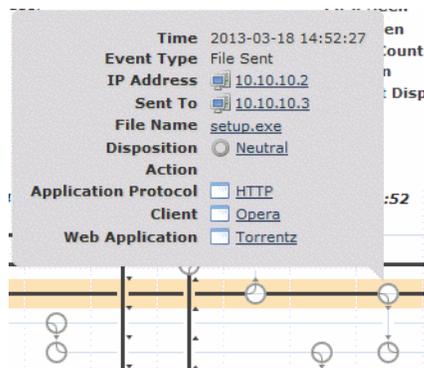


You can view all events not displayed in the File Summary event view by clicking the arrow icon (|<+>|). The first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. If endpoint-based malware events are not displayed, you must switch to the Malware Events table to view these.

Each data point represents an event plus the file disposition, as described in the legend below the map. For example, a Malware Block event icon combines the Malicious Disposition icon and the Block Event icon.

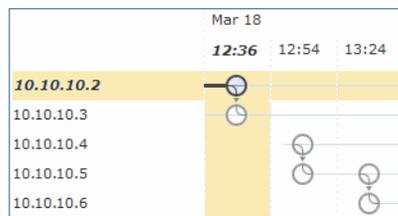
Endpoint-based malware events include one icon. A retrospective event displays an icon in the column for each host on which the file is detected. File transfer events always include two icons, one file send icon and one file receive icon, connected by a vertical line. Arrows indicate the file transfer direction from sender to receiver.

You can view summary information from the event icon by hovering your pointer over the event icon (⊞). The displayed summary information matches the information displayed in the Events table. The following screenshot shows an event icon's summary information:



If you click any event summary information link, the first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type the File Summary event view opens in a new window, displaying all file events that match on the criteria value you clicked.

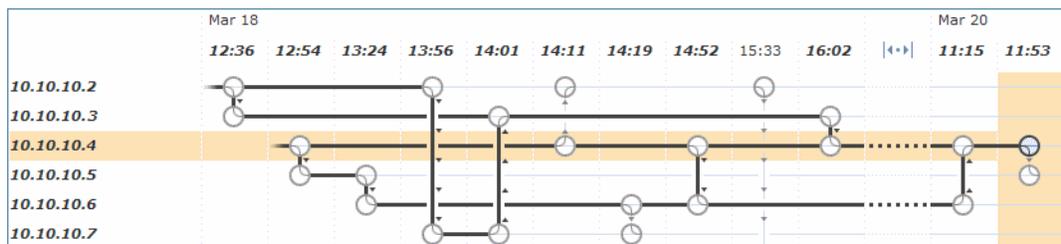
To locate the first time a file event occurred involving an IP address, click the address. This highlights a path to that data point, as well as any intervening file events and IP addresses related to the first file event. The corresponding event in the Events table is also highlighted. The map scrolls to that data point if not currently visible. The following screenshot shows the path highlighted after clicking an IP address:



To track a file's progress through the network, you can click any data point to highlight a path that includes all data points related to the selected data point. This includes data points associated with the following types of events:

- any file transfers in which the associated IP address was either sender or receiver
- any endpoint-based malware events involving the associated IP address
- if another IP address was involved, all file transfers in which that associated IP address was either sender or receiver
- if another IP address was involved, any endpoint-based malware events involving the other IP address

The following screenshot shows the path highlighted after clicking an event icon:



All IP addresses and timestamps associated with any highlighted data point are also highlighted. The corresponding event in the Events table is also highlighted. If a path includes truncated events, the path itself is highlighted with a dotted line. Truncated events might intersect the path, but are not displayed in the map.

Events Table

LICENSE: Malware or Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The Events table lists event information for each data point in the map. You can sort events in ascending or descending order by clicking the column headers. You can highlight a data point in the map by selecting the table row. The map scrolls to display the selected file event if not currently visible. For more information on the fields, see [Understanding the File Events Table](#) on page 1224.

Events

Time	Event ...	Sending IP	Receiving IP	File Name	D...	Action	Prot...	Client	We...	Description
2013-03-18 12...	Transfer	10.10.10.2	10.10.10.3	setup.exe	N...		HTTP	Inter...		Tumblr
2013-03-18 13...	Transfer	10.10.10.4	10.10.10.5	setup.exe	N...		HTTP	Firefox		Face...
2013-03-18 13...	Transfer	10.10.10.5	10.10.10.6	setup.exe	N...		HTTP	Inter...		Mega...
2013-03-18 14...	Transfer	10.10.10.2	10.10.10.7	setup.exe	N...		HTTP	Firefox		The ...

CHAPTER 32

INTRODUCTION TO NETWORK DISCOVERY

The Sourcefire 3D System uses a feature called *network discovery* to monitor traffic on your network and build a comprehensive map of your network assets.

As managed devices passively observe traffic on the network segments you specify, the system compares specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine the number and types of hosts (including network devices) on your network, as well as the operating systems, active applications, and open ports on those hosts.

You can also configure Sourcefire managed devices to monitor user activity on your network, which allows you to identify the source of policy breaches, attacks, or network vulnerabilities.

To supplement the data gathered by the system, you can import records generated by NetFlow-enabled devices, Nmap active scans, the Sourcefire host input feature, and Sourcefire User Agents that reside on a Microsoft Active Directory server and report LDAP authentications. The Sourcefire 3D System integrates these records with the information it collects via direct network traffic observation by managed devices.

The system uses all of this information to help you with forensic analysis, behavioral profiling, access control, and mitigating and responding to the vulnerabilities and exploits to which your organization is susceptible.

For more information, see:

- [Understanding Discovery Data Collection](#) on page 1251
- [Understanding NetFlow](#) on page 1271

- [Creating a Network Discovery Policy](#) on page 1275
- [Obtaining User Data from LDAP Servers](#) on page 1298

Understanding Discovery Data Collection

LICENSE: FireSIGHT

Discovery data includes information on your network's hosts and the operating systems, active applications, and user activity on those hosts.

To begin collecting discovery data, you must first apply an access control policy. The access control policy (see [Using Access Control Policies](#) on page 424) defines the traffic that you permit, and therefore the traffic you can monitor with network discovery. Note that this means if you block certain traffic using access control, the system cannot examine that traffic for host, user, or application activity. For example, if you block access to social networking applications, the system does not provide you with any discovery data on social network applications.

After you apply an access control policy, you must configure and apply a network discovery policy, which specifies the network segments and ports you want to monitor with your managed devices, and the kinds of data you want to collect. When you apply the network discovery policy, the system begins generating discovery data, which you can then view and analyze using the Defense Center web interface.

The system stores network discovery data in the Defense Center database; for information on storage limits, see [Configuring Database Event Limits](#) on page 1978. In addition to the database limits, the total number of detected hosts and users the Defense Center can store depends on your FireSIGHT license.

After you reach the licensed user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database. On the other hand, after you reach the licensed host limit, you can configure the system either to stop adding new hosts to the database, or to replace the hosts that have remained inactive for the longest time.

To supplement the data gathered by the system, you can import records generated by NetFlow-enabled devices, Nmap active scans, the Sourcefire host input feature, and Sourcefire User Agents that reside on a Microsoft Active Directory server and report LDAP authentications. The Sourcefire 3D System integrates these records with the information it collects via direct network traffic observation by managed devices.

For more information, see:

- [Understanding Host Data Collection](#) on page 1252
- [Understanding User Data Collection](#) on page 1252
- [Understanding Application Detection](#) on page 1263

- [Importing Third-Party Discovery Data](#) on page 1269
- [Uses for Discovery Data](#) on page 1270

Understanding Host Data Collection

LICENSE: FireSIGHT

As the system passively monitors the traffic that travels through your network, it system compares specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine the following information about the hosts on your network, including:

- the number and types of hosts (including network devices such as bridges, routers, load balancers, and NAT devices)
- basic network topology data, including the number of hops from the discovery point on the network to the hosts
- the operating systems running on the hosts

If the system cannot identify the operating system of a host, you can use the custom fingerprinting feature to create custom client or server fingerprints. The system uses these fingerprints to identify new hosts. You can map fingerprints to systems in the vulnerability database (VDB) to allow the appropriate vulnerability information to be displayed whenever a host is identified using the custom fingerprint. For more information, see [Using Custom Fingerprinting](#) on page 1649.

You can also add or update host and operating system data through the host input feature. In addition, if you create a NetFlow-enabled discovery rule with host detection enabled, hosts can be added to the network map from NetFlow data.

You can view the hosts detected by the system using the Defense Center web interface:

- For information on viewing and searching for hosts using the event viewer, see [Working with Hosts](#) on page 1401.
- For information on viewing the network map, which is a detailed representation of your network assets and topology, see [Using the Network Map](#) on page 1314.
- For information on viewing host profiles, which are complete views of all the information available for your detected hosts, see [Using Host Profiles](#) on page 1334.

Understanding User Data Collection

LICENSE: FireSIGHT

You can use the Sourcefire 3D System 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. In other words, the system can tell you the “who” behind the “what.” For example, you could determine:

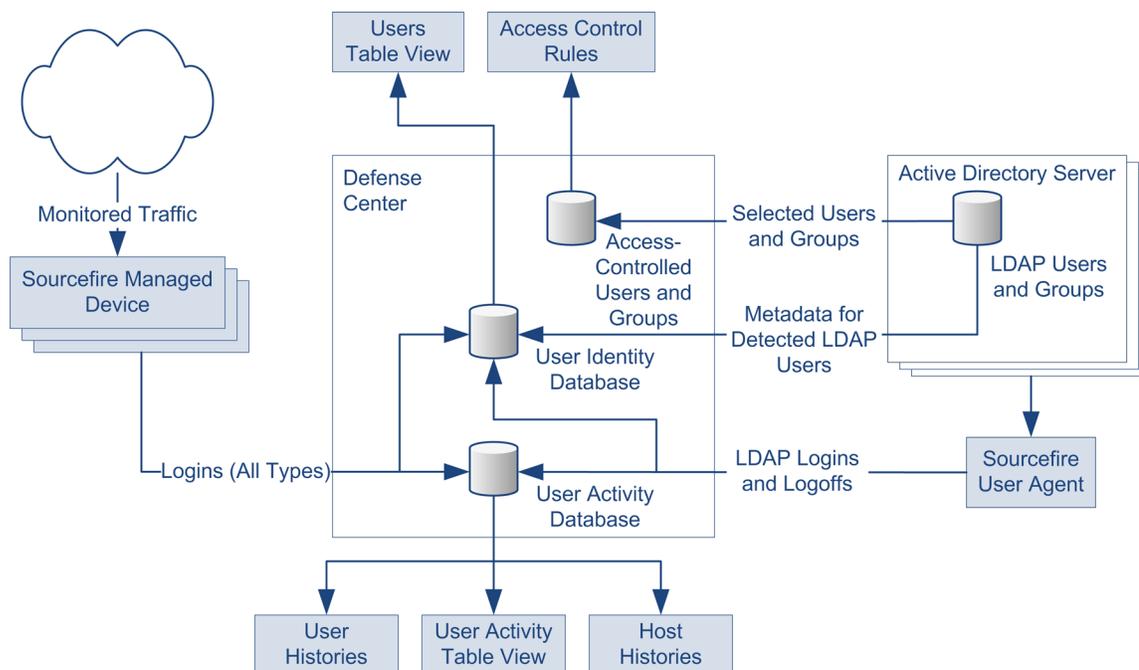
- who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact flag
- who initiated an internal attack or portscan
- who is attempting unauthorized access of a server that has high host criticality
- who is consuming an unreasonable amount of bandwidth
- who has not applied 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 take a targeted approach to mitigate risk, block users or user activity, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

The system downloads the users used in access control policies from the Microsoft Active Directory LDAP server, based on the user awareness settings in the LDAP connection. The Sourcefire User Agent then provides login data for these users and the users are added to the user database. These users are referred to as *access-controlled users*. When you author access control policies that include user conditions, you write those conditions against access-controlled users. For more information, see [Adding User Conditions](#) on page 497.

When the system detects user data from a user login, either from a Sourcefire User Agent, or from an email login over POP3, SMTP, or IMAP, the user from the login is checked against the list of users. If the login user matches an existing user reported by an agent, 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 following diagram illustrates how the Sourcefire 3D System collects and stores user data.



As shown in the diagram, there are three sources for user data, and three places that data is stored. For more information on user data collection, see:

- [Managed Devices](#) on page 1254
- [Sourcefire User Agents](#) on page 1255
- [Defense Center-LDAP Server Connections](#) on page 1258
- [Users Database](#) on page 1258
- [User Activity Database](#) on page 1259
- [Access-Controlled Users Database](#) on page 1260
- [User Data Collection Limitations](#) on page 1261

Managed Devices

LICENSE: FireSIGHT

You use the network discovery policy to configure managed devices to passively detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), and SMTP logins on the

networks you specify. Note that when you enable discovery of users in a network discovery rule, host discovery is automatically enabled.

IMPORTANT! Managed devices interpret only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.

When a device detects a login, it sends the following information to the Defense Center to be logged as user activity:

- the user name identified in the login
- the time of the login
- the IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for SMTP and Oracle logins), or the session originator (for SIP logins)
- the user's email address (for POP3, IMAP, and SMTP logins)
- the name of the device that detected the login

If the user was previously detected, the Defense Center updates that user's login history. Note that the Defense Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Defense Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user, rather, it updates the LDAP user's history.

If the user has never been detected before, the Defense Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data in those login events that the Defense Center can correlate with other login types.

The Defense Center does **not** log user activity or user identities in the following cases:

- if you configured the network discovery policy to ignore that login type, as described in [Restricting User Logging](#) on page 1285
- if a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

Sourcefire User Agents

LICENSE: FireSIGHT

If your organization uses Microsoft Active Directory LDAP servers, Sourcefire recommends that you install Sourcefire User Agents to monitor user activity via your Active Directory servers. If you want to perform user control, you **must** install and use Sourcefire User Agents; the agents associate users with IP addresses, which in turn 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.

To use an agent, you must configure a connection between each Defense Center connected to the agent and the monitored LDAP servers. This 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. For more information on configuring LDAP servers for user discovery, see [Creating LDAP Connections with the Defense Center](#) on page 1298.

Each agent can monitor logins using encrypted traffic, either through regularly scheduled polling or real-time monitoring. Logins are generated by the Active Directory server when a user logs into a computer, whether at the workstation or through a Remote Desktop login.

Agents can also monitor and report user logoffs. Logoffs are generated by the agent itself when it detects a user logged out of a host IP address. Logoffs are also generated when the agent detects that the user logged into a host has changed, before the Active Directory server reports that the user has changed. Combining logoff data with login data develops a more complete view of the users logged into the network.

Polling an Active Directory server allows an agent to retrieve batches of user activity data at the defined polling interval. Real-time monitoring transmits user activity data to the agent as soon as the Active Directory server receives the data.

You can configure the agent to exclude reporting any logins or logoffs associated with a specific user name or IP address. This can be useful, for example, to exclude repeated logins to shared servers, such as file shares and print servers, as well as exclude users logging into machines for troubleshooting purposes.

The agents send records of all detected logins and logoffs that do not contain an excluded user name or IP address to Defense Centers, which log and report them as user activity. The agents detect the Defense Center version and send the login records in the appropriate data format. This supplements any user activity detected directly by managed devices. The logins reported by User Agents associate users with IP addresses, which in turn allows access control rules with user conditions to trigger.

User Agents monitor users as they log into the network or when accounts authenticate against Active Directory credentials for other reasons. Version 2.1 of the Sourcefire User Agent detects interactive user logins to a host, Remote Desktop logins, file-share authentication, and computer account logins, as well as user logoffs and Remote Desktop sessions where the user has logged off.

The type of login detected determines how the agent reports the login and how the login appears in the host profile. An *authoritative user login* for a host causes the current user mapped to the host IP address to change to the user from the new login. Other logins either do not change the current user or only change the current user for the host if the existing user on the host did not have an authoritative user login to the host. In these cases, if the expected user is no

longer logged in, the agent generates a logoff for that user. User logins detected by network discovery only change the current user for the host if the existing user on the host did not have an authoritative user login to the host. Agent-detected logins have the following effect on the network map:

- When the agent detects an interactive login to a host by a user or a Remote Desktop login, the agent reports an authoritative user login for the host and changes the current user for the host to the new user.
- If the agent detects a login for file-share authentication, the agent reports a user login for the host, but does not change the current user on the host.
- If the agent detects a computer account login to a host, the agent generates a NetBIOS Name Change discovery event and the host profile reflects any change to the NetBIOS name.
- If the agent detects a login from an excluded user name, the agent does not report a login to the Defense Center.

When a login or other authentication occurs, the agent sends the following information to the Defense Center:

- the user's LDAP user name
- the time of the login or other authentication
- the IP address of the user's host, and the link-local address if the agent reports an IPv6 address for a computer account login

The Defense Center records login and logoff information as user activity. When a Sourcefire User Agent reports user data from a user login or logoff, the reported user is checked against the list of users. If the reported user matches an existing user reported by an agent, the reported data is assigned to the user. Reported users that do not match existing users cause a new user to be created.

Even though the user activity associated with an excluded user name is not reported, related user activity may still be reported. If the agent detects a user login to a machine, then the agent detects a second user login, and you have excluded the user name associated with the second user login from reporting, the agent reports a logoff for the original user. However, no login for the second user is reported. As a result, no user is mapped to the IP address, even though the excluded user is logged into the host.

Note the following limitations on user names detected by the agent:

- User names ending with a dollar sign character (\$) reported to a Version 5.0.2+ Defense Center update the network map, but do not appear as user logins. Agents do not report user names ending with a dollar sign character (\$) to any other versions of Defense Centers.
- Defense Center display of user names containing Unicode characters may have limitations.

The total number of detected users the Defense Center can store depends on your RNA or FireSIGHT license. After you reach the licensed user limit, in most cases the system stops adding new users to the database. To add new users, you

must either manually delete old or inactive users from the database, or purge all users from the database.

IMPORTANT! Version 1.0 (legacy) Sourcefire Agents installed on Active Directory LDAP servers can continue to send user login data from the Active Directory server to a single Defense Center. Deployment requirements and detection capabilities of legacy agents are unchanged. You must install them on the Active Directory server to connect to exactly one Defense Center. Note, however, that the User Agent Status Monitor health module does not support legacy agents and should not be enabled on Defense Centers with legacy agents connected. You should plan to upgrade your deployment to use Version 2.1 of the Sourcefire User Agent as soon as possible, in preparation for future releases when support for legacy agents will be phased out.

Defense Center-LDAP Server Connections

LICENSE: FireSIGHT

The Defense Center-LDAP server connection allows you to retrieve metadata for certain detected users. You can retrieve metadata for LDAP users, whether their logins were detected by managed devices or by a User Agent; you can also retrieve metadata for POP3 and IMAP users if those users have the same email address as an LDAP user.

If your organization uses Microsoft Active Directory servers, the connection also allows you to specify the LDAP users and groups you want to use in access control rules. If you want to perform user control, you **must** configure a connection between the Defense Center and an Active Directory server. If your organization does not use Active Directory, you can still detect user logins using managed devices, and you can still obtain metadata for some of those users from an Oracle or OpenLDAP server. However, you cannot perform user control based on those users or their activity.

From the LDAP server, the Defense Center obtains the following information and metadata about each user:

- LDAP user name
- first and last names
- email address
- department
- telephone number

Users Database

LICENSE: FireSIGHT

The users database contains a record for each user detected by either managed devices or User Agents. The total number of detected users the Defense Center can store depends on your RNA or FireSIGHT license. After you reach the

licensed limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative user logins. If you have reached the limit and the system detects an authoritative user login for a previously undetected user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new user.

You can view the contents of the users database with the Defense Center web interface. For information on viewing, search for, and deleting detected users, see [Working with Users](#) on page 1445.

User Activity Database

LICENSE: FireSIGHT

The user activity database contains records of user activity on your network, either from a connection to an Active Directory LDAP server that is also monitored by a Sourcefire User Agent, or through network discovery. 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 FireSIGHT licensed limit

You can view the user activity detected by the system using the Defense Center web interface. For information on viewing, searching for, and deleting user activity, see [Working with User Activity](#) on page 1453. If you plan to use Version 2.1 of the Sourcefire 3D System User Agent to send LDAP login data to your Defense Centers, you must configure a connection for each agent on each Defense Center where you want the agent to connect. That connection allows the agent to establish a secure connection with the Defense Center, over which it can send login data. If the agent is configured to exclude specific user names, login data for those user names are not reported to the Defense Center.

In addition, if you are planning to implement user access control, you must set up a connection to each Microsoft Active Directory server where you plan to collect data, with user awareness parameters configured.

Whenever possible the Sourcefire 3D System correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event.

The system also uses user activity to generate *host histories*, which track the hosts that each user has logged into, and *user histories*, which track the users that have logged into each individual host. The system provides a graphical representation of the last twenty-four hours of each user's activity and the last twenty-four hours of the logins to each host. For more information, see [Understanding User Details and Host History](#) on page 1449 and [Working with](#)

[User History in the Host Profile](#) on page 1357.

Access-Controlled Users Database

LICENSE: Control

The access-controlled users database contains the users and groups that you can use in access control rules, so that you can perform user control with the Sourcefire 3D System. These users can be one of two types:

- An *access-controlled user* is a user that you can add to access control rules to perform user control. You specify the groups that access-controlled users must belong to when you configure the Defense Center-LDAP server connection.
- A *non-access-controlled user* is any other detected user.

You specify the groups that access-controlled users must belong to when you configure the Defense Center-LDAP server connection, as described in [Creating LDAP Connections with the Defense Center](#) on page 1298.

If you plan to use Version 2.1 of the Sourcefire 3D System User Agent to send LDAP login and logoff data to your Version 5.x Defense Centers, you must configure a connection for each agent on each Defense Center where you want the agent to connect. That connection allows the agent to establish a secure connection with the Defense Center, over which it can send the user activity data.

If the agent is configured to exclude specific user names, user activity data for those user names are not reported to the Defense Center. These excluded user names remain in the database, but are not associated with IP addresses.

In addition, if you are planning to implement user access control, you must set up a connection to each Microsoft Active Directory server where you plan to collect data, with user awareness parameters configured.

The maximum number of users you can use in access control depends on your FireSIGHT license. When configuring the Defense Center-LDAP server connection, make sure the total number of users you include is less than your FireSIGHT user license. See [Understanding FireSIGHT Host and User License Limits](#) on page 2045 for more information.

User Data Collection Limitations

LICENSE: FireSIGHT

The following table describes the limitations of user data collection.

User Awareness Limitations

LIMITATION	DESCRIPTION
user control	To perform user control, your organization must use Microsoft Active Directory LDAP servers. The system obtains the users and groups you can use in access control rules from Active Directory, and also ties users to IP addresses with the logins and logoffs reported by Sourcefire User Agents installed on Active Directory servers.
non-Kerberos logins for LDAP connections	<p>Managed devices interpret only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications if they use other protocols, such as SSL or TLS.</p> <p>On the other hand, Sourcefire User Agents use the security logs on Active Directory servers to collect user login data and have no such limitations.</p>
login detection	<p>Version 2.1 of the Sourcefire User Agent reports user logins to hosts with IPv6 addresses to Defense Centers running Version 5.2+.</p> <p>The agent reports non-authoritative user logins and NetBIOS logins to Defense Centers running Version 5.0.1+.</p> <p>The agent reports authoritative logins from actual user names to Defense Centers running Version 4.10.x+.</p> <p>If you want to detect logins to an Active Directory server, you must configure the Active Directory server connection with the server IP address. See the <i>Sourcefire 3D System User Agent Configuration Guide</i> for more information.</p> <p>If multiple users are logged into a host using remote sessions, the agent may not detect logins from that host properly. See the <i>Sourcefire 3D System User Agent Configuration Guide</i> for more information on how to prevent this.</p>

User Awareness Limitations (Continued)

LIMITATION	DESCRIPTION
logoff detection	<p>The agent reports detected logoffs to Version 5.2+ Defense Centers.</p> <p>Logoffs may not be immediately detected. The timestamp associated with a logoff reflects when the agent detected the user was no longer mapped to the host IP address, which may not correspond to the actual time the user logged off of the host.</p> <p>Logoffs are generated by the agent itself when it detects a user logged out of a host IP address. Logoffs are also generated when the agent detects that the user logged into a host has changed, before the Active Directory server reports that the user has changed.</p>
real-time data retrieval	<p>The Active Directory server must be running Windows Server 2008 or Windows Server 2012.</p>
multiple logins to the same host by different users	<p>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 only non-authoritative logins have been logged into the host, the last non-authoritative login is considered the current user. If multiple users are logged in through remote sessions, the last user reported by the Active Directory server is the user reported to the Defense Center.</p>
multiple logins to the same host by the same user	<p>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.</p> <p>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.</p>
Unicode characters	<p>The user interface may not correctly display user names with Unicode characters.</p> <p>The agent does not report user names with Unicode characters to Version 4.10.x Defense Centers.</p>

User Awareness Limitations (Continued)

LIMITATION	DESCRIPTION
LDAP user accounts in the users database	If you remove or disable an LDAP user on your user awareness or RUA LDAP servers, or exclude the user name from being reported to the Defense Center, the Defense Center does not remove that user from the users database, and that user continues to count against your licensed limit for user listed in the database. You must manually purge the user from the database. For Version 5.x, note that the user license limit is applied in parallel for access-controlled users; the user count for access-controlled users depends on the number of users retrieved by your LDAP configuration.
AOL Instant Messenger (AIM) login detection	Managed devices can detect AIM logins using the OSCAR protocol only. While most AIM clients use OSCAR, some use TOC2.

Understanding Application Detection

LICENSE: FireSIGHT

When the Sourcefire 3D System analyzes IP traffic, it attempts to identify the commonly used applications on your network. Application awareness is crucial to performing application-based access control.

There are three types of applications that the system detects:

- *application protocols* such as HTTP and SSH, which represent communications between hosts
- *clients* such as web browsers and email clients, which represent software running on the host
- *web applications* such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic

The system identifies applications in your network traffic either using ASCII or hexadecimal patterns in the packet headers, or the port that the traffic uses. Some application detectors use both port and pattern detection to increase the likelihood of correctly identifying traffic for a particular application. In addition, Secure Socket Layers (SSL) protocol detectors use information from the secured

session to identify the application from the session. There are two sources of application detectors in the Sourcefire 3D System:

- *Sourcefire-provided detectors*, which detect web applications, clients, and application protocols

The availability of Sourcefire-provided detectors for applications (and operating systems, see [Understanding Host Data Collection](#) on page 1252) depend on the version of the Sourcefire 3D System and the version of the VDB you have installed. Release notes and advisories contain information on new and updated detectors. You can also import individual detectors authored by Sourcefire Professional Services. For a complete list of detected applications, see the [Sourcefire Support Site](#).

- *user-defined application protocol detectors*, which you can create to enhance the system’s application protocol detection capabilities

You can also detect application protocols through *implied application protocol detection*, which implies the existence of an application protocol based on the detection of a client.

The system characterizes each application that it detects using the criteria described in the following table. The system uses these characteristics to create application filters, or groups of applications. You can use these filters and filters that you create to perform access control, as well as to constrain searches, reports, and dashboard widgets. For more information, see [Working with Application Filters](#) on page 176.

Application Characteristics

CRITERION	DESCRIPTION	EXAMPLE
Risk	How likely the application is to be used for purposes that might be against your organization’s security policy. An application’s risk can range from Very Low to Very High .	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is used within the context of your organization’s business operations, as opposed to recreationally. An application’s business relevance can range from Very Low to Very High .	Gaming applications tend to have a very low business relevance.

Application Characteristics (Continued)

CRITERION	DESCRIPTION	EXAMPLE
Type	The type of application: <ul style="list-style-type: none"> • Application Protocols represent communications between hosts. • Clients represent software running on a host. • Web Applications represent the content or requested URL for HTTP traffic. 	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged high bandwidth and displays ads .

To supplement the application data gathered by the system, you can use records generated by NetFlow-enabled devices, Nmap active scans, and the Sourcefire host input feature.

For more information, see:

- [Understanding the Application Protocol Detection Process](#) on page 1265
- [Implied Application Protocol Detection from Client Detection](#) on page 1267
- [Special Considerations for Application Protocol Detection: Squid](#) on page 1268
- [Special Considerations: SSL Application Detection](#) on page 1269
- [Working with Application Detectors](#) on page 1663
- [Importing Third-Party Discovery Data](#) on page 1269
- [Understanding NetFlow](#) on page 1271

Understanding the Application Protocol Detection Process

LICENSE: FireSIGHT

When the system detects application traffic, it first determines whether the application protocol is running on a port identified by a detector that uses that specific port as its only detection criterion. If the application protocol is running on

one of those ports, the system positively identifies the application protocol using the well-known port detector.

IMPORTANT! Because you can create and activate user-defined port-based application protocol detectors on ports used by Sourcefire-provided detectors, it is possible to override Sourcefire's detection capabilities. For example, if your user-defined detector identifies all application protocol traffic on port 22 as the `myapplication` application protocol, SSH traffic on port 22 will be misidentified as `myapplication` traffic.

If the application protocol is not running on one of those ports, the system employs a more robust method to identify it based on port and pattern matches. If two detectors both positively identify the traffic, the detector that employs the longer pattern match has precedence. Similarly, detectors with multiple pattern matches have precedence over single pattern matches.

Note that the system identifies only those application protocols running on hosts in your monitored networks, as defined in the network discovery policy. For example, if an internal host accesses an FTP server on a remote site that you are not monitoring, the system does not identify the application protocol as FTP. On the other hand, if a remote or internal host accesses an FTP server on a host you are monitoring, the system can positively identify the application protocol.

An exception occurs if the system can identify the client used in connections between a monitored host accessing a non-monitored server. In that case, the system positively identifies the appropriate application protocol that corresponds with the client in the connection, but does not add the application protocol to the network map. For more information, see [Implied Application Protocol Detection from Client Detection](#) on page 1267.

The following table outlines how the Sourcefire 3D System identifies detected application protocols in the Defense Center web interface: the network map, host profiles, event views, and so on.

Sourcefire 3D System Identification of Application Protocols

APPLICATION	DESCRIPTION
the application protocol name	The Defense Center identifies an application protocol with its name if the application protocol was: <ul style="list-style-type: none"> positively identified by the system identified using NetFlow data and there is a port-application protocol correlation in <code>/etc/sf/services</code> manually identified using the host input feature identified by Nmap or another active source
pending	The Defense Center identifies an application protocol as pending if the system can neither positively nor negatively identify the application. <p>Most often, the system needs to collect and analyze more connection data (from which applications are identified) before it can identify a pending application.</p> <p>In the Application Details and Servers tables and in the host profile, the pending status appears only for application protocols where specific application protocol traffic was detected (rather than implied by detected client or web application traffic).</p>
unknown	The Defense Center identifies an application protocol as unknown if the application: <ul style="list-style-type: none"> does not match any of the system's detectors the application protocol was identified using NetFlow data, but there is no port-application protocol correlation in <code>/etc/sf/services</code>
blank	All available detected data has been examined and no application protocol was identified. In the Application Details and Servers tables and in the host profile, the application protocol is left blank for non-HTTP generic client traffic with no detected application protocol.

Implied Application Protocol Detection from Client Detection

LICENSE: FireSIGHT

If the system can identify the client used in a connection between a monitored host accessing a non-monitored server, the Defense Center infers that the connection is using the application protocol that corresponds with the client. (Because the system tracks applications only on monitored networks, connection logs usually do not include application protocol information for connections where a monitored host is accessing a non-monitored server.)

There are several consequences of the implied detection of an application protocol from the detection of a client:

- Because the system does not generate a New TCP Port or New UDP Port event for these servers, the server does not appear in the Servers table. In addition, you cannot trigger either discovery event alerts or correlation rules using the detection of these application protocol as a criterion.
- Because the application protocol is not associated with a host, you cannot view its details in host profiles, set its server identity, or use its information in host profile qualifications for traffic profiles or correlation rules. In addition, the system does not associate vulnerabilities with hosts based on this type of detection.

You can, however, trigger correlation events on the application protocol information in a connection. You can also use the application protocol information in connection logs to create connection trackers and traffic profiles.

Host Limits and Discovery Event Logging

LICENSE: FireSIGHT

When the system detects a client, server, or web application it generates a discovery event unless the associated host has already reached its maximum number of clients, servers, or web applications.

Host profiles display up to 16 clients, 100 servers, and 100 web applications per host. See [Working with Servers in the Host Profile](#) on page 1347 and [Viewing Applications in the Host Profile](#) on page 1354 for more information.

Note that actions dependent on the detection of clients, servers, or web applications are unaffected by this limit. For example, access control rules configured to trigger on a server will still log connection events.

Special Considerations for Application Protocol Detection: Squid

LICENSE: FireSIGHT

The system positively identifies Squid server traffic when either:

- the system detects a connection from a host on your monitored network to a Squid server where proxy authentication is enabled, or
- the system detects a connection from a Squid proxy server on your monitored network to a target system (that is, the destination server where the client is requesting information or another resource)

However, the system cannot identify Squid service traffic if:

- a host on your monitored network connects to a Squid server where proxy authentication is disabled, or
- the Squid proxy server is configured to strip Via: header fields from its HTTP responses

Special Considerations: SSL Application Detection

LICENSE: FireSIGHT

The Sourcefire 3D System provides detectors that can use session information from a Secure Socket Layers (SSL) session to identify the application protocol, client application, or web application in the session.

When the system detects an encrypted connection, it marks that connection as either a generic HTTPS connection or as a more specific secure protocol, such as SMTPS, when applicable. When the system detects an SSL session, it adds **SSL c1ient** to the **Client** field in connection events for the session. If it identifies a web application for the session, the system generates discovery events for the traffic.

For SSL application traffic, managed devices running Version 5.2 or later can also detect the common name from the server certificate and match that against a client or web application from an SSL host pattern. When the system identifies a specific client, it replaces **SSL c1ient** with the name of the client. Note that managed devices running versions earlier than Version 5.2 cannot detect applications in SSL traffic, even if managed by a Version 5.2 Defense Center.

Because the SSL application traffic is encrypted, the system can use only information in the certificate for identification, not application data within the encrypted stream. For this reason, SSL host patterns can sometimes only identify the company that authored the application, so SSL applications produced by the same company may have the same identification.

In some instances, such as when an HTTPS session is launched from within an HTTP session, managed devices running Version 5.2 or later detect the server name from the client certificate in a client-side packet.

To enable SSL application identification, you must create access control rules that monitor responder traffic. Those rules must have either an application condition for the SSL application or URL conditions using the URL from the SSL certificate. For network discovery, the responder IP address does not have to be in the networks to monitor in the network discovery policy; the access control policy configuration determines whether the traffic is identified. You can filter by the **SSL protocol** tag, in the application detectors list or when adding application conditions in access control rules, to identify detectors for SSL applications.

Importing Third-Party Discovery Data

LICENSE: FireSIGHT

You can use Nmap active scans to add information about operating systems, applications, and vulnerabilities, supplementing the data gathered by the system. For more information on Nmap scanning and scan results, see [Understanding Nmap Scans](#) on page 1693.

You can also use the host input feature to supplement the information that the system gathers from monitoring network traffic, either by configuring a third-party application to interact with the Sourcefire 3D System via an API, or by manually

adding data. You can create product, vulnerability, and fix mappings to map third-party data to Sourcefire definitions, enabling impact correlation for operating systems and servers. For more information on the host input feature and mapping third-party data, see the *Sourcefire 3D System Host Input API Guide* and [Importing Host Input Data](#) on page 1680.

The system reconciles the collected data about operating system and server identities and determines each identity based on fingerprint source priority values, identity conflict resolution settings, and time of collection.

You can also configure your network map to use data from NetFlow-enabled devices to enhance your network map and event tables. For more information, see [Understanding NetFlow](#) on page 1271.

Uses for Discovery Data

LICENSE: FireSIGHT

Logging discovery data allows you to take advantage of many features in the Sourcefire 3D System, including:

- viewing the network map, which is a detailed representation of your network assets and topology that you can view by grouping hosts and network devices, host attributes, application protocols, or vulnerabilities; see [Using the Network Map](#) on page 1314
- viewing host profiles, which are complete views of all the information available for your detected hosts; see [Using Host Profiles](#) on page 1334
- viewing dashboards, which (among other capabilities) can provide you with an at-a-glance view of your network assets and user activity; see [Using Dashboards](#) on page 66
- viewing detailed information on the discovery events and user activity logged by the system; see [Working with Discovery Events](#) on page 1377
- creating reports based on discovery data; see [Working with Reports](#) on page 1724
- performing application and user control, that is, writing access control rules using application and user conditions; see [Understanding and Writing Access Control Rules](#) on page 472
- associating hosts and any servers or clients they are running with the exploits to which they are susceptible, which allows you to identify and mitigate vulnerabilities, evaluate the impact that intrusion events have on your network, and tune intrusion rule states so that they provide maximum protection for your network assets; see [Working with Vulnerabilities in the Host Profile](#) on page 1363, [Using Impact Flags to Evaluate Events](#) on page 641, and [Managing FireSIGHT Rule State Recommendations](#) on page 764
- alerting you via email, SNMP trap, or syslog when the system generates either an intrusion event with a specific impact flag, or a specific type of discovery event; see [Configuring External Alerting](#) on page 524

- monitor your organization's compliance with a white list of allowed operating systems, clients, application protocols, and protocols; see [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530
- creating correlation policies with rules that trigger and generate correlation events when the system generates discovery events or detects user activity; see [Configuring Correlation Policies and Rules](#) on page 1459
- if you log NetFlow connections, using that connection data; see [Uses for Connection Data in the Sourcefire 3D System](#) on page 554

Understanding NetFlow

LICENSE: FireSIGHT

NetFlow is an embedded instrumentation within Cisco IOS Software that characterizes network operation. Standardized through the RFC process, NetFlow is available not only on Cisco networking devices, but can also be embedded in Juniper, FreeBSD, and OpenBSD devices.

NetFlow-enabled devices are widely used to capture and export data about the traffic that passes through those devices. NetFlow-enabled devices have a database called the NetFlow cache that stores records of the flows that pass through the devices. A flow, called a *connection* in the Sourcefire 3D System, is a sequence of packets that represents a session between a source and destination host, using specific ports, protocol, and application protocol.

For the networks you specify, Sourcefire managed devices detect the records exported by NetFlow-enabled devices, generate connection events based on the data in those records, and finally send those events to the Defense Center to be logged in the database. You can also configure the system to add host and application protocol information to the database, based on the information in NetFlow connections.

You can use this discovery and connection data to supplement the data gathered directly by your managed devices. This is especially useful if you have NetFlow-enabled devices deployed on networks that your managed devices cannot monitor.

You configure NetFlow data collection, including connection logging, using rules in the network discovery policy. Contrast this with connection logging for connections detected by Sourcefire managed devices, which you configure per access control rule, as described in [Logging Connection, File, and Malware Information](#) on page 515. Because NetFlow data collection is linked to networks rather than access control rules, you do not have as much granular control over which connections you want to log. Also, the system automatically saves all NetFlow-based connection events to the Defense Center connection event database; you cannot send them to the system log or an SNMP trap server.

For more information, see:

- [Differences Between NetFlow and FireSIGHT Data](#) on page 1272
- [Preparing to Analyze NetFlow Data](#) on page 1274
- [Uses for Discovery Data](#) on page 1270
- [Uses for Connection Data in the Sourcefire 3D System](#) on page 554

Differences Between NetFlow and FireSIGHT Data

LICENSE: FireSIGHT

With one exception (TCP flags), the information available in NetFlow records is more limited than the information generated by monitoring network traffic using managed devices. Because the system cannot directly analyze the traffic represented by NetFlow data, when the system processes NetFlow records it uses various methods to convert that data into connection logs as well as into host and application protocol records.

There are several differences between converted NetFlow data and the discovery and connection data gathered directly by your managed devices. You should keep the differences in mind when performing analysis that requires:

- statistics on the number of detected connections
- operating system and other host-related information (including vulnerabilities)
- application data, including client information, web application information, and vendor and version server information
- knowing which host in a connection is the initiator and which is the responder

TIP! For each field in a connection event, the [Connection Data Based on Logging and Detection Methods table](#) on page 552 indicates the available data depending on whether the connection was detected directly by Sourcefire managed devices, or if the connection event is based on NetFlow data.

Number of Connection Events Generated Per Monitored Session

For connections detected directly by managed devices, depending on the access control rule action, you can log a bidirectional connection event at the beginning or end of a connection, or both.

However, because NetFlow-enabled devices export unidirectional connection data, the system always generates at least two connection events for each connection detected by NetFlow-enabled devices, depending on how you configured the devices. This also means that a summary's connection count is incremented by two for every connection based on NetFlow data, providing an inflated count of the number of connections that are actually occurring on your network.

Note that if you configure your NetFlow-enabled devices to output records only when the connection ends, the system generates two connection events for that session. On the other hand, if you configure your NetFlow-enabled devices to output records at a fixed interval even if a connection is still ongoing, the system generates a connection event for each record exported by the device. For example, if you configure your NetFlow-enabled devices to output records for long-running connections every five minutes, and a particular connection lasts twelve minutes, the system generates six connection events for that session:

- one pair of events for the first five minutes
- one pair for the second five minutes
- a final pair when the connection is terminated

For this reason, Sourcefire **strongly** recommends that you configure your NetFlow-enabled devices to output records only when monitored sessions close.

Host and Operating System Data

Although you can configure the network discovery policy to add hosts to the network map based on NetFlow records, the host profile does not include any operating system or NetBIOS data for the hosts involved in the connection, nor can the system identify if the hosts are network devices (bridges, routers, NAT devices, or load balancers). You can, however, manually set a host's operating system identity using the host input feature.

Application Data

For connections detected directly by managed devices, the system can identify application protocols, clients, and web applications by examining the packets in the connection.

When the system processes NetFlow records, the system uses a port correlation in `/etc/sf/services` to extrapolate application protocol identity. However, there is no vendor or version information for those application protocols, nor do connection logs contain information on client or web applications used in the session. You can, however, manually provide this information using the host input feature.

Note that a simple port correlation means that application protocols running on non-standard ports may be unidentified or misidentified. Additionally, if no correlation exists, the system marks the application protocol as **unknown** in connection logs.

Vulnerability Mappings

The Sourcefire 3D System cannot determine which vulnerabilities might affect hosts added to the network map based on NetFlow records, unless you use the host input feature to manually set either a host's operating system identity or an application protocol identity. Note that because there is no client information in NetFlow connections, you cannot associate client vulnerabilities with NetFlow hosts.

Initiator and Responder Information in Connections

For connections detected directly by managed devices, the system can identify which host is the initiator, or source, and which is the responder, or destination. However, NetFlow data does not contain initiator or responder information.

When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known:

- If both or neither port being used is a well-known port, the system considers the host using the lower-number port to be the responder.
- If only one of the hosts is using a well-known port, the system considers that host to be the responder.

For this purpose, a well-known port is any port that is either numbered from 1 to 1023, or that contains application protocol information in `/etc/sf/services` on the managed device.

Preparing to Analyze NetFlow Data

LICENSE: FireSIGHT

Before you configure the Sourcefire 3D System to analyze NetFlow data, you must enable the NetFlow feature on the routers or other NetFlow-enabled devices you plan to use, and configure the devices to export NetFlow version 5 data to a destination network where the sensing interface of a managed device is connected.

Note that the system can parse both NetFlow version 5 and NetFlow version 9 records. Your NetFlow-enabled devices **must** use one of those versions if you want to use them with your Sourcefire 3D System deployment. In addition, the system requires that specific fields be in the templates and records that your NetFlow-enabled devices broadcast. If your NetFlow-enabled devices are using version 9, which you can customize, you **must** make sure that the templates and records that the devices broadcast contain the following fields, in any order:

- IN_BYTES (1)
- IN_PKTS (2)
- PROTOCOL (4)
- TCP_FLAGS (6)
- L4_SRC_PORT (7)
- IPV4_SRC_ADDR (8)
- L4_DST_PORT (11)
- IPV4_DST_ADDR (12)
- LAST_SWITCHED (21)
- FIRST_SWITCHED (22)

- IPV6_SRC_ADDR (27)
- IPV6_DST_ADDR (28)

Because the Sourcefire 3D System uses managed devices to analyze NetFlow data, your deployment must include at least one managed device that can monitor your NetFlow-enabled devices. At least one sensing interface on that managed device must be connected to a network where it can collect the data that your NetFlow-enabled devices export. Because the sensing interfaces on managed devices do not usually have IP addresses, the system does not support the direct collection of NetFlow records.

In addition, Sourcefire **strongly** recommends that you configure your NetFlow-enabled devices to output records only when monitored sessions close. If you configure your NetFlow-enabled devices to output records at fixed intervals, analyzing the connection data derived from the NetFlow records may be more complicated; see [Number of Connection Events Generated Per Monitored Session](#) on page 1272.

Finally, note that the Sampled NetFlow feature available on some NetFlow-enabled devices collects NetFlow statistics on only a subset of packets that pass through the devices. Although enabling this feature can improve CPU utilization on the NetFlow-enabled device, it may affect the data you are collecting for analysis by the system.

Creating a Network Discovery Policy

LICENSE: FireSIGHT

The network discovery policy on the Defense Center controls how the system collects data on your organization's network assets and which network segments and ports are monitored.

Discovery rules within the policy specify what networks and ports the Sourcefire 3D System monitors to generate discovery data based on network data in traffic, and what zones the policy is applied to. Within a rule, you can configure whether hosts, applications, and users are discovered. You can create rules to exclude networks and zones from discovery. When you create a rule for discovery from a NetFlow device, you can choose to just log connections.

The network discovery policy has a single default rule in place, configured to discover applications in any IPv4 traffic on the 0.0.0.0/0 network. Note that you must have applied an access control policy to the targeted device before you can apply a network discovery policy. The rule does not exclude any networks, zones, or ports, host and user discovery is not configured, and a NetFlow device is not configured. Note that the policy is applied to any managed devices by default when they are registered to the Defense Center. To begin collecting host or data, you must add or modify discovery rules and reapply the policy to a device.

Remember that the access control policy (see [Using Access Control Policies](#) on page 424) defines the traffic that you permit, and therefore the traffic you can monitor with network discovery. Note that this means if you block certain traffic

using access control, the system cannot examine that traffic for host, user, or application activity. For example, if you block access to social networking applications in the access control policy, the system will not provide you with any discovery data on those applications.

If you want to adjust the scope of network discovery, you can create additional discovery rules and modify or remove the default rule. You can configure discovery of data from NetFlow devices and can restrict the protocols for traffic where user data is discovered on your network.

If you want to use the Sourcefire 3D System to perform intrusion detection and prevention but do not need to take advantage of discovery data, you can optimize performance by disabling new discovery. First, make sure that your applied access control policies do not contain rules with user, application, or URL conditions. Then, remove all rules from your network discovery policy and apply it to your managed devices. For more information on configuring access control rules, see [Understanding and Writing Access Control Rules](#) on page 472.

If you enable user discovery in your discovery rules, you can detect users through user login activity in traffic over a set of application protocols. You can disable discovery in particular protocols across all rules if needed. Disabling some protocols can help avoid reaching the user limit associated with your FireSIGHT license, reserving available user count for users from the other protocols.

Advanced network discovery settings allow you to manage what data is logged, how discovery data is stored, what vulnerability mappings are used for impact assessment, and what happens when sources offer conflicting discovery data. You can also add NetFlow devices and sources for host input.

For more information, see:

- [Working with Discovery Rules](#) on page 1276
- [Restricting User Logging](#) on page 1285
- [Configuring Advanced Network Discovery Options](#) on page 1287
- [Applying the Network Discovery Policy](#) on page 1297

Working with Discovery Rules

LICENSE: FireSIGHT

Discovery rules allow you to tailor the information discovered for your network map to include only the specific data you want. Rules in your network discovery policy are evaluated sequentially. Note that while you can create rules with overlapping monitoring criteria, doing so may affect your system performance.

When you exclude a host or a network from monitoring, the host or network does not appear in the network map and no events are reported for it. Sourcefire recommends that you exclude load balancers (or specific ports on load balancers) and NAT devices from monitoring. These devices may create excessive and misleading events, filling the database and overloading the Defense Center. For example, a monitored NAT device might exhibit multiple updates of its operating

system in a short period of time. If you know the IP addresses of your load balancers and NAT devices, you can exclude them from monitoring.

TIP! The system can identify many load balancers and NAT devices by examining your network traffic. To determine which hosts on your network are load balancers and NAT devices, apply your network discovery policy, wait for the system to populate the network map, then perform a search of hosts constraining on host type. For more information, see [Searching for Hosts](#) on page 1408.

In addition, if you need to create a custom server fingerprint, you should temporarily exclude from monitoring the IP address that you are using to communicate with the host you are fingerprinting. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address. After you create the fingerprint, you can configure your policy to monitor that IP address again. For more information, see [Fingerprinting Servers](#) on page 1655.

Sourcefire also recommends that you **not** monitor the same network segment with NetFlow-enabled devices and Sourcefire managed devices. Although ideally you should configure your network discovery policy with non-overlapping rules, the system does drop duplicate connection logs generated by managed devices. Note that you **cannot** drop duplicate connection logs for connections detected by both a managed device and a NetFlow-enabled device.

For more information, see the following sections:

- [Understanding Device Selection](#) on page 1277
- [Understanding Actions and Discovered Assets](#) on page 1278
- [Understanding Monitored Networks](#) on page 1279
- [Understanding Zones in Network Discovery Policies](#) on page 1279
- [Understanding Port Exclusions](#) on page 1280
- [Adding a Discovery Rule](#) on page 1280
- [Creating Network Objects](#) on page 1284
- [Creating Port Objects](#) on page 1285

Understanding Device Selection

LICENSE: FireSIGHT

If you select a NetFlow device in a discovery rule, the rule is limited to discovery of NetFlow data for the specified networks. Select the NetFlow device before you configure other aspects of rule behavior, as the available rule actions change when you select a NetFlow device. In addition, you cannot configure port exclusions for NetFlow traffic.

Before you can select a NetFlow device in a network discovery rule, you must configure a connection to the NetFlow device in the network discovery advanced settings. For more information, see [Adding NetFlow-Enabled Devices](#) on

page 1292.

Understanding Actions and Discovered Assets

LICENSE: FireSIGHT

When you configure a discovery rule, you must select an action for the rule. The action determines what assets are discovered or excluded when the system processes the rule. However, note that the affect of a rule action depends on whether you are using the rule to discover data from a managed device or from a NetFlow-enabled device.

Note that if you create a network discovery policy without any rules that discover hosts or users, applying the policy disables new discovery for the appliance. To optimize performance when using managed devices only for intrusion prevention, remove all discovery rules from your policy and apply it to the active devices.

The following table describes what assets are discovered by rules with the specified action settings in those two scenarios.

Discovery Rule Actions

ACTION	MANAGED DEVICE	NETFLOW
Exclude	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.	
Discover: Hosts	Adds hosts to the network map based on discovery events. (Optional, unless user discovery is enabled, then required.)	Adds hosts to the network map based on NetFlow records. (Required)
Discover: Applications	Adds applications to the network map based on application detectors. Note that you cannot discover hosts or users in a rule without also discovering applications. (Required)	Adds application protocols to the network map based on NetFlow records and the port-application protocol correlation in <code>/etc/sf/services</code> . (Optional)

Discovery Rule Actions (Continued)

ACTION	MANAGED DEVICE	NETFLOW
Discover: Users	Adds users to the users table and logs user activity based on activity detected in traffic matching the user protocols configured in the network discovery policy. (Optional)	n/a
Log NetFlow Connections	n/a	Logs NetFlow connections only. Does not discover hosts or applications.

Understanding Monitored Networks

LICENSE: FireSIGHT

A discovery rule causes discovery of monitored assets only in traffic to and from hosts in the specified networks. For a discovery rule, discovery occurs for connections that have at least one IP address within the networks specified, with events generated only for IP addresses within the networks to monitor. The default discovery rule discovers applications only on the 0.0.0.0/0 and ::/0 networks.

For rules with a specified NetFlow device and the **Log Network Connections** option enabled, connections to and from IP addresses in the specified networks are also logged. Note that network discovery rules provide the only way to log NetFlow network connections.

You can also use network object or object groups to specify the networks to monitor. If you modify a network object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

Understanding Zones in Network Discovery Policies

LICENSE: FireSIGHT

For performance reasons, you should configure each discovery rule so that the zones in the rule include the sensing interfaces on your managed devices that are physically connected to the networks-to-monitor in the rule.

Unfortunately, you may not always be kept informed of network configuration changes. A network administrator may modify a network configuration through routing or host changes without informing you, which may make it challenging to stay on top of proper network discovery policy configurations. If you do not know how the sensing interfaces on your managed devices are physically connected to your network, leave the zone configuration as the default, which is to apply the discovery rule to all zones in your deployment. (If no zones are excluded, the discovery policy is applied to all zones.)

Understanding Port Exclusions

LICENSE: FireSIGHT

Just as you can exclude hosts from monitoring (see [Understanding Actions and Discovered Assets](#) on page 1278), you can exclude specific ports from monitoring.

For example, load balancers can report multiple applications on the same port in a short period of time. You can configure your network discovery policy so that it excludes that port from monitoring, such as excluding port 80 on a load balancer that handles a web farm.

As another scenario, your organization may use a custom client that uses a specific range of ports. If the traffic from this client generates excessive and misleading events, you can exclude those ports from monitoring. Similarly, you may decide that you do not want to monitor DNS traffic. In that case, you could configure your policy so that it does not monitor port 53.

When adding ports to exclude, you can decide whether to use a reusable port object from the Available Ports list, add ports directly to the source or destination exclusion lists, or create a new reusable port and then move it into the exclusion lists.

Note that you cannot configure NetFlow-enabled devices to exclude ports from monitoring.

Adding a Discovery Rule

LICENSE: FireSIGHT

You can configure discovery rules to tailor the discovery of host and application data to your needs. Note that when you modify an object referenced in a rule, you must reapply the network discovery policy for those changes to take effect.

To add a discovery rule:

ACCESS: Admin/Discovery Admin

1. Check your access control policies to ensure that you are logging connections as needed for the traffic where you want to discover network data. For more information, see [Logging Connection, File, and Malware Information](#) on page 515 and [Logging Connections for the Default Action](#) on page 431. To discover the most data, log at the end of the connection for traffic you want to discover.

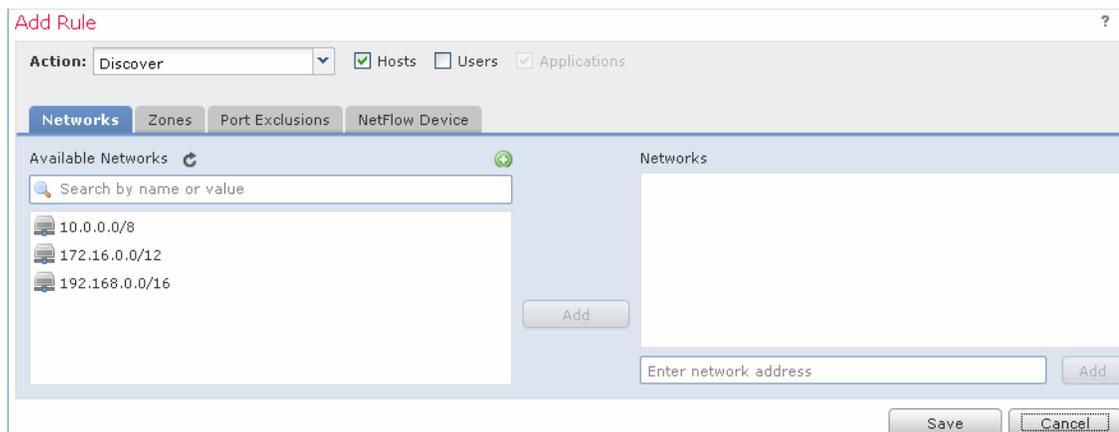
2. Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

Networks	Zones	Source Port Excl...	Destination Port Excl...	NetFlow Device	Action
0.0.0.0/0	any	none	none	none	✓ Discover: Applicatio  

3. Click **Add Rule**.

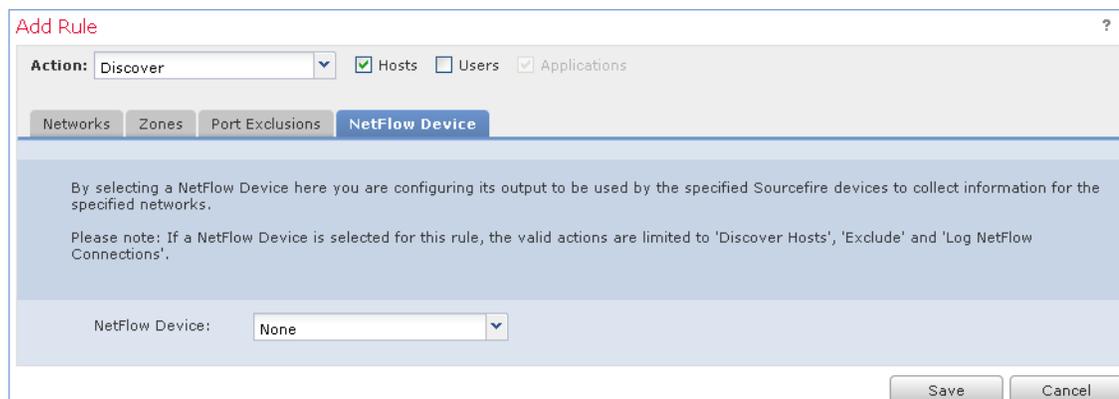
The Add Rule pop-up window appears.



4. You have two options:

- If you plan to use the rule to monitor NetFlow traffic, within the Add Rule pop-up window, click **NetFlow Device**.

The NetFlow Device page appears.



Note that the NetFlow page is available only if you have added a NetFlow device to the discovery policy. For more information, see [Adding NetFlow-Enabled Devices](#) on page 1292.

- If you plan to use the rule to monitor managed devices, skip to step 6.

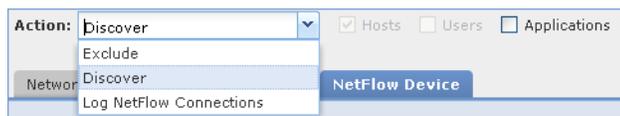
For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272 and [Understanding Device Selection](#) on page 1277

5. Select the IP address for the NetFlow device you want to use from the drop-down list.

6. Set the action for the rule:
 - To exclude all traffic that matches the rule from network discovery, select **Exclude**. Note that the Port Exclusions tab is disabled when you select this rule action.
 - To discover the selected types of data in traffic that matches the rule, select **Discovery** and select or clear the appropriate data type check boxes.

If monitoring managed device traffic, application logging is required. If monitoring users, host logging is required. If monitoring NetFlow traffic, note that you cannot log users and that logging applications is optional.

- If monitoring NetFlow traffic, to use the rule to log connections in NetFlow traffic, select **Log NetFlow Connections**. Note that this option only appears after you have selected a NetFlow device in the rule.



IMPORTANT! The system detects connections in NetFlow traffic based on network discovery policy settings. Connection logging in managed device traffic is configured in the access control policy. For more information, see [Logging Connections for the Default Action](#) on page 431.

For more information on rule actions and discovery of assets, see [Understanding Actions and Discovered Assets](#) on page 1278

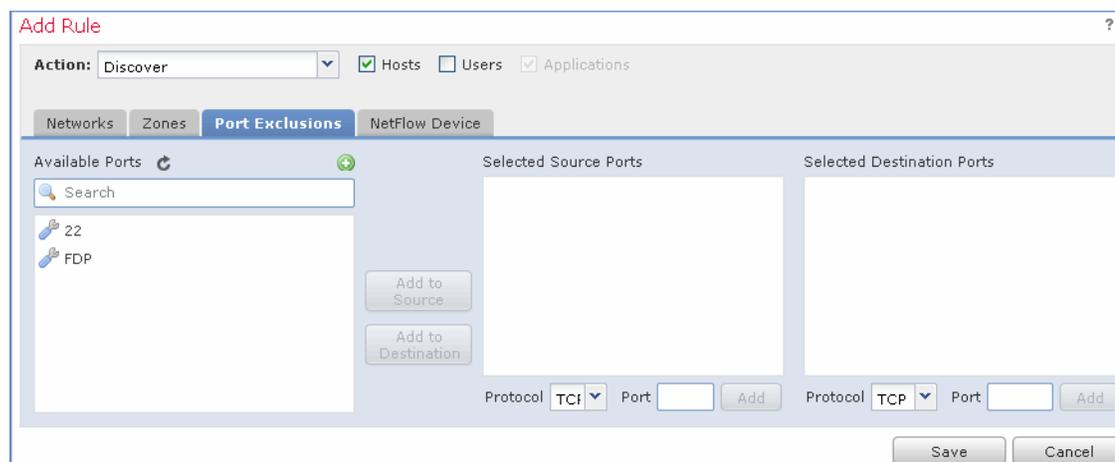
7. Every discovery rule must include at least one network. Optionally, to restrict the rule action to specific networks, click the **Networks** tab, select a network from the **Available Networks** list, and click **Add**, or type the network below the Networks list and click **Add**.

For information on network monitoring, see [Understanding Monitored Networks](#) on page 1279. For information on adding network objects to the Available Networks list, see [Creating Network Objects](#) on page 1284. Note that If you modify a network object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

8. Optionally, to restrict the rule actions to traffic in specific zones, click **Zones**, select a zone or zones from the **Available Zones** list, and click **Add**.

For information on selecting zones for monitoring, see [Understanding Zones in Network Discovery Policies](#) on page 1279.

- To exclude ports from monitoring, click **Port Exclusions**.
The Port Exclusions page appears.



- To exclude specific source ports from monitoring, you have two options:
 - Select a port or ports from the **Available Ports** list and click **Add to Source**.
 - To exclude traffic from a specific source port without adding a port object, under the **Selected Source Ports** list, select the appropriate protocol from the **Protocol** drop-down list, type a port number from 1 to 65535 into the **Port** field, and click **Add**.
For information on excluding ports from monitoring, see [Understanding Port Exclusions](#) on page 1280. For information on adding port objects to the Available Ports list, see [Creating Port Objects](#) on page 1285. Note that if you modify a port object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.
- To exclude specific destination ports from monitoring, you have two options:
 - Select a port or ports from the **Available Ports** list and click **Add to Destination**.
 - To exclude traffic from a specific destination port without adding a port object, under the **Selected Destination Ports** list, select the appropriate protocol from the **Protocol** drop-down list, type a port number from 1 to 65535 into the **Port** field, and click **Add**.
- If you are finished editing the rule, click **Save** to return to the discovery policy rule list.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Creating Network Objects

LICENSE: FireSIGHT

The list of available networks that appears in a discovery rule contains reusable network object and groups that can be used anywhere in the Sourcefire 3D System. You can add new network objects to the list. Note that when you modify an object referenced in a rule, you must reapply the network discovery policy for those changes to take effect.

To create a new network object:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**.
The Network Discovery Policy page appears.
2. Click **Add Rule**.
The Add Rule pop-up window appears.
3. On the Networks page, click the add icon (+).
The Network Objects pop-up window appears.



4. Type a **Name** for the network object.
5. For each IP address, CIDR block, and prefix length you want to add to the network object, type its value and click **Add**.
6. Click **Save** to add the network object to the Available Networks list.

TIP! If the network does not immediately appear on the list, click the refresh icon (↻).

Creating Port Objects

LICENSE: FireSIGHT

The list of available ports that appears in a discovery rule contains reusable port objects and groups that can be used anywhere in the Sourcefire 3D System. You can add new port objects to the list. Note that when you modify an object referenced in a rule, you must reapply the network discovery policy for those changes to take effect.

To create a new port object:

ACCESS: Admin/Discovery Admin

1. Click **Port Exclusions**.
The Port Exclusions page appears.
2. To add a port to the Available Ports list, click the add object icon (+).
The Port Objects pop-up window appears.



3. Supply a **Name** for the port object.
4. In the **Protocol** field, specify the protocol of the traffic you want to exclude. Select **TCP**, **UDP**, or **Other** and choose an option from the drop-down list to select a protocol or **All**.
5. In the **Port(s)** field, enter the ports you want to exclude from monitoring. You can specify a single port, a range of ports using the dash (-), or a comma-separated list of ports and port ranges. Allowed port values are from 1 to 65535.
6. Click **Save** to add the port to the Available Ports list.

TIP! If the port does not immediately appear on the list, click the refresh icon (↻).

Restricting User Logging

LICENSE: FireSIGHT

When you apply a network discovery policy with rules that discover users, users are discovered in traffic that uses the AIM, IMAP, LDAP, Oracle, POP3, and SIP protocols. These users are added to the users table, accessible through the

Analysis menu. You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information.

The total number of detected users the Defense Center can store depends on your FireSIGHT license. After you reach the licensed limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database. Restricting protocol detection helps minimize user name clutter and preserve FireSIGHT user licenses.

For example, obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.

As another example, AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Defense Center cannot correlate these users with other types of users.

Keep in mind that only managed devices can detect non-LDAP user logins. If you are using only Sourcefire User Agents installed on Microsoft Active Directory servers to detect user activity, restricting non-LDAP logins has no effect. Also, you cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

You can choose whether or not to record failed login attempts for failed user logins detected in LDAP, POP3, or IMAP traffic. A failed login attempt does not add a new user to the list of users in the database. Note that the Sourcefire User Agent does not report failed login activity. The user activity type for detected failed login activity is Failed User Login.

To restrict the protocols where user logins are detected:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery.**

The Network Discovery Policy page appears.

2. Click **User**.

The User page appears.

Networks	User	Advanced
Protocol Detection		
aim		Yes
imap		Yes
ldap		Yes
oracle		Yes
pop3		Yes
sip		Yes
Capture Failed Login Attempts		Yes

3. Select check boxes for protocols where you want to detect logins or clear check boxes for protocols where you do not want to detect logins.
4. Click **Save** to save the network policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Configuring Advanced Network Discovery Options

LICENSE: FireSIGHT

The Advanced tab of the network discovery policy allows you to configure policy-wide settings for what events are detected, how long discovery data is retained and how often it is updated, what vulnerability mappings are used for impact correlation, and how operating system and server identity conflicts are resolved. In addition, you can add host input sources and NetFlow devices to allow import of data from other sources.

Note that the database event limits for discovery and user activity events are set in the system policy. For more information, see [Configuring Database Event Limits](#) on page 1978.

To configure advanced settings:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**.

The Network Discovery Policy page appears.

2. Click **Advanced**.

The Advanced page appears.

The screenshot shows the 'Advanced' configuration page for a Network Discovery Policy. It is divided into several sections, each with a title and a pencil icon for editing:

- General Settings**:
 - Capture Banners: No
 - Update Interval: 3600
- Identity Conflict Settings**:
 - Generate Identity Conflict: No
 - Automatically Resolve Conflicts: (Disabled)
- Vulnerabilities to use for Impact Assessment**:
 - Use Network Discovery Vulnerability Mappings: Yes
 - Use Third-Party Vulnerability Mappings: Yes
- NetFlow Devices**:
 - NetFlow Device: (Empty table with a plus icon)
- Network Discovery Data Storage Settings**:
 - When Host Limit Reached: Drop hosts
 - Host Timeout (minutes): 10080
 - Server Timeout (minutes): 10080
 - Client Timeout (minutes): 10080
- Event Logging Settings**:
 - All events enabled.
- OS and Server Identity Sources**:

Name	Type	Timeout
Nmap	Scanner	0 hours

3. Edit advanced settings as needed:

- [Configuring General Settings](#) on page 1288
- [Configuring Identity Conflict Resolution](#) on page 1289
- [Enabling Vulnerability Impact Assessment Mappings](#) on page 1291
- [Adding NetFlow-Enabled Devices](#) on page 1292
- [Configuring Data Storage](#) on page 1293
- [Configuring Discovery Event Logging](#) on page 1295
- [Adding Identity Sources](#) on page 1295

4. When you finish configuring settings, click **Save** to save the policy.

5. When the policy is complete and saved, apply the policy to put the updated settings into effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Configuring General Settings

LICENSE: FireSIGHT

The general settings control how often the system updates information in the network map and whether server banners are captured during discovery.

Capture Banners

Select this check box if you want the system to store header information from network traffic that advertises server vendors and versions (“banners”). This information can provide additional context to the information gathered. You can access server banners collected for hosts by accessing server details.

Update Interval

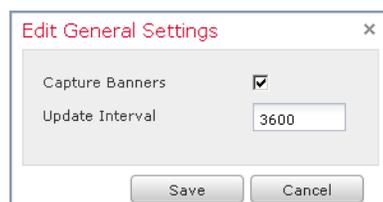
The interval at which the system updates information (such as when any of a host's IP addresses was last seen, when an application was used, or the number of hits for an application). The default setting is 3600 seconds (1 hour).

Note that setting a lower interval for update timeouts provides more accurate information in the host display, but generates more network events.

To update general settings:

ACCESS: Admin/Discovery Admin

1. Click the edit icon (✎) next to **General Settings**.
The General Settings pop-up window appears.



2. Update the settings as needed.
3. Click **Save** to save the general settings and return to the Advanced tab of the network discovery policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Configuring Identity Conflict Resolution

LICENSE: FireSIGHT

The system matches fingerprints for operating systems and servers against patterns in traffic to determine what operating system and which applications are running on a particular host. To provide the most reliable operating system and server identity information, the system collates fingerprint information from several sources.

The system uses all passive data to derive operating system identities and assign a confidence value. For more information on current identities and how the system selects the current identity, see [Enhancing Your Network Map](#) on page 1644.

By default, unless there is an identity conflict, identity data added by a scanner or third-party application overrides identity data detected by the Sourcefire 3D System. You can use the Identity Sources settings to rank scanner and third-party application fingerprint sources by priority. The system retains one identity for each

source, but only data from the highest priority third-party application or scanner source is used as the current identity. Note, however, that user input data overrides scanner and third-party application data regardless of priority.

An identity conflict occurs when the system detects an identity that conflicts with an existing identity that came from either the active scanner or third-party application sources listed in the Identity Sources settings or from a Sourcefire 3D System user. By default, identity conflicts are not automatically resolved and you must resolve them through the host profile or by rescanning the host or re-adding new identity data to override the passive identity. However, you can set your system to always automatically resolve the conflict by keeping the passive identity or to always resolve it by keeping the active identity.

Generate Identity Conflict Event

Enable this option to generate an event when an identity conflict occurs on a host in the network map.

Automatically Resolve Conflicts

You have the following options:

- To force manual conflict resolution of identity conflicts, select **Disabled** from the **Automatically Resolve Conflicts** drop-down list.
- To use the passive fingerprint when an identity conflict occurs, select **Identity** from the **Automatically Resolve Conflicts** drop-down list.
- To use the current identity from the highest priority active source when an identity conflict occurs, select **Keep Active** from the **Automatically Resolve Conflicts** drop-down list.

To update identity conflict resolution settings:

ACCESS: Admin/Discovery Admin

1. Click the edit icon (✎) next to **Identity Conflict Settings**.

The Edit Identity Conflict Settings pop-up window appears.



The screenshot shows a dialog box titled "Edit Identity Conflict Settings". It contains two settings: "Generate Identity Conflict" with an unchecked checkbox, and "Automatically Resolve Conflicts" with a dropdown menu currently set to "(Disabled)". At the bottom are "Save" and "Cancel" buttons.

2. Update the settings as needed.

3. Click **Save** to save the identity conflict settings and return to the **Advanced** tab of the network discovery policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Enabling Vulnerability Impact Assessment Mappings

LICENSE: FireSIGHT

You can configure how the Sourcefire 3D System performs impact flag correlation with intrusion events.

Your options are as follows:

- Select **Use Network Discovery Vulnerability Mappings** if you want to use Sourcefire vulnerability information to perform impact flag correlation.
- Select **Use Third-Party Vulnerability Mappings** if you want to use third-party vulnerability references to perform impact flag correlation. For more information, see [Mapping Third-Party Vulnerabilities](#) on page 1687 or the *Sourcefire 3D System Host Input API Guide*.

You can select either or both of the check boxes. If the system generates an intrusion event and the host involved in the event has servers or an operating system with vulnerabilities in the selected vulnerability mapping sets, the intrusion event is marked with the Vulnerable (level 1: red) impact flag. For any servers which do not have vendor or version information, note that you need to configure vulnerability mapping in the system policy. For more information, see [Mapping Vulnerabilities for Servers](#) on page 1996.

If you clear both check boxes, intrusion events will **never** be marked with the Vulnerable (level 1: red) impact flag. For more information, see [Using Impact Flags to Evaluate Events](#) on page 641.

To update vulnerability settings:

ACCESS: Admin/Discovery Admin

1. Click the edit icon (🔧) next to **Vulnerabilities to use for Impact Assessment**. The Edit Vulnerability Settings pop-up window appears.



2. Update the settings as needed.

3. Click **Save** to save the vulnerability settings and return to the Advanced tab of the network discovery policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Adding NetFlow-Enabled Devices

LICENSE: FireSIGHT

If you have enabled the NetFlow feature on your NetFlow-enabled devices, you can use the connection data exported by these devices to supplement the connection data collected by Sourcefire devices.

Before you can use them in discovery rules, you must configure the NetFlow-enabled devices you plan to use (see [Preparing to Analyze NetFlow Data](#) on page 1274), then add them to the network discovery policy.

For more information on using NetFlow data with the Sourcefire 3D System, including information on additional prerequisites, see [Understanding NetFlow](#) on page 1271.

To add NetFlow-enabled devices for connection data collection:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**.
The Network Discovery Policy page appears.
2. Click **Advanced**.
The Advanced page appears.
3. Click the add icon (+) next to NetFlow Devices.
The Add NetFlow Device pop-up window appears.



4. In the **IP Address** field, enter the IP address of the NetFlow-enabled device you want to use to collect connection data.

5. To add additional NetFlow-enabled devices, repeat steps 3 and 4.

TIP! To remove a NetFlow-enabled device, click the delete icon () next to the device you want to remove. Keep in mind that if you use a NetFlow-enabled device in a discovery rule, you must delete the rule before you can delete the device from the Advanced page. For more information, see [Working with Discovery Rules](#) on page 1276.

6. Click **Save**.
The device appears on the list of NetFlow-enabled devices.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Configuring Data Storage

LICENSE: FireSIGHT

Data storage settings control the kinds of data stored in the database, and therefore determine the data that the Sourcefire 3D System can use. These settings also control how long data is retained in the network map.

The following options comprise the network discovery data storage settings.

When Host Limit Reached

You can control how hosts are handled when the Defense Center reaches its host limit (as determined by the FireSIGHT license) and the network map is full. This option is especially valuable if you want to prevent spoofed hosts from taking the place of valid hosts in the network map. To drop old hosts, select **Drop hosts** from the **When Host Limit Reached** drop-down list. To drop new hosts, select **Don't insert new hosts** from the **When Host Limit Reached** drop-down list. For more information, see [Understanding FireSIGHT Host and User License Limits](#) on page 2045.

Host Timeout

The amount of time that passes, in minutes, before the system drops a host from the network map due to inactivity. The default setting is 10080 minutes (7 days). Individual host IP and MAC addresses can time out individually, but a host does not disappear from the network map unless all of its associated addresses have timed out.

To avoid premature timeout of hosts, make sure that the host timeout value is longer than the update interval in the network discovery policy. For more information on the update interval, see [Configuring General Settings](#) on page 1288.

Server Timeout

The amount of time that passes, in minutes, before the system drops a server from the network map due to inactivity. The default setting is 10080 minutes (7 days).

To avoid premature timeout of servers, make sure that the service timeout value is longer than the update interval in the network discovery policy. For more information, see [Configuring General Settings](#) on page 1288.

Client Application Timeout

The amount of time that passes, in minutes, before the system drops a client from the network map due to inactivity. The default setting is 10080 minutes (7 days).

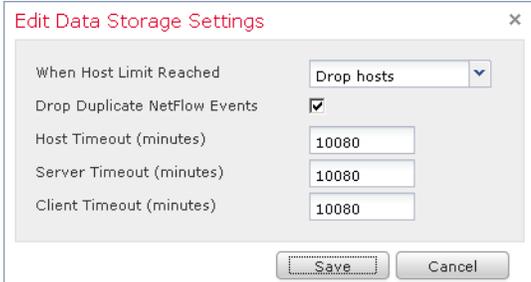
You should make sure that the client timeout value is longer than the update interval in the network discovery policy. For more information, see [Configuring General Settings](#) on page 1288.

To update data storage settings:

ACCESS: Admin/Discovery Admin

1. Click the edit icon (✎) next to **Data Storage Settings**.

The Data Storage Settings pop-up window appears.



2. Update the settings as needed.
3. Click **Save** to save the data storage settings and return to the Advanced tab of the network discovery policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Configuring Discovery Event Logging

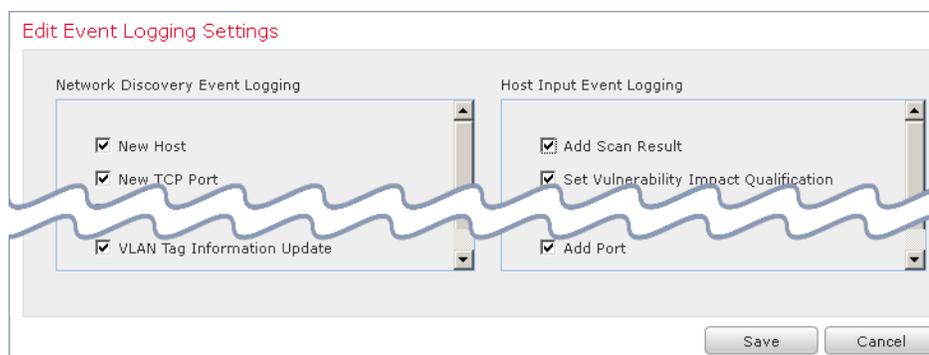
LICENSE: FireSIGHT

The Event Logging Settings control whether discovery and host input events are logged. If you do not log an event, you cannot retrieve it in event views or use it to trigger correlation rules.

To set event logging settings:

ACCESS: Admin/Discovery Admin

1. Click the edit icon (✎) next to **Event Logging Settings**.
The Event Logging Settings pop-up window appears.



2. Select or clear the check boxes next to the discovery and host input event types you want to log in the database. See [Understanding Discovery Event Types](#) on page 1389 and [Understanding Host Input Event Types](#) on page 1394 for information about each event type.
3. Click **Save** to save the event logging settings and return to the Advanced tab of the network discovery policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Adding Identity Sources

LICENSE: FireSIGHT

You can add new active sources through this page, or change the priority or timeout settings for existing sources. Note that adding a scanner to this page does not add the full integration capabilities that exist for the Nmap scanners, but does allow integration of imported third-party application or scan results. If you import data from a third-party application or scanner, remember to make sure that you map vulnerabilities from the source to the vulnerabilities in the network map. For more information, see [Mapping Third-Party Vulnerabilities](#) on page 1687.

To add identity sources:

ACCESS: Admin/Discovery Admin

1. Click the edit icon (✎) next to **OS and Server Identity Sources**.
The Edit OS and Server Identity Sources pop-up window appears.



2. To add a new source, click **Add Source**.
The Add Identity Source pop-up window appears.

Name: Identity Source name
Type: [dropdown]
Timeout: [text] Hours [dropdown]

3. Type a **Name** for the source.
4. Select **Scanner** or **Application** from the **Type** drop-down list.
5. To indicate the duration of time that should elapse between the addition of an identity to the network map by this source and the deletion of that identity, select **Hours**, **Days**, or **Weeks** from the **Timeout** drop-down list and type the appropriate duration.

TIP! To delete a source that you added, click the delete icon (🗑) next to the source.

6. Optionally, to promote a source and cause the operating system and application identities to be used in favor of sources below it in the list, select the source and click the up arrow.
7. Optionally, to demote a source and cause the operating system and application identities to be used only if there are no identities provided by sources above it in the list, select the source and click the down arrow.

8. Click **Save** to save the identity source settings and return to the Advanced tab of the network discovery policy.

IMPORTANT! You must apply the network discovery policy for your changes to take effect. For more information, see [Applying the Network Discovery Policy](#) on page 1297.

Applying the Network Discovery Policy

LICENSE: FireSIGHT

By default, the network discovery policy is applied to any targeted zones on managed devices when they are registered with the Defense Center. Applying the network discovery policy allows the system to begin monitoring your network according to your specifications. If you change the network discovery policy, you must reapply it before your changes take effect.

When you reapply the network discovery policy:

- the system deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks
- the affected managed devices discard any discovery data that has not yet been sent to the Defense Center

When you apply a network discovery policy, make sure that you have already applied an access control policy to the targeted zones on managed devices. If an access control policy has not been applied, the network discovery policy apply fails. Note that you cannot apply a network discovery policy on a Defense Center where no FireSIGHT license is installed.

If you modify a network or port object used in the network discovery policy, you must reapply the policy for those changes to take effect for discovery.

Note that you cannot apply a network discovery policy to stacked devices running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails).

To apply the network discovery policy:

ACCESS: Admin/Security Approver

1. Select **Policies > Network Discovery**.
The Network Discovery Policy page appears.
2. Click **Apply**.
A message appears, confirming that you want to apply the policy to all zones targeted by access control policies on the Defense Center.
3. Click **Yes** to apply the policy.

Obtaining User Data from LDAP Servers

LICENSE: FireSIGHT

The Sourcefire 3D System can obtain both user identity and user activity information from your organization's LDAP servers.

Sourcefire User Agents allow you to monitor users when they authenticate against Active Directory credentials on Microsoft Active Directory servers. You can install an agent on any Microsoft Windows 7 or Microsoft Windows Server 2008 device with TCP/IP access to the Microsoft Active Directory servers you want to monitor. Each agent can monitor logins on up to five servers.

The agents send records of those logins to the Defense Center, which logs and reports them as user activity. This supplements any user activity detected directly by managed devices. More important, the logins reported by User Agents associate users with IP addresses, which in turn allows access control rules with user conditions to trigger.

You can configure a connection between the Defense Center and LDAP servers. This connection not only allows you to retrieve metadata for the users whose logins were detected by User Agents, but also is used to specify the users and groups you want to use in access control rules.

IMPORTANT! Legacy agents, which you install on your Microsoft Active Directory servers, also monitor users when they authenticate against Active Directory credentials. However, you should plan to transition to Version 2.0 of the Sourcefire User Agent as soon as possible in preparation for end of support for legacy agents in future releases.

For more information, see:

- [Understanding User Data Collection](#) on page 1252
- [Adding User Conditions](#) on page 497
- [Creating LDAP Connections with the Defense Center](#) on page 1298
- [Enabling and Disabling User Awareness LDAP Connections](#) on page 1306
- [Performing an On-Demand User Data Retrieval for Access Control](#) on page 1307
- [Configuring Defense Center-User Agent Connections](#) on page 1307

Creating LDAP Connections with the Defense Center

LICENSE: FireSIGHT

If you want to perform user control (that is, write access control rules with user conditions), you must configure a connection between the Defense Center and at least one of your organization's Microsoft Active Directory servers. This configuration, called an *LDAP connection* or a *user awareness authentication object*, contains connection settings and authentication filter settings for the

server. The connection's user and group access control parameters specify the users and groups you can use in access control rules.

IMPORTANT! If you want to perform user control, you **must** use Microsoft Active Directory. The system uses Sourcefire User Agents running on Active Directory servers to associate users with IP addresses, which is what allows access control rules to trigger.

Note that you can also create authentication objects to manage external authentication to the Sourcefire 3D System's web interface; see [Managing Authentication Objects](#) on page 1854. Those objects are similar to the authentication objects you create for user control, and you configure them in a similar way.

After you create an LDAP connection for user control, the Defense Center queries the LDAP server on a schedule that you specify. If you add new users or remove users from the LDAP server, you must wait until the Defense Center performs its scheduled update for those changes to take effect for access control. Alternately, you can perform an on-demand query.

The Defense Center-LDAP server connection also allows you to retrieve metadata for users, both access-controlled and non-access-controlled, whose logins were detected by User Agents, as well as for certain users whose activity was detected directly by managed devices. The Defense Center regularly queries the LDAP server to obtain metadata for new LDAP, POP3, and IMAP users whose activity was detected since the last query. If a user already exists in the Defense Center's Users database, the Defense Center updates the metadata if it has not been updated in the last 12 hours.

The Defense Center uses the email addresses in POP3 and IMAP logins to correlate with users on the LDAP server. For example, if a managed device detects a POP3 login for a user with the same email address as an LDAP user, the Defense Center associates the LDAP user's metadata with that user. Note that it may take several minutes for the Defense Center to update with user metadata after the system detects a new user login.

The Defense Center obtains the following information and metadata about each user:

- LDAP user name
- first and last name
- email address

- department
- telephone number

IMPORTANT! If you remove a user that has been detected by the system from your LDAP servers, the Defense Center 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 Defense Center next updates its list of access-controlled users.

For more information, see:

- [Preparing to Connect to an LDAP Server](#) on page 1300
- [Creating an LDAP Connection for User Control](#) on page 1301

Preparing to Connect to an LDAP Server

LICENSE: FireSIGHT

The Sourcefire 3D System supports connections to LDAP servers running the following:

- Microsoft Active Directory on Windows Server 2003 and Windows Server 2008
- Oracle Directory Server Enterprise Edition 7.0 on Windows Server 2003 and Windows Server 2008
- OpenLDAP on Linux

You must have TCP/IP access from the Defense Center to the LDAP servers. In addition, your LDAP servers must use the LDAP field names shown in the following table. For example, the system maps the givenname metadata for a particular user on an LDAP server to the first name of that user in the Defense Center database. If you rename the field on the LDAP server, the Defense Center cannot populate its database with the information in that field.

Mapping LDAP Fields to Sourcefire Fields

DEFENSE CENTER FIELD	MICROSOFT ACTIVE DIRECTORY	ORACLE DIRECTORY SERVER	OPENLDAP
Username	samaccountname	cn uid	cn uid
First Name	givenname	givenname	givenname
Last Name	sn	sn	sn

Mapping LDAP Fields to Sourcefire Fields (Continued)

DEFENSE CENTER FIELD	MICROSOFT ACTIVE DIRECTORY	ORACLE DIRECTORY SERVER	OPENLDAP
Email	mail userprincipalname (if mail has no value)	mail	mail
Department	department distinguishedname (if department has no value)	department	ou
Phone	telephonenumber	n/a	telephonenumber

Creating an LDAP Connection for User Control

LICENSE: FireSIGHT

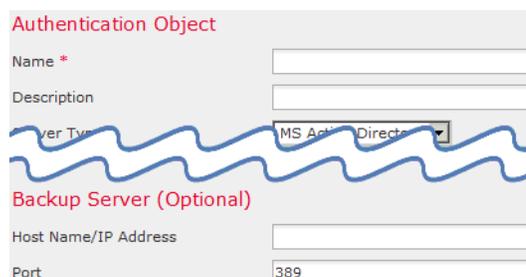
You configure a connection between the Defense Center and an LDAP server by creating a *user awareness* authentication object. This object contains connection settings and authentication filter settings for the LDAP server from which you want to retrieve user information. It also specifies the users and groups you can use in access control rules. The method you use to create a user awareness authentication object is similar to creating an external authentication object, as described in [Managing Authentication Objects](#) on page 1854.

TIP! To delete an LDAP connection, click the delete icon (🗑️) and confirm that you want to delete it. To modify a connection, click the edit icon (✏️) and see the procedure in this section for settings you can configure. If the connection is enabled, your changes take effect when the Defense Center next queries the LDAP server.

The following list contains the information you must provide when creating an LDAP connection. You should work closely with your LDAP administrators to obtain the information.

Server Type, IP Address, and Port

You must specify the server type, IP address or hostname, and port for a primary, and optionally a backup, LDAP server. If you want to perform user control, you **must** use a Microsoft Active Directory server.



The screenshot shows a configuration form with the following fields:

- Authentication Object** (Section Header)
- Name * (Text input)
- Description (Text input)
- Server Type (Dropdown menu, currently showing "MS Active Directory")
- Backup Server (Optional)** (Section Header)
- Host Name/IP Address (Text input)
- Port (Text input, containing "389")

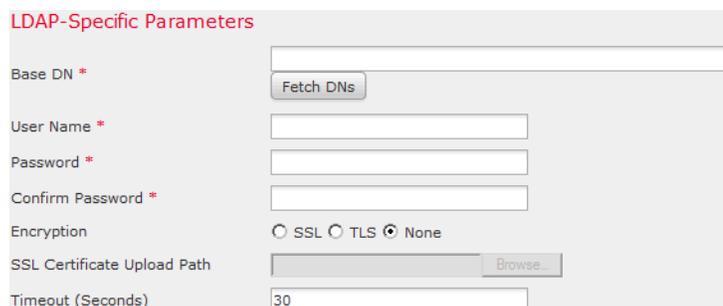
LDAP-Specific Parameters

When the Defense Center searches the LDAP directory server to retrieve user information on the authentication server, it needs a starting point for that search. You can specify the *namespace*, or directory tree, that the local appliance should search by providing a base distinguished name, or *base DN*. 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`. Note that after you identify a primary server, you can automatically retrieve a list of available base DNs from the server and select the appropriate base DN.

You must supply user credentials for a user with appropriate rights to the user information you want to retrieve. Remember that the distinguished name for the user you specify must be unique to the directory information tree for the directory server.

You can also specify an encryption method for the LDAP connection. Note that if you are using a certificate to authenticate, the name of the LDAP server in the certificate **must** match the host name that you specified in the Defense Center web interface. For example, if you use `10.10.10.250` when configuring the LDAP connection but `computer1.example.com` in the certificate, the connection fails.

Finally, you must specify the timeout period after which attempts to contact an unresponsive LDAP server roll over to the backup connection.



The screenshot shows a configuration form with the following fields:

- LDAP-Specific Parameters** (Section Header)
- Base DN * (Text input with a "Fetch DNs" button)
- User Name * (Text input)
- Password * (Text input)
- Confirm Password * (Text input)
- Encryption (Radio buttons: SSL, TLS, None)
- SSL Certificate Upload Path (Text input with a "Browse..." button)
- Timeout (Seconds) (Text input, containing "30")

User and Group Access Control Parameters

If you enable an authentication object for user awareness, you must specify the groups you want to use in access control.

Including a group automatically includes all of that group's members, including members of any sub-groups. However, if you want to use the sub-group in access control rules, you must explicitly include the sub-group. You can also exclude groups and individual users. Excluding a group excludes all the members of that group, even if the users are members of an included group.

The maximum number of users you can use in access control depends on your FireSIGHT license. When choosing which users and groups to include, make sure the total number of users is less than your FireSIGHT user license.

IMPORTANT! If you do not specify any groups to include, the system retrieves user data for all the groups that match the LDAP parameters you provided. For performance reasons, Sourcefire recommends that you explicitly include only the groups that represent the users you want to use in access control. Note that you **cannot** include the Users or Domain Users groups.

You must also specify how often the Defense Center queries the LDAP server to obtain new users to use in access control.

The screenshot shows the 'User/Group Access Control Parameters' configuration window. At the top, the title is 'User/Group Access Control Parameters'. Below the title, there is an 'Enable' checkbox which is checked. To the right of the checkbox are two list boxes: 'Available Groups' containing '<none>' and 'Groups to Include' containing '<all>'. Between these two boxes are '>' and '<' buttons. To the right of the 'Groups to Include' box is a note: 'Users in Groups to Include will be used for User Awareness.' Below these elements is a section for 'Update Schedule' with a blue button labeled 'Update Schedule', a 'Start at (Hours)' dropdown set to '0', and an 'Update Interval (Hours)' dropdown set to '24'. At the bottom left, there is a note: '*Required Field'.

To create an LDAP connection for user control:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Users**.

The Users Policy page appears.



2. Click **Add LDAP Connection**.

The Create User Awareness Authentication Object page appears.

3. Type a **Name** and **Description** for the object.

4. Select the LDAP **Server Type**.

If you want to perform user control, you **must** use a Microsoft Active Directory server. If you use any other type of LDAP server, you are limited to retrieving metadata for some users whose activity was detected directly by managed devices (as opposed to by Sourcefire User Agents).

IMPORTANT! Sourcefire User Agents cannot transmit Active Directory user names ending with the \$ character to the Defense Center. If your Active Directory server contains such user names, you must edit those names to remove the final \$ character if you want to monitor them.

5. Specify an **IP Address** or **Host Name** for a primary and, optionally, a backup LDAP server.

6. Specify the **Port** that your LDAP servers use for authentication traffic.

7. Specify the **Base DN** for the LDAP directory you want to access.

For example, to authenticate names in the Security organization at the Example company, type `ou=security,dc=example,dc=com`.

TIP! To fetch a list of all available domains, click **Fetch DNs** and select the appropriate base distinguished name from the drop-down list.

8. Specify the distinguished **User Name** and **Password** that you want to use to validate access to the LDAP directory. Confirm the password.
For example, if you are connecting to an OpenLDAP server where user objects have a **uid** attribute and the object for the administrator in the Security division at our example company has a **uid** value of **NetworkAdmin**, you would type **uid=NetworkAdmin,ou=security,dc=example,dc=com**.
9. Choose an **Encryption** method. If you are using encryption, you can add an **SSL Certificate**.
The host name in the certificate **must** match the host name of the LDAP server you specified in step 5.
10. Specify the **Timeout** period (in seconds) timeout period after which attempts to contact an unresponsive primary LDAP server roll over to the backup connection.
11. Optionally, before you specify user awareness settings for the object, test the connection by clicking **Test**.
12. You have two options, depending on the type of LDAP server you selected in step 4:
 - If you are connecting to an Active Directory server, you can enable **User/Group Access Control Parameters** to specify users to use in access control. Continue with the next step.
 - If you are connecting to any other kind of server, or do not want to perform user control, skip to step 17.
13. Click **Fetch Groups** to populate the available groups list using the LDAP parameters you provided.
14. Specify the users you want to use in access control by using the right and left arrow buttons to include and exclude groups.
Including a group automatically includes all of that group's members, including members of any sub-groups. However, if you want to use the sub-group in access control rules, you must explicitly include the sub-group. Excluding a group excludes all the members of that group, even if the users are members of an included group.
15. Specify any particular **User Exclusions**.
Excluding a user prevents you from writing an access control rule using that user as a condition. Separate multiple users with commas. You can also use an asterisk (*) as a wildcard character in this field.

16. Specify how often you want the Defense Center to query the LDAP server to obtain new user and group information.

By default, the Defense Center queries the server once a day at midnight:

- Use the **Start At** drop-down list to specify when you want the query to occur. **0** represents midnight, **1** represents 1:00 AM, and so on.
- Use the **Update Interval** drop-down list to specify how often, in hours, you want to query the server.

17. Click **Save**.

If you added or made changes to user and group access control parameters, confirm that you want to implement your changes. The object is saved. Note that you must enable the connection before the Defense Center can query the LDAP server; see the next section, [Enabling and Disabling User Awareness LDAP Connections](#).

Enabling and Disabling User Awareness LDAP Connections

LICENSE: FireSIGHT

Only enabled LDAP connections allow the Defense Center to query the LDAP servers. To stop queries, you can temporarily disable LDAP connections rather than deleting them.

When you enable an LDAP connection where you have specified user and group access control parameters, you can force the Defense Center to query the server immediately, or you can wait until the first scheduled query occurs, as defined by the access control parameters in the LDAP connection. Note that you can also perform an on-demand query; see the next section, [Performing an On-Demand User Data Retrieval for Access Control](#).

The maximum number of users the Defense Center can retrieve from the server depends on your FireSIGHT license. If your access control parameters are too broad, the Defense Center obtains information on as many users as it can and reports the number of users it failed to retrieve in the task queue.

To enable or disable an LDAP connection:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Users**.
The Users Policy page appears.
2. Next to the LDAP connection you want to enable or disable, click the slider.
If the connection was enabled, it is disabled. If it was disabled, it is enabled.

3. If you are enabling the connection and your connection has user and group access control parameters, choose whether you want to immediately query the LDAP server to obtain user and group information.

If you do not immediately query the LDAP server, the query occurs at the scheduled time.

The query begins. You can monitor its progress in the task queue (**System > Monitoring > Task Status**).

Performing an On-Demand User Data Retrieval for Access Control

LICENSE: FireSIGHT

If you change the user and group access control parameters in an LDAP connection, or if you change the users or groups on your LDAP server and want your changes to be immediately available for access control, you can force the Defense Center to perform an on-demand user data retrieval from an LDAP server.

The maximum number of users the Defense Center can retrieve from the server depends on your FireSIGHT license. If the access control parameters in your LDAP connection are too broad, the Defense Center obtains information on as many users as it can and reports the number of users it failed to retrieve in the task queue.

To perform an on-demand user data retrieval:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Users**.
The Users Policy page appears.
2. Next to the LDAP connection you want to use to query the LDAP server, click the download icon ().
The query begins. You can monitor its progress in the task queue (**System > Monitoring > Task Status**).

Configuring Defense Center-User Agent Connections

LICENSE: FireSIGHT

If you use Microsoft Active Directory LDAP servers, Sourcefire recommends that you connect Sourcefire User Agents to your Active Directory servers. User Agents monitor users as they log into the network or when accounts authenticate against Active Directory credentials for other reasons (for example, your organization may use services or applications that rely on Active Directory for centralized authentication).

The agents send records of those logins and logoffs to the Defense Center, which logs and reports them as user activity. The Defense Center uses this data in two main ways:

- to supplement user activity detected directly by managed devices, as defined in your network discovery policy
- to associate users with IP addresses, which in turn allows access control rules with user conditions to trigger

IMPORTANT! If you want to perform user control, you **must** install and use Sourcefire User Agents. However, User Agents only detect LDAP logins. If you want to detect other types of logins, you must use managed devices; see [Understanding Actions and Discovered Assets](#).

You can use Version 2.1 of the Sourcefire User Agent to report user logins and logoffs to any Version 5.x Sourcefire 3D System Defense Center. If you have agents prior to Version 2.1, you can continue to use those agents to report Active Directory server login data to your Defense Centers. Note, however, that support for legacy agents will be phased out in future releases. Sourcefire recommends that you transition to Version 2.1 of the Sourcefire User Agent as soon as possible.

You can use the Sourcefire User Agent Status Monitor health module to monitor the heartbeat of agents connected to a Defense Center. For more information, see [Configuring User Agent Status Monitoring](#) on page 2143.

To use a Sourcefire User Agent, first configure the Defense Center to connect to the Windows host where you plan to install the agent. Then, install and configure the agent.

User Agents can connect to up to five Defense Centers at a time. In a high availability deployment, connect agents to both the primary Defense Center and the secondary Defense Center. To do so you must make sure agents can communicate with both the primary Defense Center and the secondary Defense Center.

For more information, see:

- [Configuring the Defense Center to Connect to a Sourcefire User Agent](#) on page 1309
- [Installing a Sourcefire User Agent](#) on page 1310
- [Configuring User and Security Permissions](#) on page 1312
- [Configuring a Sourcefire User Agent](#) on page 1312

Configuring the Defense Center to Connect to a Sourcefire User Agent

LICENSE: FireSIGHT

The first step in collecting LDAP user login information using Sourcefire User Agents is to configure each Defense Center to allow connections from the agents you plan to connect to your Active Directory servers.

TIP! To delete the Defense Center-User Agent connection, click the delete icon () and confirm that you want to delete it.

To configure the Defense Center to connect to a Sourcefire User Agent:

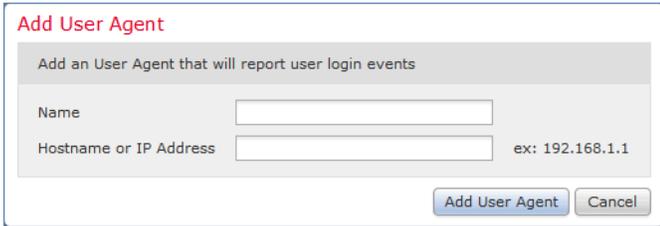
ACCESS: Admin/Discovery Admin

1. Select **Policies > Users**.

The Users Policy page appears.

2. Click **Add Sourcefire User Agent**.

The Add User Agent pop-up window appears.



The screenshot shows a dialog box titled "Add User Agent" with the subtitle "Add an User Agent that will report user login events". It contains two input fields: "Name" and "Hostname or IP Address". The "Hostname or IP Address" field has an example "ex: 192.168.1.1". At the bottom right are "Add User Agent" and "Cancel" buttons.

3. Type a descriptive name for the agent in the **Name** field.
4. Type the IP address or host name of the computer where the agent will reside in the **Hostname or IP Address** field.
5. Click **Add User Agent**.

The Defense Center can now connect to a User Agent on the configured host.

If you want to perform user control (that is, write access control rules with user conditions), you must configure and enable a connection between the Defense Center and at least one of your organization's Microsoft Active Directory servers. This configuration, called an *LDAP connection* or a *user awareness authentication object*, contains connection settings and authentication filter settings for the server. The connection's user and group access control parameters specify the users and groups you can use in access control rules. See [Creating an LDAP Connection for User Control](#) on page 1301 for more information.

Continue with the next section, [Installing a Sourcefire User Agent](#) on page 1310.

Installing a Sourcefire User Agent

LICENSE: FireSIGHT

After you configure the Defense Center to connect to the Windows computer where you plan to install each agent, install and configure the agents. Set up the Windows computer with the following prerequisites:

- The computer is running Windows Vista, Windows 7, Windows 8, Windows Server 2003, Windows Server 2008, or Windows Server 2012. The computer does not have to be an Active Directory server.
- The computer has Microsoft .NET Framework Version 4.0 Client Profile and Microsoft SQL Server Compact (SQL CE) Version 3.5 installed. The framework is available from Microsoft as the .NET Framework Version 4.0 Client Profile redistributable package (`dotNetFx40_Client_x86_x64.exe`). The SQL CE is available from Microsoft as an executable file (`SSCERuntime-ENU.exe`).

IMPORTANT! If you do not have both the .NET Framework and SQL CE installed, when you open the agent executable file (`Sourcefire_User_Agent_2.1.0-build_number_Setup.exe`), it prompts you to download the appropriate files.

- The computer has TCP/IP access to the Active Directory servers you want to monitor, and uses the same version of the Internet Protocol as the Active Directory servers. If the agent is monitoring the Active Directory servers real-time, the computer's TCP/IP access must be on at all times to retrieve login data.
- The computer has TCP/IP access to the Defense Centers where you want to report data and an IPv4 address.
- The computer has an IPv6 address, if you want to detect logoffs from hosts with IPv6 addresses, or an IPv4 address, if you want to detect logoffs from hosts with IPv4 addresses.
- The computer does not have a legacy agent or Version 2.0.x agent already installed. As these agents do not automatically uninstall, to uninstall an existing agent, open **Add/Remove Programs** in the control panel.

Once you set up the computer with the prerequisites, install the agent.

The agent runs as a service using the **Local system** account. If the Windows computer where the agent is running is connected to the network, the service continues to poll and send user data, even if a user is not actively logged into the system.

IMPORTANT! Do not make changes to the service configuration; the agent does not function correctly using a different account.

In a high availability configuration, add both Defense Centers to the agent to enable update of user login data to both the primary and the secondary so the data remains current on both.

To install a Sourcefire User Agent:

ACCESS: Any

1. Download the Sourcefire User Agent setup file (`Sourcefire_User_Agent_2.1.0-build_number_Setup.zip`, where *build_number* represents the number of the agent build) from the [Sourcefire Support Site](#).

IMPORTANT! Download the setup file directly from the Support Site and do not transfer it by email. If you transfer the setup file by email, it may become corrupted.

2. Copy the setup file to the Windows computer where you want to install the agent and unpack the file.
The agent requires 3 MB free on the hard drive for installation. Sourcefire recommends you allocate 4 GB on the hard drive for the agent local database.
3. Open the setup executable file (`Sourcefire_User_Agent_2.1.0-build_number_Setup.exe`).

TIP! If you are using an account that is not a member of the Administrators group and do not have permissions to install new applications on the Windows computer, you must elevate to a user that does belong to the group to have the appropriate permissions to start the installation. To access the escalation option, right click the `Sourcefire_User_Agent_2.1.0-build_number_Setup.exe` file and select **Run As**. Select an appropriate user and supply the password for that user.

4. If you do not have both Microsoft .NET Framework Version 4.0 Client Profile and SQL CE Version 3.5 installed on the Windows computer where you install the agent, you are prompted to download the appropriate files. Download and install the files.
The setup wizard appears.
5. Follow the prompts in the wizard to install the agent.
The agent is installed. The Sourcefire User Agent starts as a service on the Windows system. Continue with [Configuring User and Security Permissions](#) on page 1312.

Configuring User and Security Permissions

After you prepare the computer with all agent prerequisites, configure user permissions and Windows security permissions to allow the agent to communicate with the Active Directory server, access the security logs to retrieve login data, and optionally, retrieve logoff data. Optionally, enable idle session timeouts in the group policy to help prevent the agent from detecting and reporting extraneous logins due to multiple sessions on a host. For more information, see the *Sourcefire 3D System User Agent Configuration Guide*.

Continue with [Configuring a Sourcefire User Agent](#) on page 1312.

Configuring a Sourcefire User Agent

LICENSE: FireSIGHT

Once the agent is installed, you can configure it to receive data from Active Directory servers, report the information to Defense Centers, exclude specific user names and IP addresses from the reporting, and log status messages to a local event log or the Windows application log.

To configure the agent:

ACCESS: Any

- ▶ On the computer where you installed the agent, select **Start > Programs > Sourcefire > Configure Sourcefire User Agent**.

The Sourcefire Active Directory User Agent dialog box appears, with the General tab active.

The [User Agent Configuration Actions](#) table describes the actions you can take when configuring the agent and where to configure them. For more information, see the *Sourcefire 3D System User Agent Configuration Guide*.

User Agent Configuration Actions

To...	YOU CAN...
change the agent name, change the logoff check frequency, and start and stop the service	select the General tab.
add, modify, or remove Active Directory servers, enable real-time Active Directory server data retrieval, and modify the Active Directory server polling interval and maximum poll length	select the Active Directory Servers tab.
add or remove Defense Centers	select the Sourcefire DCs tab.
add, modify, or remove user names excluded from reporting	select the Excluded Usernames tab.
add, modify, or remove IP addresses excluded from reporting	select the Excluded Addresses tab.
view, export, and clear the event log, log to Windows application logs, and modify how long messages should be kept	select the Logs tab.
perform troubleshooting and maintenance tasks, as directed by Sourcefire Support	select the Logs tab, enable Show Debug Messages in Log , then select the Maintenance tab.
save changes to the agent settings	click Save . A message displays below Save stating when you have unsaved changes.
close the agent without saving changes to the agent settings	click Cancel .

CHAPTER 33

USING THE NETWORK MAP

The Sourcefire 3D System passively collects traffic traveling over the network, decodes the data, and then compares it to established operating system and fingerprints. From this information, the system builds a *network map*, which is a detailed representation of your network.

The network map allows you to use the Defense Center to view your network topology in terms of hosts and network devices (bridges, routers, NAT devices, and load balancers). It is a useful tool for a quick, overall view of your network. The network map also allows you to drill down on associated host attributes, applications, clients, and vulnerabilities. In other words, you can select different views of the network map to suit the analysis you perform.

You can augment the information your system collects by adding operating system, application, client, protocol, or host attribute information from a third-party application using the host input feature. You can also actively scan hosts in the network map using Nmap and add the scan results to your network map.

You can use the custom topology feature to help you organize and identify subnets in the views of the network map. For example, if each department in your organization uses a different subnet, you can assign familiar labels to those subnets using the custom topology feature.

For more information, see the following sections:

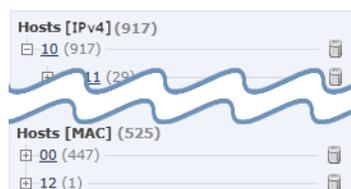
- [Understanding the Network Map](#) on page 1315
- [Working with the Hosts Network Map](#) on page 1316
- [Working with the Network Devices Network Map](#) on page 1318
- [Working with the Mobile Devices Network Map](#) on page 1319

- [Working with the Applications Network Map](#) on page 1321
- [Working with the Vulnerabilities Network Map](#) on page 1323
- [Working with the Host Attributes Network Map](#) on page 1325
- [Working with Custom Network Topologies](#) on page 1326

Understanding the Network Map

LICENSE: FireSIGHT

Each view of the network map has the same format: a hierarchical tree with expandable categories and sub-categories. When you click a category, it expands to show you the sub-categories beneath it. You can select different views of the network map depending on the kind of analysis you are performing.



The Defense Center gathers data from all security zones where discovery policies are applied (including zones that process data from NetFlow-enabled devices). If multiple devices detect the same network asset, the Defense Center combines the information into a composite representation of the asset.

Although you can configure your network discovery policy to add data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

From any network map, you can view any host's *host profile*, which provides a complete view of all the information collected by the system for that host. The host profile contains general information, such as the host name, operating system, and all associated IP addresses, as well as more specific information including detected protocols, applications, and clients that are running on the host. The host profile also includes information about the vulnerabilities associated with the host and its detected assets. For more information on host profiles, see [Using Host Profiles](#) on page 1334.

You can delete an item from the network map if you are no longer interested in investigating it. You can delete hosts and applications from the network map; you can also delete or deactivate vulnerabilities. If the system detects activity associated with a deleted host, it re-adds the host to the network map. Similarly, deleted applications are re-added to the applications network map if the system detects a change in the application (for example, if an Apache web server is upgraded to a new version). Vulnerabilities are reactivated on specific hosts if the system detects a change that makes the host vulnerable.

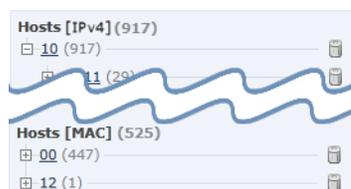
You can also use the network map to deactivate vulnerabilities network-wide, which means that you deem these hosts, which the system has judged to be vulnerable, to be safe from that particular attack or exploit.

TIP! If you want to permanently exclude a host or subnet from the network map, modify the network discovery policy. You may wish to exclude load balancers and NAT devices from monitoring. They may create excessive and misleading events, filling the database and overloading the Defense Center. See [Understanding Host Data Collection](#) on page 1252 for more information.

Working with the Hosts Network Map

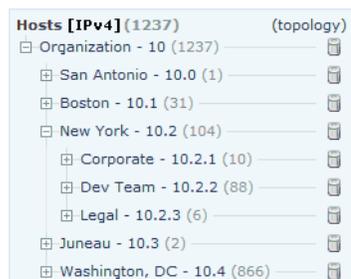
LICENSE: FireSIGHT

Use the hosts network map to view the hosts on your network, organized by subnet in a hierarchical tree, as well as to drill down to the host profiles for specific hosts. This network map view provides a count of all unique hosts detected by the system, regardless of whether the hosts have one IP address or multiple IP addresses.



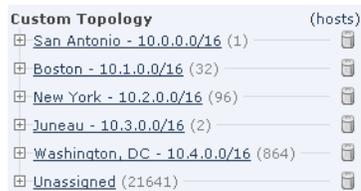
Although you can configure your network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for hosts added to the network map using NetFlow data, unless you provide it using the host input feature.

By creating a custom topology for your network, you can assign meaningful labels to your subnets, such as department names, that appear in the hosts network map.



You can also view the hosts network map according to the organization you specified in the custom topology; see [Working with Custom Network Topologies](#)

on page 1326.



You can delete entire networks, subnets, or individual hosts from the hosts network map. For example, if you know that a host is no longer attached to your network, you can delete it from the network map to simplify your analysis. If the system afterwards detects activity associated with the deleted host, it re-adds the host to the network map. If you want to permanently exclude a host or subnet from the network map, modify the network discovery policy. See [Creating a Network Discovery Policy](#) on page 1275 for more information.

IMPORTANT! Sourcefire **strongly** recommends that you do **not** delete network devices from the network map, because the system uses their locations to determine network topology (including generating network hops and TTL values for monitored hosts). Although you cannot delete network devices from the network devices network map, make sure you do not delete them from the hosts network map.

To view the hosts network map:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Hosts > Network Map**, then select the **Hosts** tab.

The hosts network map appears, displaying a host count and a list of host IP addresses and MAC addresses. Each address or partial address is a link to the next level.

2. Drill down to the specific IP address or MAC address of the host you want to investigate.

For example, to view a host with the IP address 192.168.40.11, click **192**, then **192.168**, then **192.168.40**, then **192.168.40.11**. When you click **192.168.40.11**, the host profile appears. For more information on host profiles, see [Using Host Profiles](#) on page 1334.

To filter by IP or MAC addresses, type an address in the search field. To clear the search, click the clear icon (✕).

3. Optionally, to delete a subnet, IP address, or MAC address, click the delete icon (🗑) next to the element you want to delete, then confirm that you want to delete the host or subnet.

The host is deleted. If the system rediscovers the host, it re-adds the host to the network map.

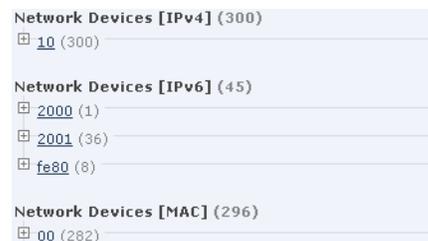
4. Optionally, switch between the hosts view and the topology view of the hosts network map:
 - To switch to a view of the hosts network map organized by your custom topology, on the hosts view (the default), click **(topology)** at the top of the network map.
 - To switch to a view of the hosts network map organized by subnet, on the topology view, click **(hosts)** at the top of the network map.

For information on configuring custom topologies, see [Working with Custom Network Topologies](#) on page 1326.

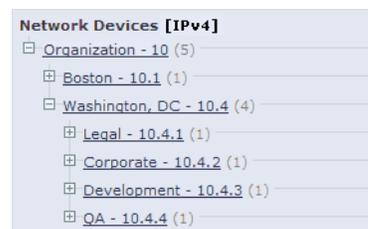
Working with the Network Devices Network Map

LICENSE: FireSIGHT

Use the network devices network map to view the network devices (bridges, routers, NAT devices, and load balancers) that connect one segment of your network to another, as well as to drill down to the host profiles of those network devices. The network devices network map is separated into two sections: IP and MAC. The IP section lists network devices identified by an IP address; the MAC section lists network devices identified by a MAC address. This network map view also provides a count of all unique network devices detected by the system, regardless of whether the devices have one IP address or multiple IP addresses.



If you create a custom topology for your network, the labels you assign to your subnets appear in the network devices network map.



The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their types (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge

- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a network device communicates using CDP, it may have one or more IP addresses. If it communicates using STP, it may only have a MAC address.

You cannot delete network devices from the network map, because the system uses their locations to determine network topology (including generating network hops and TTL values for monitored hosts).

The host profile for a network device has a Systems section rather than an Operating Systems section, which includes a Hardware column that reflects the hardware platform for any mobile devices detected behind the network device. If a value for a hardware platform is listed under Systems, that system represents a mobile device or devices detected behind the network device. Note that mobile devices may or may not have hardware platform information, but hardware platform information is never detected for systems that are not mobile devices.

To view the network devices network map:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Hosts > Network Map > Network Devices**.

The network devices network map appears, displaying a count of unique network devices and a list of network device IP addresses and MAC addresses. Each address or partial address is a link to the next level of addresses or to the host profile for an individual host.

2. Drill down to the specific IP address or MAC address of the network device you want to investigate.

The host profile for the network device appears. For more information on host profiles, see [Using Host Profiles](#) on page 1334.

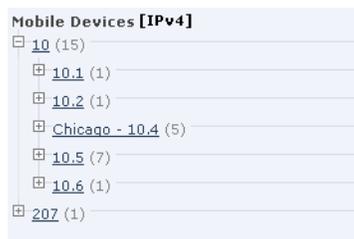
3. Optionally, to filter by IP or MAC addresses, type an address in the search field. To clear the search, click the clear icon (✕).

Working with the Mobile Devices Network Map

LICENSE: FireSIGHT

Use the mobile devices network map to view mobile devices attached to your network, and to drill down to the host profiles for those devices. This network map view also provides a count of all unique mobile devices detected by the

system, regardless of whether the devices have one IP address or multiple IP addresses.



The methods the system uses to distinguish mobile devices include:

- analysis of user agent strings in HTTP traffic from the mobile device's mobile browser
- monitoring of HTTP traffic of specific mobile applications

If you create a custom topology for your network, the labels you assign to your subnets appear in the mobile devices network map.

To view the mobile devices network map:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Hosts > Network Map**, then select the **Mobile Devices** tab.

The mobile devices network map appears, displaying a count of unique mobile devices and a list of mobile device IP addresses. Each address or partial address is a link to the next level.

2. Drill down to the specific IP address of the mobile device you want to investigate.

For example, to view a device with the IP address 10.11.40.11, click **10**, then **10.11**, then **10.11.40**, then **10.11.40.11**. When you click **10.11.40.11**, the host profile appears. For more information on host profiles, see [Using Host Profiles](#) on page 1334.

To filter by IP or MAC addresses, type an address in the search field. To clear the search, click the clear icon (✕).

3. Optionally, to delete a subnet or IP address, click the delete icon (🗑) next to the element you want to delete, then confirm that you want to delete the device or subnet.

The device is deleted. If the system rediscovers the device, it re-adds the device to the network map.

Working with the Applications Network Map

LICENSE: FireSIGHT

Use the applications network map to view the applications on your network, organized in a hierarchical tree by application name, vendor, version, and finally by the hosts running each application.



The applications that the system detects may change with system software and VDB updates, and also if you import any add-on detectors. The release notes or advisory text for each system or VDB update contains information on any new and updated detectors. For a comprehensive up-to-date list of detectors, see the Support Site.

From the applications network map, you can view the host profile of each host that runs a specific application as well as delete any application category, any application running on all hosts, or any application running on a specific host. For example, you can delete an application from the network map if you know it is disabled on the host and you want to make sure the system does not use it for impact flag qualification.

Deleting an application from the network map does not remove it from your network. A deleted application reappears in the network map if your system detects a change in the application (for example, if an Apache web server is upgraded to a new version) or if you restart your system's discovery function.

Depending on what you delete, the behavior differs:

- If you delete an application category, the application category is removed from the network map. All applications that reside beneath the category are removed from any host profile that contains the applications.

For example, if you delete **http**, all applications identified as **http** are removed from all host profiles and **http** no longer appears in the applications view of the network map.

- If you delete a specific application, vendor, or version, the affected application is removed from the network map and from any host profiles that contain it.
For example, if you expand the **http** category and delete **Apache**, all applications listed as Apache with any version listed beneath Apache are removed from any host profiles that contain them. Similarly, if instead of deleting **Apache**, you delete a specific version (**1.3.17** for example), only the version you selected will be deleted from affected host profiles.
- If you delete a specific IP address, the IP address is removed from the application list and the application itself is removed from the host profile of the IP address you selected.
For example, if you expand **http, Apache, 1.3.17 (Win32)**, and then delete **172.16.1.50:80/tcp**, the Apache 1.3.17 (Win32) application is deleted from the host profile of IP address 172.16.1.50.

To view the applications network map:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Hosts > Network Map > Applications**.
The applications network map appears.
2. Drill down to the specific application you want to investigate.
For example, if you want to view a specific type of web server like Apache, click **http**, then click **Apache**, and then click the version of the Apache web server you want to view.
To filter by IP or MAC addresses, type an address in the search field. To clear the search, click the clear icon (✕).
3. Click a specific IP address under the application you selected.
The host profile of the host running the application appears with the applications section expanded. For more information about the applications section of the host profile, see [Working with Servers in the Host Profile](#) on page 1347.

Servers (2) ▼

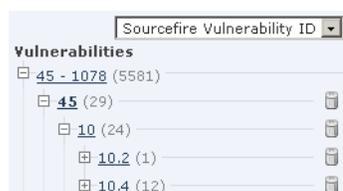
Protocol	Port	Application Protocol	Vendor and Version
tcp	8305	<input type="checkbox"/> SF MGMT	
tcp	22	<input type="checkbox"/> SSH	OpenSSH 4.3

4. Optionally, to delete any application category, any application running on all hosts, or any application running on a specific host, click the delete icon () next to the element you want to delete, then confirm that you want to delete it.
The application is deleted. If the system rediscovers the application, it is re-added to the network map.

Working with the Vulnerabilities Network Map

LICENSE: FireSIGHT

Use the vulnerabilities network map to view the vulnerabilities that the system has detected on your network, organized by Sourcefire vulnerability ID (SVID), Bugtraq ID, CVE ID, or Snort ID. The vulnerabilities are arranged by identification number, with affected hosts listed beneath each vulnerability.



From the vulnerabilities network map, you can view the details of specific vulnerabilities; you can also view the host profile of any host subject to a specific vulnerability. This can help you evaluate the threat posed by that vulnerability to specific affected hosts.

If you deem that a specific vulnerability is not applicable to the hosts on your network (for example, you have applied a patch), you can deactivate the vulnerability. Deactivated vulnerabilities still appear on the network map, but the IP addresses of their previously affected hosts appear in gray italics. The host profiles for those hosts show deactivated vulnerabilities as invalid, though you can manually mark them as valid for individual hosts; see [Setting Vulnerabilities for Individual Hosts](#) on page 1368 for more information.

If there is an identity conflict for an application or operating system on a host, the system lists the vulnerabilities for both potential identities. When the identity conflict is resolved, the vulnerabilities remain associated with the current identity. For more information, see [Understanding Current Identities](#) on page 1646 and [Understanding Identity Conflicts](#) on page 1648.

By default, the vulnerability network map displays the vulnerabilities of a detected application only if the packet contains the application's vendor and version. However, you can configure the system to list the vulnerabilities for applications lacking vendor and version data by enabling the vulnerability mapping setting for the application in the system policy. For information on setting the vulnerability mapping for an application, see [Mapping Vulnerabilities for Servers](#) on page 1996.

The numbers next to a vulnerability ID (or range of vulnerability IDs) represent two counts:

- The first number is a count of non-unique hosts that are affected by a vulnerability or vulnerabilities. If a host is affected by more than one vulnerability, it is counted multiple times. Therefore, it is possible for the count to be higher than the number of hosts on your network. Deactivating

a vulnerability decrements this count by the number of hosts that are potentially affected by the vulnerability. If you have not deactivated any vulnerabilities for any of the potentially affected hosts for a vulnerability or range of vulnerabilities, this count is not displayed.

- The second number is a similar count of the total number of non-unique hosts that the system has determined are *potentially* affected by a vulnerability or vulnerabilities.

Deactivating a vulnerability renders it inactive only for the hosts you designate. You can deactivate a vulnerability for all hosts that have been judged vulnerable or for a specified individual vulnerable host. If the system subsequently detects the vulnerability on a host where it has not been deactivated (for example, on a new host in the network map), the system activates the vulnerability for that host. You have to explicitly deactivate the newly discovered vulnerability. Also, if the system detects an operating system or application change for a host, it may reactivate associated deactivated vulnerabilities.

To view the vulnerabilities network map:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Hosts > Network Map > Vulnerabilities**.

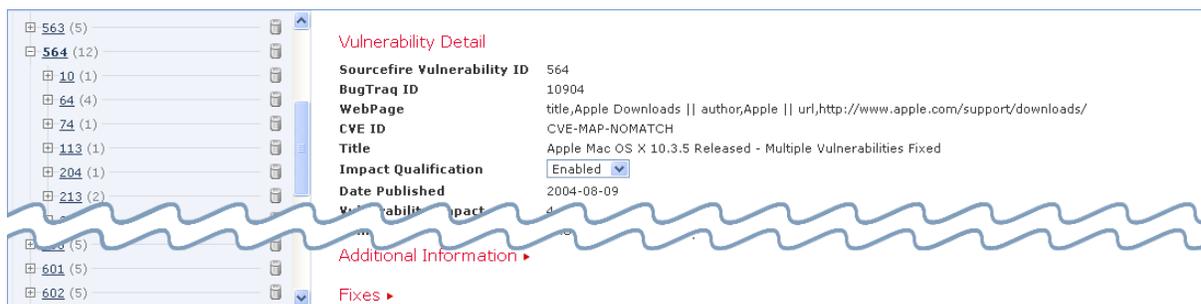
The vulnerabilities network map appears.

2. From the **Type** drop-down list, select the class of vulnerability you want to view. By default, vulnerabilities are displayed by Sourcefire vulnerability ID (SVID).

3. Drill down to the specific vulnerability you want to investigate.

To filter by IP or MAC addresses, type an address in the search field. To clear the search, click the clear icon (**X**).

The vulnerability details appear. For details on the information provided, see [Viewing Vulnerability Details](#) on page 1365.



In addition, on the network map, the Defense Center displays the IP addresses of affected hosts. You can click any IP address to display the host profile for that host.

4. Optionally, deactivate the vulnerability:
 - To deactivate the vulnerability for all hosts affected by the vulnerability, click the delete icon () next to the vulnerability number.
 - To deactivate the vulnerability for an individual host, click the delete icon () next to the host's IP address.

The vulnerability is deactivated. The applicable hosts' IP addresses appear in gray italics in the network map. In addition, host profiles for those hosts show deactivated vulnerabilities as invalid.

TIP! See [Setting Vulnerabilities for Individual Hosts](#) on page 1368 for more information on reactivating vulnerabilities.

Working with the Host Attributes Network Map

LICENSE: FireSIGHT

Use the host attributes network map to view the hosts on your network organized by their host attributes. When you select the host attribute you want to use to organize your hosts, the Defense Center lists the possible values for that attribute in the network map and groups hosts based on their assigned values. You can also view the host profile of any host assigned a specific host attribute value.

The host attributes network map can organize hosts based on user-defined host attributes. For any of these attributes, the network map displays hosts that do not have a value assigned as Unassigned.



For more information, see [Working with User-Defined Host Attributes](#) on page 1370.

In addition, the host attributes network map can organize hosts based on the host attributes that correspond to any compliance white lists you have created. Each compliance white list that you create automatically creates a host attribute with the same name as the white list.



Possible white list host attribute values are:

- **Compliant**, for hosts that are compliant with the white list
- **Non-Compliant**, for hosts that violate the white list
- **Not Evaluated**, for hosts that are not valid targets of the white list or have not been evaluated for any reason

For more information on compliance white lists, see [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530.

IMPORTANT! You cannot organize hosts using predefined host attributes, such as host criticality, on the host attributes network map.

To view the host attributes network map:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Hosts > Network Map > Host Attributes**.

The host attributes network map appears.

2. From the **Attribute** drop-down list, select a host attribute.

The Defense Center lists the values for the host attribute and indicates, in parentheses, the number of hosts assigned that value.

To filter by IP or MAC addresses, type an address in the search field. To clear the search, click the clear icon (✕).

3. Click any host attribute value to view hosts assigned the value.
4. Click a host IP address to view the host profile for that host.

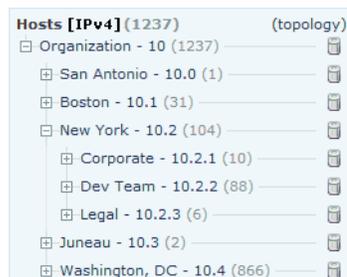
Working with Custom Network Topologies

LICENSE: FireSIGHT

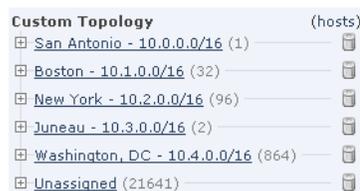
Use the custom topology feature to help you organize and identify subnets in your hosts and network devices network maps.

For example, if each department within your organization uses a different subnet, you can label those subnets using the custom topology feature. Then, when you

view the hosts or network devices network map, the labels you assign to your subnets appear, as shown in the following graphic.



You can also view the hosts network map according to the organization you specified in the custom topology.



For more information about the hosts and network devices network maps, see [Working with the Hosts Network Map](#) on page 1316 and [Working with the Network Devices Network Map](#) on page 1318.

For more information, see the following sections:

- [Creating Custom Topologies](#) on page 1327
- [Managing Custom Topologies](#) on page 1332

Creating Custom Topologies

LICENSE: FireSIGHT

To create a custom topology, you must specify its networks. You can do this using any or all of three strategies:

- by importing the Sourcefire-discovered topology, which adds networks using a “best guess” at how your network is deployed based on the hosts and network devices the system has detected
- by importing networks from a network discovery policy, which adds the networks that you configured the Sourcefire 3D System to monitor in a network discovery policy
- by adding networks to your topology manually, if the other two methods create an inaccurate or incomplete representation of your deployment

You must save and activate the topology before using it with the network map.

To create a custom topology:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Topology**.
The Custom Topology page appears.
2. Click **Create Topology**.
The Create Topology page appears.
3. Provide basic topology information, such as the topology name and description.
See [Providing Basic Topology Information](#) on page 1329.
4. Add networks to your topology. You can use any or all of the following strategies:
 - To add networks to your topology by importing the Sourcefire-discovered topology, follow the procedure in [Importing a Discovered Topology](#) on page 1329.
 - To add networks to your topology by importing them from a network discovery policy, follow the procedure in see [Importing Networks from a Network Discovery Policy](#) on page 1330.
 - To add networks to your topology manually, follow the procedure in [Manually Adding Networks to Your Custom Topology](#) on page 1331.
5. Refine your topology:
 - To remove a network from your custom topology, click **Delete** next to the network you want to remove.
 - To rename a network, click **Rename** next to the network. In the pop-up window that appears, type the new name in the **Name** field and click **Rename**. This name labels the network in the network map.
6. Click **Save**.
The topology is saved.

IMPORTANT! You must activate the topology before you can use it in the network map. For more information, see [Managing Custom Topologies](#) on page 1332.

Providing Basic Topology Information

LICENSE: FireSIGHT

You must give each custom topology a name, and, optionally, a short description.



To provide basic topology information:

ACCESS: Admin

1. On the Edit Topology page, in the **Name** field, type a name for the topology.
2. Optionally, in the **Description** field, type a description for the topology.
3. Optionally, continue with the procedures in the following sections, depending on how you want to build your custom topology:
 - [Importing a Discovered Topology](#) on page 1329
 - [Importing Networks from a Network Discovery Policy](#) on page 1330
 - [Manually Adding Networks to Your Custom Topology](#) on page 1331

Importing a Discovered Topology

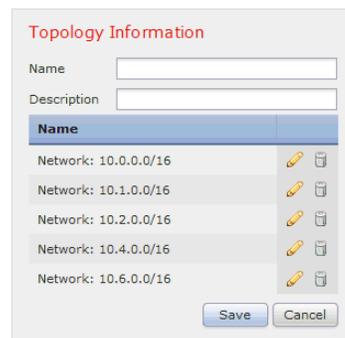
LICENSE: FireSIGHT

One way you can add networks to your custom topology is to import the topology discovered by your Sourcefire 3D System. This discovered topology is the system's "best guess" at how your network is deployed based on the hosts and network devices it has detected.

To import a discovered topology:

ACCESS: Admin

1. On the Edit Topology page, click **Import Discovered Topology**.
2. The discovered networks populate the page.



Name	
Network: 10.0.0.0/16	 
Network: 10.1.0.0/16	 
Network: 10.2.0.0/16	 
Network: 10.4.0.0/16	 
Network: 10.6.0.0/16	 

3. Optionally, continue with the procedures in the following sections, depending on how you want to build your custom topology:
 - [Importing a Discovered Topology](#) on page 1329
 - [Importing Networks from a Network Discovery Policy](#) on page 1330
 - [Manually Adding Networks to Your Custom Topology](#) on page 1331

Importing Networks from a Network Discovery Policy

LICENSE: FireSIGHT

One way you can add networks to your custom topology is to import the networks that you configured the Sourcefire 3D System to monitor in a network discovery policy. For information on discovery policies, see [Creating a Network Discovery Policy](#) on page 1275.

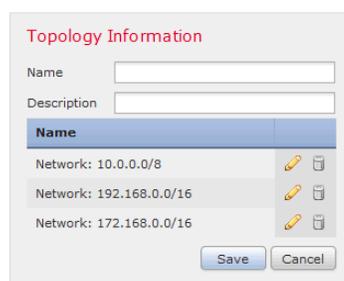
To import networks from a network discovery policy:

ACCESS: Admin

1. On the Edit Topology page, click **Import Policy Networks**.
A pop-up window appears.



2. From the drop-down list, choose the network discovery policy you want to use and click **Load**.
3. The monitored networks in your network discovery policy populate the page.
For example, if you configured your network discovery policy to monitor the 10.0.0.0/8, 192.168.0.0/16, and 172.12.0.0/16 networks, those networks appear on the page.



4. To add networks from a different network discovery policy, repeat steps 1 and 2.

5. Optionally, follow the procedures in the following sections, depending on how you want to build your custom topology:
 - [Importing a Discovered Topology](#) on page 1329
 - [Manually Adding Networks to Your Custom Topology](#) on page 1331

Manually Adding Networks to Your Custom Topology

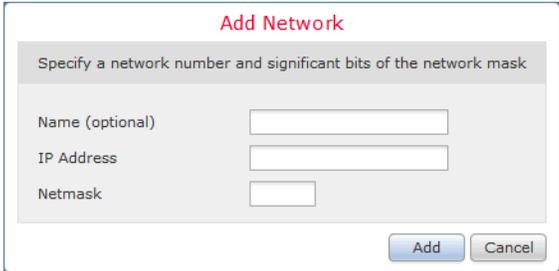
LICENSE: FireSIGHT

If importing the Sourcefire-discovered topology and importing networks from your network discovery policy creates an inaccurate or incomplete representation of your network deployment, you can add networks to your custom topology manually.

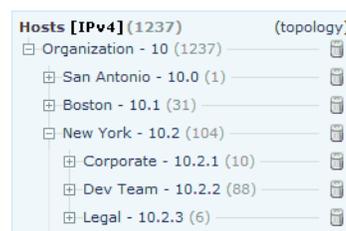
To add a network to a custom topology manually:

ACCESS: Admin

1. On the Edit Topology page, click **Add Network**.
A pop-up window appears.



2. Optionally, name the network by typing a name in the **Name** field.
This name labels the networks in the hosts and network devices network maps after you activate the topology.



For more information, see [Working with the Hosts Network Map](#) on page 1316 and [Working with the Network Devices Network Map](#) on page 1318.

3. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in CIDR notation) that represent the network you want to add to your topology.

For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

4. Click **Add**.
The network is added to your topology.
5. To add additional networks to your topology, repeat steps 1 through 4.

TIP! To delete a network from your topology, click **Delete** next to the network you want to delete, and confirm that you want to delete the network, as well as all links to the network.

6. Optionally, follow the procedures in the following sections, depending on how you want to build your custom topology:
 - [Importing a Discovered Topology](#) on page 1329
 - [Importing Networks from a Network Discovery Policy](#) on page 1330

Managing Custom Topologies

LICENSE: FireSIGHT

Use the Custom Topology page to manage custom topologies. You can create, modify, and delete topologies.

Topologies

Name	
Organization Topology	<input checked="" type="checkbox"/>  
Sample Topology	<input type="checkbox"/>  

A topology's status appears with its name. If the light bulb icon next to the policy name is lit, the topology is active and affects your network map. If it is dark, the topology is inactive. Only one custom topology can be active at any time. If you have created multiple topologies, activating one automatically deactivates the currently active topology.

Use the following procedures to either activate or deactivate a custom topology, modify a topology, or delete a topology.

If you delete the active topology, your changes take effect immediately; that is, your network map no longer displays your custom topology.

To activate or deactivate a custom topology:

ACCESS: Admin

1. Select **Policies > Network Discovery > Custom Topology**.
The Custom Topology page appears.
2. You have two options:
 - To **activate** a topology, click **Activate** next to the policy.
 - To **deactivate** a topology, click **Deactivate** next to the policy.

To modify a custom topology:

ACCESS: Admin

1. Select **Policies > Network Discovery > Custom Topology**.

The Custom Topology page appears.

2. Click the edit icon () next to the topology you want to edit.

The Edit Topology page appears. See [Creating Custom Topologies](#) on page 1327 for information on the various configurations you can change.

3. Make changes as needed and click **Save**.

The topology is changed. If the topology is active, the changes you made take effect in the network map immediately.

To delete a custom topology:

ACCESS: Admin

1. Select **Policies > Network Discovery > Custom Topology**.

The Custom Topology page appears.

2. Click **Delete** next to the topology you want to delete. If the topology is active, confirm that you want to delete it.

The topology is deleted.

CHAPTER 34

USING HOST PROFILES

A host profile provides a complete view of all the information the system has gathered about a single host. You can access general host information, such as the host name and operating system, through the profile. If you need to quickly find the MAC address for a host, for example, you can look in the host profile.

Host attributes for that host are also listed in the profile. Host attributes are user-defined descriptions that you can apply to a host. For example, you might assign a host attribute that indicates the building where the host is located. From a host profile, you can view the existing host attributes applied to that host and can modify the host attribute values. As another example, you can use the *host criticality* attribute to designate the business criticality of a given host and to tailor correlation policies and alerts based on host criticality.

Host profiles also provide you with information about the servers, clients, and host protocols running on a particular host, including whether they are in compliance with a compliance white list. You can remove servers from the servers list, and view details for those servers. You can also view *connection events* for servers, log information about the session where server traffic was detected. You can also view details and connection events for clients and delete servers, clients or host protocols from the host profile.

If your Sourcefire 3D System deployment includes a Protection license, you can tailor the way the system processes traffic so it best fits the type of operating system on the host and the servers and clients the host is running. For more information, see [Using Adaptive Profiles](#) on page 1003.

You can also see user history information for a host if you have configured the system to track it. A graphic representation of the last twenty-four hours of user activity is then available.

You can modify the list of vulnerabilities for the host from the host profile. You can use this capability to track which vulnerabilities have been addressed for the host. You can also apply fixes for vulnerabilities, causing all vulnerabilities addressed by the fix to be automatically marked invalid.

You can work with the vulnerability information generated by the Sourcefire system, and also use information on vulnerabilities detected by third-party scanners, which you import onto the Defense Center using the host input feature.

Optionally, you can perform an Nmap scan from the host profile, to augment the server and operating system information in your host profile. The Nmap scanner actively probes the host to obtain information about the operating system and servers running on the host. The results of the scan are added to the list of operating system and server identities for the host.

Note that a host profile may not be available for every host on your network. Possible reasons include:

- the host was deleted from the network map because it timed out
- you have reached your FireSIGHT host license limit
- the host resides in a network segment that is not monitored by your network discovery policy

Note that the information displayed in a host profile may vary according to the type of host and the information available about the host. For example, if your system detects a host using a non-IP-based protocol like STP, SNAP, or IPX, the host is added to the network map as a MAC host and much less information is available than for an IP host.

As another example, although you can configure your network discovery policy to add hosts and server and clients to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts and servers and clients is limited. For example, no operating system data is available for these hosts, unless you provide it using a scanner or the host input feature. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

The following graphic shows an example of a host profile.

Host Profile Scan Host Generate White List Profile

IP Addresses fe80::200:5aee:feaa:20a2

NetBIOS Name

Device (Hops) linden (31)

MAC Addresses (TTL) 00:00:00:00:00:00 (Dell Inc.) (64)

Host Type Host

Last Seen 2013-02-20 16:47:59

Current User

View Context Explorer
Discovery Events
Malware Events
Intrusion Events by Source
Intrusion Events by Destination

Operating System (pending)

Users (no user history available)

Attributes

Host Protocols

The following graphic shows an example of a host profile for a MAC host.

Host Profile

IP Addresses

NetBIOS Name

Device (Hops)

MAC Addresses (TTL) 00:11:22:33:44:55

Host Type Host

Last Seen 2013-04-16 18:51:28

Operating System (pending)

Users (no user history available)

Attributes

Host Criticality None

Host Protocols

Protocol	Layer
ARP	Network

For more information about each section of the host profile, see the following:

- [Viewing Host Profiles](#) on page 1337 explains how to access a host profile.
- [Working with Basic Host Information in the Host Profile](#) on page 1338 describes the information provided in the Host section of a host profile.
- [Working with IP Addresses in the Host Profile](#) on page 1341 describes the information provided in the IP Addresses section of a host profile.

- [Working with Operating Systems in the Host Profile](#) on page 1341 describes the information provided in the Operating System or Operating System Conflicts section of a host profile and explains how to edit the operating system or resolve an operating system conflict.
- [Working with Servers in the Host Profile](#) on page 1347 describes the information provided in the Servers, Server Detail, and Server Banner sections of a host profile.
- [Working with Applications in the Host Profile](#) on page 1354 describes the information provided in the Clients section of a host profile.
- [Working with VLAN Tags in the Host Profile](#) on page 1357 describes the information provided in the VLAN Tag section of a host profile.
- [Working with User History in the Host Profile](#) on page 1357 describes the information provided in the User History section of a host profile.
- [Working with Host Attributes in the Host Profile](#) on page 1358 describes the information provided in the Attributes section of a host profile.
- [Working with the Predefined Host Attributes](#) on page 1369 explains how to set the host criticality attribute and how to add notes to a host profile.
- [Working with User-Defined Host Attributes](#) on page 1370, which provides information about creating and using user-defined host attributes.
- [Working with Host Protocols in the Host Profile](#) on page 1359 describes the information provided in the Host Protocols section of a host profile.
- [Working with White List Violations in the Host Profile](#) on page 1360 describes the information provided in the White List Violations section of a host profile.
- [Working with Malware Detections in the Host Profile](#) on page 1362 describes the information provided in the Most Recent Malware Detections section of a host profile.
- [Working with Vulnerabilities in the Host Profile](#) on page 1363, which describes the information provided in the Vulnerabilities and Vulnerability Detail sections of a host profile.

Viewing Host Profiles

LICENSE: FireSIGHT

You can access a host profile from any network map or from any event view that includes the IP addresses of hosts on monitored networks. For example, the table view of discovery events includes a link to the host profile next to every entry in the IP Address column.

To view a host profile from an event view:**ACCESS:** Admin/Any Security Analyst

- ▶ On any event view, click the host profile icon (🏠) next to the IP address of the host whose profile you want to view.
The host profile appears in a pop-up window.

To view a host profile from a network map:**ACCESS:** Admin/Any Security Analyst

- ▶ On any network map, drill down to the IP address of the host whose profile you want to view.
The host profile appears. See [Working with the Hosts Network Map](#) on page 1316 for an example of how to access a host profile from a network map.

Working with Basic Host Information in the Host Profile

LICENSE: FireSIGHT

Each host profile provides basic information about a detected host or other device.

Host Profile Scan Host Generate White List Profile

IP Addresses	10.0.0.0 🇺🇸 2001:cdba:0000:0000:0000:3257:9652
NetBIOS Name	
Device (Hops)	linden (0)
MAC Addresses (TTL)	00:00:00:00:00:00 (Intel Corporate) (255)
Host Type	Host
Last Seen	2013-02-22 14:59:42
Current User	
View	Context Explorer Discovery Events Malware Events Intrusion Events by Source Intrusion Events by Destination

Descriptions of each of the basic host profile fields follow.

IP Addresses

All IP addresses (both IPv4 and IPv6) associated with the host. IPv6 hosts often have at least two IPv6 addresses (local-only and globally routable), and may also have IPv4 addresses. IPv4-only hosts may have multiple IPv4 addresses. Where available, routable host IP addresses also include a flag icon and country code indicating the geolocation data associated with that address. For more information on this and other geolocation features, see [Using Geolocation](#) on page 1816.

Hostname

The fully qualified domain name of the host, if known.

NetBIOS Name

The NetBIOS name of the host, if available. Microsoft Windows hosts, as well as Macintosh, Linux, or other platforms configured to use NetBIOS, can have a NetBIOS name. For example, Linux hosts configured as Samba servers have NetBIOS names.

Device (Hops)

Either:

- the reporting device for the network where the host resides, as defined in the network discovery policy, or
- the device that processed the NetFlow data that added the host to the network map

The device and the number of network hops between the device that detected the host and the host itself follows the device name, in parentheses. If multiple devices can see the host, the reporting device is displayed in bold.

If this field is blank, either:

- the host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy, or
- the host was added using the host input feature and has not also been detected by the Sourcefire 3D System

MAC Addresses (TTL)

The host's detected MAC address or addresses and associated NIC vendors, with the NIC's hardware vendor and current time-to-live (TTL) value in parentheses. If the MAC address is displayed in a bold font, the MAC address is the actual MAC address of the host, detected by the system through ARP and DHCP traffic. If multiple devices detected the host, the Defense Center displays all MAC addresses and TTL values associated with the host, regardless of which device reported them.

You can click the MAC address to view a list of hosts with the same MAC address. Router host profiles typically show the hosts (IP addresses) in the network segments they route in this list, and the IP addresses of monitored routers frequently appear in this list for monitored workstations and servers. The true IP address for the MAC address is displayed in bold.

Host Type

The type of device that the system detected: host, mobile device, jailbroken mobile device, router, bridge, NAT device, or load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

The methods the system uses to distinguish mobile devices include:

- analysis of user agent strings in HTTP traffic from the mobile device's mobile browser
- monitoring of HTTP traffic of specific mobile applications

If a device is not identified as a network device or a mobile device, it is categorized as a host.

Last Seen

The date and time that any of a host's IP addresses was last detected.

Current User

The user most recently logged into this host.

Note that a non-authoritative user logging into a host only registers as the current user on the host if the existing current user is not an authoritative user. For more information, see [Users Database](#) on page 1258.

View

Links to views of event data, using the default workflow for that event type and constrained to show events related to the host; where possible, these events include all IP addresses associated with the host. For more information, see the following sections:

- Content Explorer — for more information, see [Using the Context Explorer](#) on page 117.
- Discovery Events — for more information, see [Working with Discovery Events](#) on page 1377.
- Malware Events — for more information, see [Working with Malware Events](#) on page 1229.
- Intrusion Events by Source — for more information, see [Working with Intrusion Events](#) on page 593.
- Intrusion Events by Destination — for more information, see [Working with Intrusion Events](#) on page 593.

Working with IP Addresses in the Host Profile

LICENSE: FireSIGHT

The system detects IP addresses associated with hosts and, where supported, groups multiple IP addresses used by the same host. IPv6 hosts usually have at least two IPv6 addresses: local-only and globally routable. They may also have one or more assigned IPv4 addresses. IPv4-only hosts may have multiple IPv4 addresses.

The host profile lists all detected IP addresses associated with that host. Where available, IP addresses also feature a small flag icon and ISO country code that indicate the associated country. You can click the flag icon or country code for further geolocation details. For more information, see [Using Geolocation](#) on page 1816.

Working with Operating Systems in the Host Profile

LICENSE: FireSIGHT

The system passively detects the identity of the operating system running on a host by analyzing the network and application stack in traffic generated by the host or by analyzing host data reported by the Sourcefire User Agent. The system also collates operating system information from other sources, such as the Nmap scanner or application data imported through the host input feature. The system considers the priority assigned to each identity source when determining which identity to use. By default, user input has the highest priority, followed by application or scanner sources, followed by the Sourcefire-discovered identity.

Sometimes the system supplies a general operating system definition rather than a specific one because the traffic and other identity sources do not provide sufficient information for a more focused identity. The system collates information from the sources to use the most detailed definition possible.

Descriptions of the operating system information fields displayed in the host profile follow.

Hardware

The hardware platform for a mobile device.

OS Vendor/Vendor

The operating system vendor.

OS Product/Product

The operating system determined most likely to be running on the host, based on the identity data collected from all sources.

If the operating system is **Pending**, the system has not yet identified an operating system and no other identity data is available. If the operating system is **unknown**, the system cannot identify the operating system and no other identity data is available for the operating system.

If the host's operating system is not one the system is capable of detecting, you may want to use one of the following strategies:

- create a custom fingerprint for the host, as described in [Using Custom Fingerprinting](#) on page 1649
- run an Nmap scan against the host, as described in [Scanning a Host from the Host Profile](#) on page 1376
- import data into the network map, using the host input feature described in the *Sourcefire 3D System Host Input API Guide*
- manually enter operating system information, as described in [Working with Operating Systems in the Host Profile](#) on page 1341

OS Version/Version

The operating system version. If a host is a jailbroken mobile device, **Jailbroken** is indicated in parentheses after the version.

Source

One of the following values:

- User: *user_name*
- Application: *app_name*
- Scanner: *scanner_type* (Nmap or scanner added through system policy)
- FireSIGHT

The system may reconcile data from multiple sources to determine the identity of an operating system; see [Understanding Current Identities](#) on page 1646.

Because the vulnerabilities list for the host and the event impact correlation for events targeting the host depend on the operating system, you may want to manually supply more specific operating system information. In addition, you can indicate that fixes have been applied to the operating system, such as service packs and updates, and invalidate any vulnerabilities addressed by the fixes.

For example, if the system identifies a host's operating system as Microsoft Windows 2003, but you know that the host is actually running Microsoft Windows XP Professional with Service Pack 2, you can set the operating system identity accordingly. Setting a more specific operating system identity refines the list of vulnerabilities for the host, so your impact correlation for that host is more focused and accurate.

If the system detects operating system information for a host and that information conflicts with a current operating system identity that was supplied

by an active source, an identity conflict occurs. When an identity conflict is in effect, the system uses both identities for vulnerabilities and impact correlation.

Although you can configure the network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, there is no operating system data available for these hosts, unless you set the operating system identity. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Note that if a host is running an operating system that violates a compliance white list in an activated network discovery policy, the Defense Center marks the operating system information with the white list violation icon (ⓘ). In addition, if a jailbroken mobile device violates an active white list, the icon appears next to the operating system for the device.

Systems (1) ▾ 

Hardware	OS Vendor	OS Product	OS Version	Source
ⓘ	Ubuntu	Linux	11.x, 12.x	FireSIGHT

You can set a custom display string for the host's operating system identity. That display string is then used in the host profile.

IMPORTANT! Note that changing the operating system information for a host may change its compliance with a compliance white list.

In the host profile for a network device, the label for the Operating Systems section changes to Systems and an additional Hardware column appears. If a value for a hardware platform is listed under Systems, that system represents a mobile device or devices detected behind the network device. Note that mobile devices may or may not have hardware platform information, but hardware platform information is never detected for systems that are not mobile devices.

Systems (5) ▾  

Hardware	OS Vendor	OS Product	OS Version
	Linux	Linux	2.6
iPhone	Apple	iOS	1.x, 1.0.1, 1.0.2, 1.1.0, 1.1.1, 1.1.2, 1.1.3, 1.1.4, 2.0.0, 2.0.1, 2.0.2, 2.1.0, 2.2.0, 2.2.1, 3.0.0, 3.0.1, 3.1, 4.2.1, 5.x
	Microsoft	Windows	2000, XP, Server 2003
	Apple	Mac OSX	10.5, 10.6, Server 10.5, Server 10.6
	Google	Android	2.2.1, 2.3.4, 4.0.3

Viewing Operating System Identities

LICENSE: FireSIGHT

You can view the specific operating system identities discovered or added for a host. The system uses source prioritization to determine the current identity for

the host. In the list of identities, the current identity is highlighted by boldface text.

For each operating system identity, the host profile may include the information described in [Working with Operating Systems in the Host Profile](#) on page 1341.

Note that the View button is only available if multiple operating system identities exist for the host.

To view the list of operating system identities for a host:

ACCESS: Admin/Any Security Analyst

- ▶ Click **View** in the **Operating System** or **Operating System Conflicts** section of the host profile.

The Operating System Identity Information pop-up window appears.

TIP! Click the delete icon () next to any operating system identity to remove the identity from the Operating System Identity Information pop-up window and, if applicable, to update the current identity for the operating system in the host profile. Note that you cannot delete Sourcefire-detected operating system identities.

Operating System Identity Information

Vendor	Product	Version	Source	Confidence	Created On
Microsoft, Corp.	Windows 7	Professional	User: admin	100	2011-10-14 11:46:18
Apple	Mac OS X	10.5, 10.6, Server 10.5, Server 10.6	RNA	98	2011-10-14 11:44:53

Editing an Operating System

LICENSE: FireSIGHT

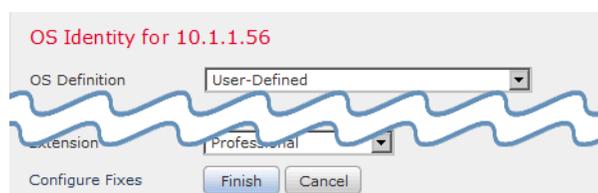
You can set the current operating system identity for a host using the Sourcefire 3D System web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation. However, note that if the system detects a conflicting operating system identity for the host after you edit the operating system, an operating system conflict occurs.

Both operating systems are then considered current until you resolve the conflict. For more information, see [Resolving Operating System Identity Conflicts](#) on page 1345.

To change the operating system identity:

ACCESS: Admin/Any Security Analyst

1. In a host profile, click **Edit** in the **Operating System** section.
A pop-up window appears where you can set the operating system identity.



2. You have several options:
 - Select **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
 - Select a variation on the current operating system identity from the **OS Definition** drop-down list, then skip to step 6.
 - Select **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.
3. Optionally, select **Use Custom Display String** and modify the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.
4. Optionally, to change to an operating system from a different vendor, select the vendor and other operating system details from the **Vendor** and **Product** drop-down lists.
5. Optionally, to configure the operating system product release level, select the applicable items in the **Major**, **Minor**, **Revision**, **Build**, **Patch**, and **Extension** drop-down lists.
6. Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.
The list of available package fixes displays.
7. Choose the applicable fixes in the drop-down list and click **Add**.
8. Optionally, add the relevant patches and extensions using the **Patch** and **Extension** drop-down lists.
9. Click **Finish** to complete the operating system identity configuration.

Resolving Operating System Identity Conflicts

LICENSE: FireSIGHT

An operating system identity conflict occurs when a new identity detected by the system conflicts with the current identity, if that identity was provided by an active source, such as a scanner, application, or user.

The list of operating system identities in conflict displays in bold in the host profile.

You can resolve an identity conflict and set the current operating system identity for a host through the system web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation.

To make one of the conflicting identities current:

ACCESS: Admin/Any Security Analyst

► You have two options:

- Click **Make Current** next to the operating system identity you want to set as the operating system for the host.
- If the identity that you *do not* want as the current identity came from an active source, delete the unwanted identity.

To resolve an operating system identity conflict:

ACCESS: Admin/Any Security Analyst

1. In a host profile, click **Resolve** in the **Operating System Conflicts** section.

A pop-up window appears where you can set the current operating system identity.



2. You have several options:

- Select **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Select a variation on one of the conflicting operating system identities from the **OS Definition** drop-down list, then skip to step 6.
- Select **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

3. Optionally, select **Use Custom Display String** and type the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

4. Optionally, to change to an operating system from a different vendor, select the vendor and other operating system details.

5. Optionally, to configure the operating system product release level, select the applicable items in the **Major**, **Minor**, **Revision**, **Build**, **Patch**, and **Extension** drop-down lists.

6. Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.
7. Add the fixes you have applied to the fixes list.
8. Click **Finish** to complete the operating system identity configuration and return to the host profile.

Working with Servers in the Host Profile

LICENSE: FireSIGHT

If the system detects servers running on a host on your monitored network or if servers are added through the host input feature or through a scanner or other active source, the Defense Center lists them in the Servers section of the host profile.

Servers (2) ▾

Protocol	Port	Application Protocol	Vendor and Version	
tcp	8305	SF MGMT		   
tcp	22	SSH	OpenSSH 4.3	   

The Defense Center lists up to 100 servers per host. After that limit is reached, new server information from any source, whether active or passive, is discarded until you delete a server from the host or a server times out. For more information, see [Host Limits and Discovery Event Logging](#) on page 1268.

If you scan a host using Nmap, Nmap adds the results of previously undetected servers running on open TCP ports to the Servers list. If you perform an Nmap scan on a host or import Nmap results, an expandable Scan Results section also appears in the host profile, listing the server information detected on the host by the Nmap scan. See [Working with Scan Results in a Host Profile](#) on page 1375 and [Setting up Nmap Scans](#) on page 1702 for more information. In addition, note that if the host is deleted from the network map, the Nmap scan results for that server for the host are discarded.

IMPORTANT! Although you can configure your network discovery policy to add server and clients to the network map based on data exported by NetFlow-enabled devices, the available information about these applications is limited. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

The process for working with servers in the host profile differs depending on how you accessed the profile:

- If you accessed the host profile by drilling down through the Servers network map, the details for that server appear with the server name highlighted in bold. If you want to view the details for any other server on the host, click the view icon (🔍) next to that server name.
- If you accessed the host profile in any other way, expand the Servers section and click the view icon (🔍) next to the server whose details you want to see.

You can also perform the following actions:

- To analyze the connection events associated with a particular server on the host, click the events icon next to the server.
The first page of your preferred workflow for connection events appears, showing connection events constrained by the port and protocol of the server, as well as the IP address of the host. If you do not have a preferred workflow for connection events, you must select one. For more information about connection data, see [Working With Connection Data](#) on page 539.
- To delete a server from the host profile, click the delete icon (🗑️) next to the server.
The server is deleted from the host profile, but will appear again if the system detects traffic from the server again. Note that deleting a server from a host may bring the host into compliance with a white list.
- To resolve a server identity conflict, click the resolve icon next to the server.
You can choose one of the conflicting identities, choose a variation on one of those identities, or set a new user-defined identity.
- To edit a server identity, click the edit icon (✏️) next to the server.
You can choose the current identity, choose a variation on that identity, or set a new user-defined identity.

Descriptions of the columns in the Servers list follow.

Protocol

The name of the protocol the server uses.

Port

The port where the server runs.

Application Protocol

One of:

- the name of the application protocol
- **pending**, if the system cannot positively or negatively identify the application protocol for one of several reasons
- **unknown**, if the system cannot identify the application protocol based on known application protocol fingerprints or if the server was added through host input by adding a vulnerability with port information without adding a corresponding server

When you hover the mouse on an application protocol name, the tags display. For information on tags, see [Understanding Application Detection](#) on page 1263.

Vendor and Version

The vendor and version identified by the Sourcefire 3D System, by Nmap, or by another active source, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Note that if the host is running a server that violates a compliance white list in an activated correlation policy, the Defense Center marks the non-compliant server with the white list violation icon (ⓘ).



Protocol	Port	Application Protocol	Vendor and Version
ⓘ udp	123	NTP	

See the following sections for more information:

- [Server Detail](#) on page 1349
- [Editing Server Identities](#) on page 1352
- [Resolving Server Identity Conflicts](#) on page 1353

Server Detail

LICENSE: FireSIGHT

The Defense Center lists up to 16 passively detected (Sourcefire- or NetFlow-detected) identities per server. A server can have multiple passive identities if the system detects multiple vendors or versions of that server. For example, a load balancer between your managed device and your web server farm may cause your system to identify multiple passive identities for HTTP if your web servers are not running the same version of the server software. Note that the Defense Center does not limit the number of server identities from active sources such as user input, scanners, or other applications.

The Defense Center displays the current identity in bold. The system uses the current identity of a server for multiple purposes, including assigning

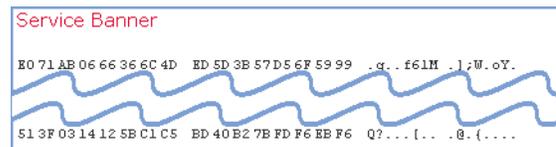
vulnerabilities to a host, impact assessment, evaluating correlation rules written against host profile qualifications and compliance white lists, and so on.

TIP! For information on changing the server identity and resolving identity conflicts from the server detail, see [Editing Server Identities](#) on page 1352 and [Resolving Server Identity Conflicts](#) on page 1353.

The server detail may also display updated sub-server information known about the selected server. The following graphic displays sub-server information for a detected application.

Application Protocol	Vendor	Version
DAV		2
mod_ssl		2.2.19
OpenSSL		0.9.8q

Finally, the server detail may display the server banner, which appears below the server details when you view a server from the host profile.



Server banners provide additional information about a server that may help you identify the server. The system cannot identify or detect a misidentified server when an attacker purposely alters the server banner string. The server banner displays the first 256 bytes of the first packet detected for the server. It is collected only once, the first time the server is detected by the system. Banner content is listed in two columns, with a hexadecimal representation on the left and a corresponding ASCII representation on the right.

IMPORTANT! To view server banners, you must enable the **Capture Banners** check box in the network discovery policy. This option is disabled by default.

Descriptions of the information provided in the server detail follow.

Protocol

The name of the protocol the server uses.

Port

The port where the server runs.

Hits

The number of times the server was detected by a Sourcefire managed device or Nmap. Note that the number of hits is 0 for servers imported through host input, unless the system detects traffic for that server.

Last Used

The time and date the server was last detected. Note that the last used time for host input data reflects the initial data import time, unless the system detects new traffic for that server. Note also that scanner and application data imported through the host input feature times out according to settings in the system policy, but user input through the Defense Center web interface does not time out.

Application Protocol

The name of the application protocol used by the server, if known.

Vendor

The server vendor. This field does not appear if the vendor is unknown.

Version

The server version. This field does not appear if the version is unknown.

Source

One of the following values:

- User: *user_name*
- Application: *app_name*
- Scanner: *scanner_type* (Nmap or scanner added through system policy)
- **FiresIGHT**, **FiresIGHT Port Match**, or **FiresIGHT Pattern Match**, for Sourcefire-detected applications
- **NetFlow**, for servers added to the network map based on NetFlow data

The system may reconcile data from multiple sources to determine the identity of a server; see [Understanding Current Identities](#) on page 1646.

To view the server detail for a server:

ACCESS: Admin/Any Security Analyst

- ▶ Click the view icon () next to a server in the **Servers** section of a host profile.

The Server Detail pop-up window appears.

Editing Server Identities

LICENSE: FireSIGHT

You can manually update the identity settings for a server on a host and configure any fixes that you have applied to the host to remove the vulnerabilities addressed by the fixes. You can also delete server identities.

Note that deleting an identity does not delete the server, even if that is the only identity. Deleting an identity does remove the identity from the Server Detail pop-up window and, if applicable, updates the current identity for the server in the host profile.

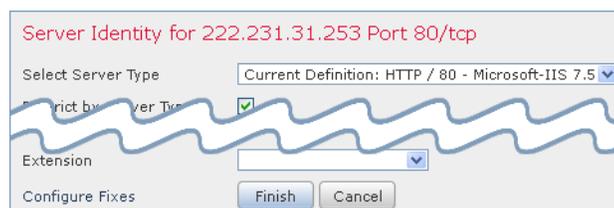
You cannot edit or delete server identities added by a Sourcefire-managed device.

To edit the server identity:

ACCESS: Admin/Any Security Analyst

1. In the **Servers** section in a host profile, click **View** to open the Server Detail pop-up window.
2. You have two options:
 - To delete a server identity, click the delete icon (🗑️) next to the server identity you want to remove.
 - To modify a server identity, click the edit icon (✏️) next to the server in the servers list.

The Server Identity pop-up window appears.

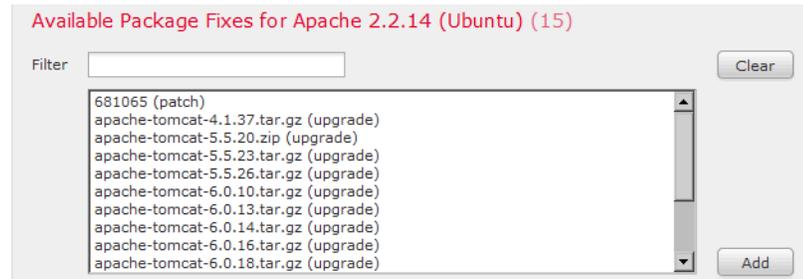


3. You have two options:
 - Select the current definition from the **Select Server Type** drop-down list.
 - Select the type of server from the **Select Server Type** drop-down list.
4. Optionally, to only list vendors and products for that server type, select the **Restrict by Server Type** check box.
5. Optionally, to customize the name and version of the server, select the **Use Custom Display String** and type a **Vendor String** and **Version String**.
6. In the **Product Mappings** section, select the operating system, product, and versions you want to use.

For example, if you want the server to map to Red Hat Linux 9, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

7. If you want to indicate that fixes for the server have been applied, click **Configure Fixes**. Otherwise, skip to step 9.

The Available Package Fixes page appears.



8. Add the patches you want to apply for that server to the fixes list.
9. Click **Finish** to complete the server identity configuration.

Resolving Server Identity Conflicts

LICENSE: FireSIGHT

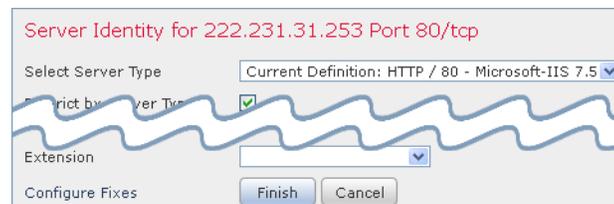
A server identity conflict occurs when an active source, such as an application or scanner, adds identity data for a server to a host, then the system detects traffic for that port that indicates a conflicting server identity.

To resolve a server identity conflict:

ACCESS: Admin/Any Security Analyst

1. Click the resolve icon next to the server in the **Servers** list.

The Server Identity pop-up window appears.

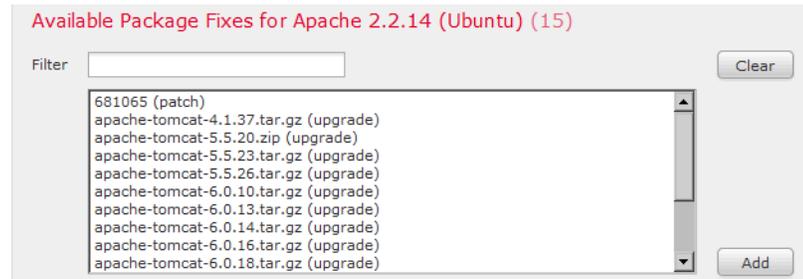


2. Select the type of server from the **Select Server Type** drop-down list.
3. Optionally, to only list vendors and products for that server type, select the **Restrict by Server Type** check box.
4. Optionally, to customize the name and version of the server, select the **Use Custom Display String** and type a **Vendor String** and **Version String**.
5. In the **Product Mappings** section, select the operating system, product, and versions you want to use.

For example, if you want the server to map to Red Hat Linux 9, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

6. If you want to indicate that fixes for the server have been applied, click **Configure Fixes**. Otherwise, skip to step 9.

The Available Package Fixes page appears.



7. Add the patches you want to apply for that server to the fixes list.
8. Click **Finish** to complete the server identity configuration and return to the host profile.

Working with Applications in the Host Profile

LICENSE: FireSIGHT

You can see the applications running on a host in the host profile. If you want to remove an application from a host profile, you can delete that application.

For more information on managing applications in the host profile, see:

- [Viewing Applications in the Host Profile](#) on page 1354
- [Deleting Applications from the Host Profile](#) on page 1356

Viewing Applications in the Host Profile

LICENSE: FireSIGHT

The system can detect a variety of clients and web applications running on the hosts on your network.

IMPORTANT! Note that you must select the **Applications** check box in discovery rules for NetFlow devices in your network discovery policy for the system to detect applications on the hosts in your monitored network. This option is enabled by default in NetFlow rules and cannot be disabled for rules used for discovery via managed devices.

The host profile displays the product and version of the detected applications on a host, any available client or web application information, and the time that the application was last detected in use.

Applications (11) ▾

Application Protocol	Client	Version	Web Application
<input type="checkbox"/> Gnutella	<input type="checkbox"/> Shareaza	2.5.0.0	
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	6.0	<input type="checkbox"/> Dropbox
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	6.0	<input type="checkbox"/> Facebook
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	6.0	<input type="checkbox"/> Google Analytics
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	6.0	<input type="checkbox"/> Microsoft Update
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	6.0	<input type="checkbox"/> Shockwave
<input type="checkbox"/> HTTP	<input type="checkbox"/> Firefox	3.5	Web Browsing
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	5.0	Web Browsing
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	7.0	Web Browsing
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	8.0	Web Browsing
<input type="checkbox"/> HTTP	<input type="checkbox"/> Internet Explorer	6.0	<input type="checkbox"/> Yahoo!

The Defense Center lists up to 16 clients running on the host. After that limit is reached, new client information from any source, whether active or passive, is discarded until you delete a client application from the host or the system deletes the client from the host profile due to inactivity (the client times out).

Additionally, for each detected web browser, the host profile displays the first 100 web applications accessed. After that limit is reached, new web applications associated with that browser from any source, whether active or passive, are discarded until either:

- the web browser client application times out, or
- you delete application information associated with a web application from the host profile

Descriptions of the application information that appears in a host profile follow.

Application Protocol

Displays the application protocol used by the application (HTTP browser, DNS client, and so on).

Client

Client information derived from payload if identified by the Sourcefire 3D System, or captured by Nmap, or by another active source, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Version

Displays the version of the client.

Web Application

For web browsers, the content detected by the system in the http traffic. Web application information indicates the specific type of content (for example, WMV or QuickTime) identified by the Sourcefire 3D System, captured by Nmap, captured by another active source, or acquired via the host input feature. The field is blank if none of the available sources provides an identification.

Note that if the host is running an application that violates a compliance white list in an activated correlation policy, the Defense Center marks the non-compliant application with the white list violation icon (🚫).

To analyze the connection events associated with a particular application on the host, click the events icon (📄) next to the application. The first page of your preferred workflow for connection events appears, showing connection events constrained by the type, product, and version of the application, as well as the IP address(es) of the host. If you do not have a preferred workflow for connection events, you must select one. For more information about connection data, see [Working With Connection Data](#) on page 539.

Deleting Applications from the Host Profile

LICENSE: FireSIGHT

You can delete an application from a host profile to remove applications that you know are not running on the host. Note that deleting an application from a host may bring the host into compliance with a white list.

IMPORTANT! If the system detects the application again, it re-adds it to the network map and the host profile.

To delete an application from a host profile:

ACCESS: Admin/Any Security Analyst

- ▶ In the **Applications** section of the host profile, click the delete icon (🗑️) next to the application you want to delete.
The application is deleted for that host.

Working with VLAN Tags in the Host Profile

LICENSE: FireSIGHT

The VLAN Tag section of the host profile appears if the host is a member of a Virtual LAN (VLAN).



VLAN ID	Type	Priority
100		

Physical network equipment often uses VLANs to create logical network segments from different network blocks. The system detects 802.1q VLAN tags and displays the following information for each:

- **VLAN ID** identifies the VLAN where the host is a member. This can be any integer between zero and 4095 for 802.1q VLANs.
- **Type** identifies the encapsulated packet containing the VLAN tag, which can be either Ethernet or Token Ring.
- **Priority** identifies the priority in the VLAN tag, which can be any integer from zero to 7, where 7 is the highest priority.

If VLAN tags are nested within the packet, the system processes and the Defense Center displays the innermost VLAN tag. The system collects and the Defense Center displays VLAN tag information only for MAC addresses that it identifies through ARP and DHCP traffic.

VLAN tag information can be useful, for example, if you have a VLAN composed entirely of printers and the system detects a Microsoft Windows 2000 operating system in that VLAN. VLAN information also helps the system generate more accurate network maps.

Working with User History in the Host Profile

LICENSE: FireSIGHT

The user history portion of the host profile provides a graphic representation of the last twenty-four hours of user activity. A typical user logs off in the evening and may share the host resource with another user. Periodic login requests, such as those made to check email, are indicated by short regular bars. A list of user identities is provided with bar graphs to indicate when the user login was detected. Note that for non-authoritative logins, the bar graph is gray.

Note that the system does associate a non-authoritative user login to a host with an IP address of that host, so the user does appear in the host's user history. However, if an authoritative user login is detected for the same host, the user associated with the authoritative user login takes over the association with the host IP address, and new non-authoritative user logins do not disrupt that user association with the host IP address. For more information on the types of users,

see [Users Database](#) on page 1258. If you configure capture of failed logins in the network discovery policy, the list includes users that failed to log into the host.

User History ▾		
Users	2011-12-04 16:56:32	2011-12-05 16:56:32
homer (POP3)		

Working with Host Attributes in the Host Profile

LICENSE: FireSIGHT

You can use *host attributes* to classify hosts in ways that are important to your network environment. Host attribute values can be positive integers, strings, or URLs. You can also create a list of string values and assign them automatically based on host IP addresses. For information about creating and managing user-defined host attributes, see [Working with User-Defined Host Attributes](#) on page 1370.

The Sourcefire 3D System includes two predefined host attributes: Host Criticality and Notes. See [Working with the Predefined Host Attributes](#) on page 1369 for information about working with these predefined host attributes.

In addition, each compliance white list that you create automatically creates a host attribute with the same name as the white list. Its possible values are Compliant (for hosts that are compliant with the white list), Non-Compliant (for hosts that violate the white list), or Not Evaluated (for hosts that are not valid targets of the white list or have not been evaluated for any reason). You **cannot** manually change the value of a white list host attribute. For more information on white lists, see [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530.

Users (no user history available)	
Attributes ▾ 	
Host Criticality	Low
Location	New York
Organization	Sales
Notes	04/01/2011 Infected by Sasser. Disinfected.

Assigning Host Attribute Values

LICENSE: FireSIGHT

You can specify positive integers, strings, or URLs as values for existing host attributes.

TIP! You can quickly assign host attributes for a host by clicking the **Edit** link in the **Attributes** section of the host profile page. This launches a pop-up window containing fields for all the host attributes.

To assign a host attribute value:

ACCESS: Admin/Any Security Analyst

1. Open a host profile.
2. Under **Attributes**, click the name of the host attribute to which you want to assign a value.
A pop-up window appears.
3. Enter a value for the attribute or select a value from the drop-down list.
4. Click **Save**.
The host attribute value is saved.

Working with Host Protocols in the Host Profile

LICENSE: FireSIGHT

You can view the protocols running on a host through the host profile. If needed, you can also delete host protocols for a particular host from the profile.

Each host profile contains information about the protocols detected in the network traffic associated with the host.

Host Protocols ▾

Protocol	Layer
tcp	Transport
IP	Network

Descriptions of the protocol and network layer information follow.

Protocol

The name of a protocol used by the host.

Layer

The network layer where the protocol runs (**Network** or **Transport**).

Note that if the host is running a protocol that violates a compliance white list in an activated correlation policy, the Defense Center marks the non-compliant protocol with the white list violation icon (⚠).

Host Protocols ▾

Protocol	Layer
⚠ igmp	Transport

You can delete a protocol from a host profile to remove protocols you know are not running on the host. Note that deleting a protocol from a host may bring the host into compliance with a compliance white list.

IMPORTANT! If the system detects the protocol again, it re-adds it to the network map and the host profile.

To delete a protocol from a host profile:

ACCESS: Admin/Any Security Analyst

- In the **Protocols** section of the host profile, click the delete icon () next to the protocol you want to delete.

The protocol is deleted for that host.

Working with White List Violations in the Host Profile

LICENSE: FireSIGHT

A *compliance white list* (or *white list*) is a set of criteria that allows you to specify the operating systems, application protocols, clients, web applications, and protocols that are allowed to run on a specific subnet.

If you add a white list to an active correlation policy, when the system detects that a host is violating the white list, the Defense Center logs a white list event—which is a special kind of correlation event—to the database. Each of these white list events is associated with a *white list violation*, which indicates how and why a particular host is violating a white list. If a host violates one or more white lists, you can view these violations in its host profile in two ways.

First, the host profile lists all of the individual white list violations associated with the host.

White List Violations (3) ▼

Type	Reason	White List
Web Application	Shockwave	My White List
Web Application	Shockwave	My White List
Client	Ares - Ares	My White List

Descriptions of the white list violation information in the host profile follow.

Type

The type of the violation, that is, whether the violation occurred as a result of a non-compliant operating system, application, server, or protocol.

Reason

The specific reason for the violation. For example, if you have a white list that allows only Microsoft Windows hosts, the host profile displays the current operating system running on the host (such as `Linux Linux 2.4, 2.6`)

White List

The name of the white list associated with the violation.

Second, in the sections associated with operating systems, applications, protocols, and servers, the Defense Center marks non-compliant elements with the white list violation icon (🚫). For example, for a white list that allows only Microsoft Windows hosts, the host profile displays the white list violation icon next to the operating system information for that host.

Systems (1) ▾ 🛠

Hardware	OS Vendor	OS Product	OS Version	Source
🚫	Ubuntu	Linux	11.x, 12.x	FireSIGHT

Note that you can use a host's profile to create a shared host profile for compliance white lists. For more information, see the next section, [Creating a White List Host Profile from a Host Profile](#).

Creating a White List Host Profile from a Host Profile

LICENSE: FireSIGHT

Shared host profiles for compliance white lists specify which operating systems, application protocols, clients, web applications, and protocols are allowed to run on target hosts across multiple white lists. That is, if you create multiple white lists but want to use the same host profile to evaluate hosts running a particular operating system across the white lists, use a shared host profile.

You can use a host profile of any host with a known IP address to create a shared host profile that your compliance white lists can use. However, note that you cannot create a shared host profile based on an individual host's host profile if the system has not yet identified the operating system of the host.

To create a shared host profile for compliance white lists based on a host profile:

ACCESS: Admin

1. Access a host profile from any network map or any event view.
For more information, see [Viewing Host Profiles](#) on page 1337.
2. Click **Generate White List Profile**.

The Edit Shared Profiles page appears. The fields on the page are pre-populated based on the information in the host profile you accessed.

3. Modify and save the shared host profile according to your specific needs.
For more information on creating shared host profiles for compliance white lists, see [Working with Shared Host Profiles](#) on page 1564.

Working with Malware Detections in the Host Profile

LICENSE: FireSIGHT and Malware

The Most Recent Malware Detections section lists the most recent malware events where the host sent or received a malware file, up to 100 events. The host profile lists both network-based and endpoint-based malware events.

Most Recent Malware Detections (2) ▼					
Time	Host Role	Threat Name	File Name	File Type	
2012-09-25 21:31:45	Receiver	Sober-tpd	Worm.Sober	MSEXE	
2012-09-25 21:18:50	Receiver	Sober-tpd	Worm.Sober	MSEXE	

If the host is involved in a file event where the file is then retrospectively identified as malware, the original events where the file was transmitted appear in the malware detections list after the malware identification occurs. When a file identified as malware is retrospectively determined not to be malware, the malware events related to that file no longer appear in the list. For example, if a file has a disposition of **malware** and that disposition changes to **Clean**, the event for that file is removed from the malware detections list on the host profile. For more information on malware events, see [Working with Malware Events](#) on page 1229.

Description of the columns in the Most Recent Malware Detections sections of the host profile follow.

Time

The date and time the event was generated.

For an event where the file was retrospectively identified as malware, note that this is the time of the original event, not the time when the malware was identified.

Host Role

The host's role in the transmission of detected malware, either sender or receiver. Note that for endpoint-based malware events, the host is always the receiver.

Threat Name

The name of the detected malware.

File Name

The name of the malware file.

File Type

The type of file; for example, PDF or MSEXE.

When viewing malware detections in the host profile, you can view malware events for that host in the event viewer. To view events, click the malware icon



Working with Vulnerabilities in the Host Profile

LICENSE: FireSIGHT

The Vulnerabilities sections of the host profile list the vulnerabilities that affect that host.

The Sourcefire Vulnerabilities section lists vulnerabilities based on the operating system, servers, and applications that the system detected on the host.

Vulnerabilities (464)

Name	Remote	Component	Port
BSD SecureLevel Time Setting Security Restriction Bypass Vulnerability		Linux 2.6	
eCryptfs 'parse_tag_3_packet()' Packet Heap Based Buffer Overflow Vulnerability	Yes	Linux 2.6	
Linux IPsec 390 Kbytes SA Cache Instructional Privilege Escalation Vulnerability		Linux 2.6	
Linux Open vSwitch x86_64 Kernel Redefined User Overflow Vulnerability		Linux 2.6	
Unix and Unix-based select() System Call Overflow Vulnerability		Linux 2.6	
Util-Linux Umount Filesystem NULL Pointer Dereference Vulnerability		Linux 2.6	

Displaying row 1 of 1 rows << Page 1 of 1 >>

If there is an identity conflict for either the identity of the host's operating system or one of the application protocols on the host, the system lists vulnerabilities for both identities until the conflict is resolved.

Because there is no operating system information available for hosts added to the network map based on NetFlow data, the Defense Center cannot determine which vulnerabilities may affect those hosts, unless you use the host input feature to manually set the hosts' operating system identity.

Server vendor and version information is often not included in traffic. By default, the system does not map the associated vulnerabilities for the sending and receiving hosts of such traffic. However, using the system policy, you can configure the system to map vulnerabilities for specific application protocols that do not have vendor or version information. For more information, see [Mapping Vulnerabilities for Servers](#) on page 1996.

If you use the host input feature to add third-party vulnerability information for the hosts on your network, additional Vulnerabilities sections appear. For example, if you import vulnerabilities from a QualysGuard Scanner, host profiles on your include a QualysGuard Vulnerabilities section.

You can associate third-party vulnerabilities with operating systems and application protocols, but not clients. For information on importing third-party vulnerabilities, see the *Sourcefire 3D System Host Input API Guide*.

Description of the columns in the Vulnerabilities sections of the host profile follow.

Name

The name of the vulnerability.

Remote

Indicates whether the vulnerability can be remotely exploited. If this column is blank, the vulnerability definition does not include this information.

Component

The name of the operating system, application protocol, or client associated with the vulnerability.

Port

A port number, if the vulnerability is associated with an application protocol running on a specific port.

Keep in mind that for third-party vulnerabilities, the information in the corresponding Vulnerabilities section in the host profile is limited to the information that you provided when you imported the vulnerability data using the host input feature.

When viewing vulnerabilities in the host profile, you can:

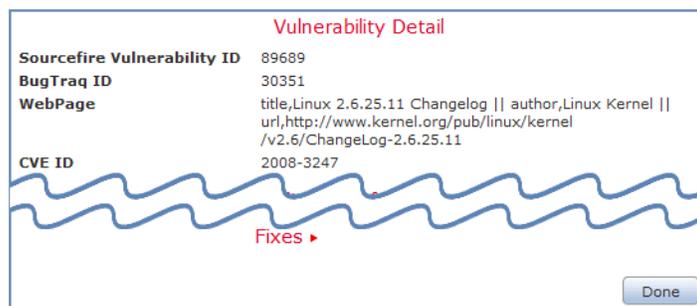
- sort the columns in the **Vulnerabilities** sections by clicking a column heading. To reverse the sort, click again.
- view technical details about a vulnerability, including known solutions, by clicking the name of the vulnerability. See [Viewing Vulnerability Details](#) on page 1365 for more information. Note that you can also access vulnerability details from the vulnerability event views or the Vulnerabilities network map.
- prevent a vulnerability from being used to evaluate impact flag correlations. See [Setting the Vulnerability Impact Qualification](#) on page 1367 for more information.
- download patches to mitigate the vulnerabilities discovered on the hosts on your network. See [Downloading Patches for Vulnerabilities](#) on page 1368 for more information.
- mark hosts as not vulnerable to individual vulnerabilities if you know that the hosts have been patched. See [Setting Vulnerabilities for Individual Hosts](#) on page 1368 for more information.

Viewing Vulnerability Details

LICENSE: FireSIGHT

Vulnerability details include a technical description of the vulnerability and known solutions.

To access the vulnerability details for a specific vulnerability, select **Analysis > Vulnerabilities** or **Analysis > Third-Party Vulnerabilities** and click the view icon (🔍) next to the SVID. You can also access vulnerability details from the network map and the host profile.



Descriptions of the fields on the Vulnerability Detail page follow.

Sourcefire Vulnerability ID

The identification number (SVID) that the system uses to track vulnerabilities.

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If the vulnerability does not have an associated SID, this field does not appear.

BugTraq ID

The identification number associated with the vulnerability in the Bugtraq database (<http://www.securityfocus.com/bid>).

CVE ID

The identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (<http://www.cve.mitre.org/>).

Title

The title of the vulnerability.

Impact Qualification

Use the drop-down list to enable or disable a vulnerability. The Defense Center ignores disabled vulnerabilities in its impact flag correlations.

The setting you specify here determines how the vulnerability is treated on a system-wide basis and is not limited to the host profile where you select the value. See [Setting the Vulnerability Impact Qualification](#) on page 1367 for information about using this feature to enable and disable a vulnerability.

Date Published

The date that the vulnerability was published.

Vulnerability Impact

The severity assigned to the vulnerability in the Bugtraq database on a scale of 1 to 10, with 10 being the most severe. The vulnerability impact is determined by the writer of the Bugtraq entry, who determines the vulnerability impact level based on his or her best judgment, guided by SANS Critical Vulnerability Analysis (CVA) criteria.

Remote

Indicates whether the vulnerability is remotely exploitable.

Available Exploits

Indicates whether there are known exploits for the vulnerability.

Description

Summary description of the vulnerability.

Technical Description

Detailed technical description of the vulnerability.

Solution

Information about repairing the vulnerability.

Additional Information

Click the arrow to view additional information (if available) about the vulnerability, such as known exploits and their availability, exploit scenarios, and mitigation strategies.

Fixes

Provides links to downloadable patches for the selected vulnerability.

TIP! If direct links to fix or patch downloads appear, right-click the link and save it to your local computer.

Setting the Vulnerability Impact Qualification

LICENSE: FireSIGHT

If the system reports a vulnerability that is not applicable to your network, you can prevent it from being used to evaluate impact flag correlations. Note that if you deactivate a vulnerability in a host profile, it deactivates it for all hosts on your network. You can, however, reactivate it at any time.

When a conflict exists for the identity of the host's operating system or one of the applications on the host, the system lists vulnerabilities for both conflicting identities until the conflict is resolved. For more information, see [Understanding Identity Conflicts](#) on page 1648 and [Resolving Operating System Identity Conflicts](#) on page 1345.

Note also that the system does not recommend a rule state for an intrusion rule based on a vulnerability that you disable using the Impact Qualification feature. For more information, see [Managing FireSIGHT Rule State Recommendations](#) on page 764.

TIP! You can also deactivate vulnerabilities from the network map and from vulnerability event views. For more information, see [Working with the Vulnerabilities Network Map](#) on page 1323 and [Deactivating Sourcefire Vulnerabilities](#) on page 1438.

To change the use of a vulnerability across the system:

ACCESS: Admin/Any Security Analyst

1. Access the host profile of a host affected by the vulnerability you want to deactivate.
2. Expand the **Vulnerabilities** section.
3. Click the name of the vulnerability you want to enable or disable.
A pop-up window appears with the vulnerability details. For more information, see [Viewing Vulnerability Details](#) on page 1365.
4. Select **Disabled** or **Enabled** from the **Impact Qualification** drop-down list to specify how the vulnerability is used.
5. Confirm that you want to change the Impact Qualification for all hosts on the network map.
The vulnerability is enabled or disabled.

6. Click **Done** to close the vulnerability details pop-up window.

Downloading Patches for Vulnerabilities

LICENSE: FireSIGHT

If they are available, you can download patches to mitigate the vulnerabilities discovered on the hosts on your network.

To download a patch for a vulnerability:

ACCESS: Admin/Any Security Analyst

1. Access the host profile of a host for which you want to download a patch.
2. Expand the **Vulnerabilities** section.
3. Click the name of the vulnerability you want to patch.
The Vulnerability Detail page appears.
4. Expand the **Fixes** section.
The list of downloadable patches for the vulnerability appears.

Fixes ▼		
Patch 4416		Download
Patch T64v50ab17-c0018301-13396-es-20020226.tar		Download
Patch T64v40gb17-c0010301-13400-es-20020226.tar		Download
Patch DUV40FB18-C0067301-13427-ES-20020228.tar	Prerequisite: 4.0F with Patch Kit 7 (BL18) installed	Download
APAR IY25504		Download
Patch T64v51b18-c0102001-13428-es-20020228.tar		Download
Patch CDE_libDtSvc_efix.tar.Z		Download
Patch T64V51AB1-C0011201-13438-ES-20020228.tar	Prerequisite: 5.1A with Patch Kit 1 (BL1) installed	Download

5. Click **Download** next to the patch you want to download.
A download page from the patch vendor appears.
6. Download the patch and apply it to your affected systems.

Setting Vulnerabilities for Individual Hosts

LICENSE: FireSIGHT

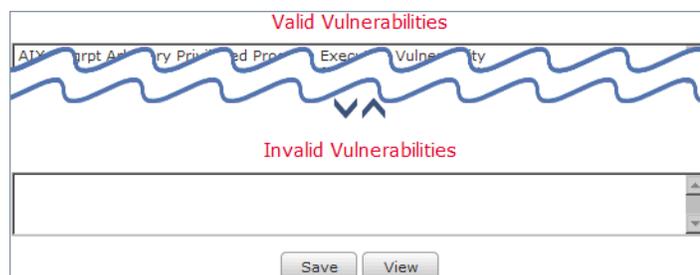
You can use the host vulnerability editor to activate or deactivate vulnerabilities on a host-by-host basis. When you deactivate a vulnerability for a host, it is still used for impact flag correlations for that host, but the impact flag is automatically reduced one level.

To activate or deactivate a vulnerability for a single host:

ACCESS: Admin/Security Analyst

1. Open a host profile.
2. Next to **Vulnerabilities**, click **Edit**.

The Host Vulnerabilities editor page appears.



TIP! To view details about a vulnerability, select it and click **View**. For more information, see [Viewing Vulnerability Details](#) on page 1365.

3. You have two options:
 - To deactivate a vulnerability, select it from the **Valid Vulnerabilities** list, then click the down arrow.
 - To activate a vulnerability, select it from the **Invalid Vulnerabilities** list, then click the up arrow.

TIP! Use Ctrl or Shift while clicking to select multiple vulnerabilities. You can click and drag to select multiple adjacent vulnerabilities; you can also double-click any vulnerability to move it from list to list.

4. Click **Save**.
Your changes are saved.

Working with the Predefined Host Attributes

LICENSE: FireSIGHT

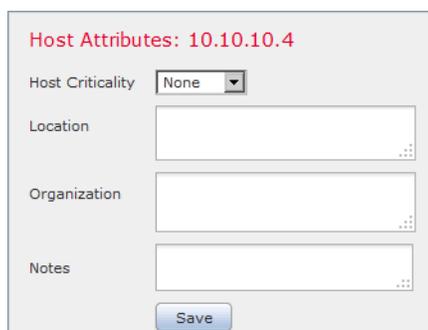
There are two predefined host attributes that you can assign to each host: host criticality and host-specific notes. Use the host criticality attribute to designate the business criticality of a given host and to tailor correlation policies and alerts based on host criticality. For example, if you consider your organization's mail servers more critical to your business than a typical user workstation, you can assign a value of High to your mail servers and other business-critical devices and Medium or Low to other hosts. You can then create a correlation policy that launches different alerts based on the criticality of an affected host.

Use the Notes feature to record information about the host that you want other analysts to view. For example, if you have a computer on the network that has an older, unpatched version of an operating system that you use for testing, you can use the Notes feature to indicate that the system is intentionally unpatched.

To set pre-defined host attributes in the host profile:

ACCESS: Admin/Security Analyst

1. Open the host profile for the host for which you want to set a business criticality.
2. Next to **Attributes**, click the pencil icon (✎).
The Host Attributes pop-up window appears.



The screenshot shows a pop-up window titled "Host Attributes: 10.10.10.4". It contains a "Host Criticality" dropdown menu with "None" selected. Below this are three text input fields labeled "Location", "Organization", and "Notes". At the bottom of the window is a "Save" button.

3. From the **Host Criticality** drop-down list, select the value you want to apply: **None**, **Low**, **Medium**, or **High**.
4. Click **Save**.
Your selection is saved.

Working with User-Defined Host Attributes

LICENSE: FireSIGHT

The Sourcefire 3D System includes two predefined host attributes, host criticality and host notes, that you can use to indicate the business criticality of the hosts on your network. If you have other criteria that you would like to use to identify your hosts, you can create user-defined host attributes.

User-defined host attributes appear in the host profile page, where you can assign values on a per-host basis. You can then use those attributes in correlation policies and searches. You can also view the attributes on the host attribute table view of events and generate reports based on them.

IMPORTANT! Host attributes are defined globally rather than per policy. After you create a host attribute, it is available regardless of the policy applied.

Some examples of user-defined host attributes include:

- assigning physical location identifiers to hosts, such as a facility code, city, or room number.
- assigning a Responsible Party Identifier that indicates which system administrator is responsible for a given host. You can then craft correlation rules and policies to send alerts to the correct system administrator when problems related to a host are detected.

Host attributes can be text strings or values selected from predefined lists of text or ranges of numbers. You can also automatically assign values to hosts from a predefined list based on the hosts' IP addresses. You can use this feature to automatically assign values to new hosts when they appear on your network for the first time.

Host attributes can be one of the following types:

Text

Allows you to manually assign a text string up to 255 characters to a host.

Integer

Allows you to specify the first and last number of a range of positive integers, then manually assign one of these numbers to a host.

List

Allows you to create a list of string values, then manually assign one of the values to a host. You can also automatically assign values to hosts based on the host's IP addresses.

IMPORTANT! If you auto-assign values based on one IP address of a host with multiple IP addresses, those values will apply across all addresses associated with that host. Keep this in mind when you view the Host Attributes table.

URL

Allows you to manually assign a URL value to a host.

Note that each compliance white list that you create automatically creates a host attribute with the same name as the white list. Its possible values are **Compliant** (for hosts that are compliant with the white list), **Non-Compliant** (for hosts that violate the white list), and **Not Evaluated** (for hosts that are not valid targets of the white list or have not been evaluated for any reason). You **cannot** manually change the value of a white list host attribute. For more information on white lists, see [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530.

See the following sections for more information:

- [Creating User-Defined Host Attributes](#) on page 1372
- [Editing a User-Defined Host Attribute](#) on page 1374
- [Deleting a User-Defined Host Attribute](#) on page 1375

Creating User-Defined Host Attributes

LICENSE: FireSIGHT

The following procedure explains how to create a user-defined host attribute.

IMPORTANT! Host attributes are defined globally rather than per policy. After you create a host attribute, it is available regardless of the policy applied.

To create a new host attribute:

ACCESS: Admin/Discovery Admin

1. Select **Analysis > Hosts > Host Attributes**.

The Host Attributes page appears.

2. Click **Host Attribute Management**.

The Host Attribute Management page appears.

3. Click **Create Attribute**.

The Create Attribute page appears.



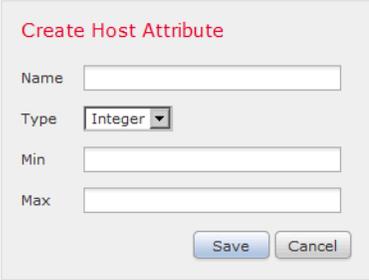
4. In the **Name** field, type a name for the host attribute using alphanumeric characters and spaces.
5. Select the type of attribute that you want to create from the **Type** drop-down list, as described in [Working with Host Attributes in the Host Profile](#) on page 1358:
 - If you are creating a **Text** or **URL** host attribute, continue with step 6.
 - If you are creating an **Integer** host attribute, see [Creating Integer Host Attributes](#) on page 1373.
 - If you are creating a **List** host attribute, see [Creating List Host Attributes](#) on page 1373.
6. Click **Save**.

The new user-defined host attribute is saved.

Creating Integer Host Attributes

LICENSE: FireSIGHT

When you define an integer-based host attribute, you must specify the range of numbers that the attribute accepts.



The 'Create Host Attribute' dialog box contains the following fields and controls:

- Name:** A text input field.
- Type:** A dropdown menu with 'Integer' selected.
- Min:** A text input field for the minimum value.
- Max:** A text input field for the maximum value.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom right.

To create an integer-based host attribute:

ACCESS: Admin/Discovery Admin

1. In the **Min** field, enter the minimum integer value that can be assigned to a host.
2. In the **Max** field, enter the maximum integer value that can be assigned to a host.
3. Click **Save**.

The new integer-based host attribute is saved.

Creating List Host Attributes

LICENSE: FireSIGHT

When you define a list-based host attribute, you must supply each of the values for the list. These values can contain alphanumeric characters, spaces, and symbols.

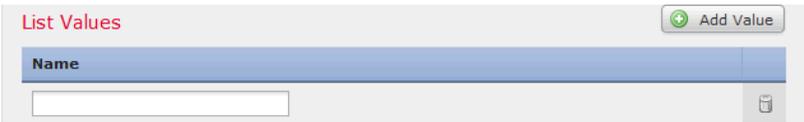
When you create a value for the host attribute, you can also auto-assign it to a range of IP addresses so that when a new host is discovered, it is automatically assigned a value for the host attribute.

To create a list-based host attribute:

ACCESS: Admin/Discovery Admin

1. To add a value to the list, click **Add Value**.

The List Values section expands.



The 'List Values' section shows a table with one row and one column labeled 'Name'. Below the table is an input field for entering a value. To the right of the table is a trash icon. Above the table is a '+ Add Value' button.

2. In the **Name** field, enter the first value you want to add, using alphanumeric characters, symbols, and spaces.

- Optionally, to auto-assign the attribute value you just added to your hosts, click **Add Networks**.

The Auto-Assign Networks section expands.

- Select the value you added from the **Value** drop-down list.
- In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in CIDR notation) that represent the IP address range where you want to auto-assign this value.

For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

- Repeat steps 1 through 5 to add additional values to the list and assign them automatically to new hosts that fall within an IP address range.

TIP! If you do not want to auto-assign list values to the hosts in specific IP ranges, you can manually assign them as described in [Working with the Predefined Host Attributes](#) on page 1369.

Editing a User-Defined Host Attribute

LICENSE: FireSIGHT

When you modify an existing user-defined host attribute, you can change the definition of a value, but you cannot change the attribute type (text, list, integer, URL). In addition, you **cannot** modify compliance white list host attributes.

To edit an existing user-defined host attribute:

ACCESS: Admin/Discovery Admin

- Select **Analysis > Hosts > Host Attributes**.

The Host Attributes page appears.

- Click **Host Attribute Management**.

The Host Attribute Management page appears.

Name	Type	Auto-Assign	
Default White List	White List	<input type="checkbox"/>	
Example Attribute 1	Text	<input type="checkbox"/>	 
Example Attribute 2	Integer	<input type="checkbox"/>	 

- Click the edit icon (✎) next to the host attribute you want to edit.
The host attribute page appears with the selected attribute's settings.
- Modify any of the settings that you want and click **Save**.
See [Creating User-Defined Host Attributes](#) on page 1372 for more information about the attribute types that you can edit and the values those attributes can contain.

Deleting a User-Defined Host Attribute

LICENSE: FireSIGHT

Delete a user-defined host attribute to remove it from every host profile where it is used. Note that you **cannot** delete compliance white list host attributes.

To delete a host attribute:

ACCESS: Admin/Discovery Admin

- Select **Analysis > Hosts > Host Attributes**.
The Host Attributes page appears.
- Click **Host Attribute Management**.
The Host Attribute Management page appears.
- Click the delete icon (🗑) next to the host attribute you want to delete.
The selected host attribute is removed from the system.

Working with Scan Results in a Host Profile

LICENSE: FireSIGHT

When you scan a host using Nmap, or when you import results from an Nmap scan, those results appear in the host profile for any hosts included in the scan.

The information that Nmap collects about the host operating system and any servers running on open unfiltered ports is added directly into the Operating System and Servers sections of the host profile, respectively. In addition, Nmap adds a list of the scan results for that host in the Scan Results section.

Scan Results (3) ▼		
Nmap	Linux 2.6.24 - 2.6.31	89%
Nmap	22/tcp ssh	open, syn-ack, OpenSSH, 5.6, protocol 2.0
Nmap	443/tcp http	open, syn-ack, Apache httpd

Each result indicates the source of the information, the number and type of the scanned port, the name of the server running on the port, and any additional information detected by Nmap, such as the state of the port or the vendor name for the server. If you scan for UDP ports, servers detected on those ports only appear in the Scan Results section.

Note that you can run an Nmap scan from the host profile. For more information, see the next section, [Scanning a Host from the Host Profile](#).

Scanning a Host from the Host Profile

LICENSE: FireSIGHT

You can perform a Nmap scan against a host from the host profile. After the scan completes, server and operating system information for that host are updated in the host profile. Any additional scan results are added to the Scan Results section of the host profile.

WARNING! Nmap-supplied server and operating system data remains static until you run another Nmap scan or override it with higher priority host input. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep Nmap-supplied operating system and server data up to date. For more information, see [Automating Nmap Scans](#) on page 1936.

To scan a host from the host profile:

ACCESS: Admin

1. In the host profile, click **Scan Host**.

The Scan Host pop-up window appears.



2. Click **Scan** next to the scan remediation you want to use to scan the host.

The host is scanned and the results are added to the host profile.

CHAPTER 35

WORKING WITH DISCOVERY EVENTS

Discovery events alert you to the activity on your network and provide you with the information you need to respond appropriately. They are triggered by the changes that your managed devices detect in the network segments they monitor. Your *network discovery policy* specifies the kinds of data the system collects, the monitored network segments, and the specific hardware interfaces that your system uses to monitor traffic. For more information on network discovery, see [Understanding Discovery Data Collection](#) on page 1251.

As a simple example of a discovery event, you may have conference rooms or spare work spaces where visiting employees attach to your network. You would expect to see New Host events generated on these segments on a regular basis, and you would not suspect malicious intent. However, if you see a New Host event on a network segment that is locked down, then you can escalate your response accordingly.

User discovery events provide information about users logged into the hosts on your network. You can view events that catalog user activity on the network and drill down to view information on a particular user. For example, if you want to see what user is associated with a new host, you can check the host profile to find out what users have been detected in traffic going to or from that host.

Discovery events provide you with much greater depth of insight into the activity on your network and with much more granularity than this simple example shows. For each monitored host, you can configure the system to detect related application protocols, network protocols, clients, users, and potential vulnerabilities. The system can also provide information on vulnerabilities detected by third-party scanners that you import onto the Defense Center using the host input feature. In addition, you can track any changes in host criticality, host attribute, or vulnerability settings that users enter via the user interface.

The system provides a set of predefined workflows that you can use to analyze the discovery events that your system generates. You can also create custom workflows that display only the information that matches your specific needs.

To collect and store network discovery data for analysis, make sure that your network discovery policy is configured to discover the appropriate data on the networks and zones where your Sourcefire-managed devices and NetFlow-enabled devices monitor traffic. To exclude monitored areas from discovery, configure that in the network discovery policy. Note that an access control policy must be applied to the managed device before you can apply a network discovery policy. For more information, see [Creating a Network Discovery Policy](#) on page 1275.

For more information, see:

- [Viewing Discovery Event Statistics](#) on page 1378
- [Viewing Discovery Performance Graphs](#) on page 1384
- [Understanding Discovery Event Workflows](#) on page 1386
- [Working with Discovery and Host Input Events](#) on page 1388
- [Working with Hosts](#) on page 1401
- [Working with Host Attributes](#) on page 1412
- [Working with Servers](#) on page 1418
- [Working with Applications](#) on page 1424
- [Working with Application Details](#) on page 1429
- [Working with Sourcefire Vulnerabilities](#) on page 1434
- [Working with Third-Party Vulnerabilities](#) on page 1441
- [Working with Users](#) on page 1445
- [Working with User Activity](#) on page 1453

Viewing Discovery Event Statistics

LICENSE: FireSIGHT

The Discovery Statistics Summary page displays a summary of the hosts, events, protocols, application protocols, and operating systems detected by the system:

- The statistics summary provides general statistics about the total events, application protocols, hosts, network devices, and information about your host limit usage; see [Statistics Summary](#) on page 1379.
- The event breakdown provides statistics about the types of events occurring on the system; see [Event Breakdown](#) on page 1381.
- The protocol breakdown provides statistics about the protocols that detected hosts are using. See [Protocol Breakdown](#) on page 1382.

- The application protocol breakdown provides statistics about the application protocols running on the network; see [Application Protocol Breakdown](#) on page 1382.
- The operating system breakdown lists the operating systems that are running on the network and how many hosts are using each operating system; see [OS Breakdown](#) on page 1383.

The page lists statistics for the last hour and the total accumulated statistics. You can select statistics for a particular device, or all devices. You can also view events that match the entries on the page by clicking the event, server, operating system, or operating system vendor listed within the summary.

To view the Event Statistics Summary:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Summary > Discovery Statistics**.

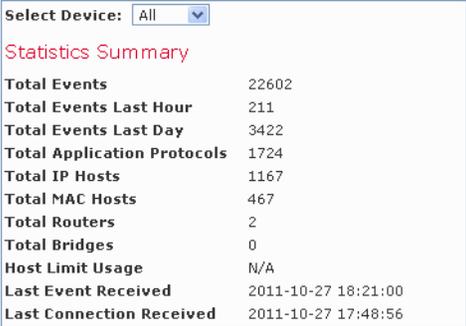
The statistics summary page appears.

2. From the **Select Device** list, select the device whose statistics you want to view. Select **All** to view statistics for all devices managed by the Defense Center.

Statistics Summary

LICENSE: FireSIGHT

The statistics summary provides general statistics about the total events, application protocols, hosts, network devices, and information about your host limit usage.



Statistics Summary	
Total Events	22602
Total Events Last Hour	211
Total Events Last Day	3422
Total Application Protocols	1724
Total IP Hosts	1167
Total MAC Hosts	467
Total Routers	2
Total Bridges	0
Host Limit Usage	N/A
Last Event Received	2011-10-27 18:21:00
Last Connection Received	2011-10-27 17:48:56

Descriptions of the rows of the Statistics Summary section follow.

Total Events

Total number of discovery events stored on the Defense Center.

Total Events Last Hour

Total number of discovery events generated in the last hour.

Total Events Last Day

Total number of discovery events generated in the last day.

Total Application Protocols

Total number of application protocols from servers running on detected hosts.

Total IP Hosts

Total number of detected hosts identified by unique IP address.

Total MAC Hosts

Total number of detected hosts not identified by IP address.

Note that the Total MAC Hosts statistic remains the same whether you are viewing discovery statistics for all devices or for a specific device. This is so because managed devices discover hosts based on their IP addresses. This statistic gives the total of all hosts that are identified by other means and is independent of a given managed device.

Total Routers

Total number of detected nodes identified as routers.

Total Bridges

Total number of detected nodes identified as bridges.

Host Limit Usage

Total percentage of the host limit currently in use. The host limit is defined by your FireSIGHT license. Note that the host limit usage only appears if you are viewing statistics for all managed devices. For more information on monitoring host usage, see [Configuring FireSIGHT Host Usage Monitoring](#) on page 2129.

IMPORTANT! If the host limit is reached and a host is deleted, the host will not reappear on the network map until you restart network discovery on all managed devices configured to perform discovery.

Last Event Received

The date and time that the most recent discovery event occurred.

Last Connection Received

The date and time that the most recent connection was completed.

Event Breakdown

LICENSE: FireSIGHT

The Event Breakdown section lists a count of each type of network discovery and host input event that occurred within the last hour, as well as a count of the total number of each event type stored in the database. For full descriptions of each event type, see [Understanding Discovery Event Types](#) on page 1389 and [Understanding Host Input Event Types](#) on page 1394.

Event Breakdown

Event	Total Last Hour	Total
Add Application Protocol	0	0
Add Client	0	0
Add Host	0	0
AN T...orma...edat	0	0
Vulnerability Set Invalid	0	0
Vulnerability Set Valid	0	0

You can also use the Events Breakdown section to view details on discovery and host input events.

To view network discovery and host input events by type:

ACCESS: Admin/Any Security Analyst

- ▶ Click the type of event you want to view.

The first page of the default discovery events workflow appears, constrained by the event type you picked. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

For information on working with discovery events, see [Working with Discovery and Host Input Events](#) on page 1388.

Protocol Breakdown

LICENSE: FireSIGHT

The Protocol Breakdown section lists the protocols currently in use by detected hosts. It displays each detected protocol name, its “layer” in the protocol stack, and the total number of hosts that communicate using the protocol.

Protocol Breakdown

Protocol	Layer	Total
IP	Network	1167
tcp	Transport	1002
udp	Transport	854
ARP	Network	542
icmp	Transport	431
igmp	Transport	308
sctp	Transport	1

Application Protocol Breakdown

LICENSE: FireSIGHT

The Application Protocol Breakdown section lists the application protocols that are currently in use by detected hosts. It lists the protocol name, the total number of hosts running the application protocol in the past hour, and the total number of hosts that have been detected running the protocol at any point.

Application Protocol Breakdown

Application	Total Last Hour	Total
<input type="checkbox"/> Unknown	0	1044
<input type="checkbox"/> SSH	26	324
<input type="checkbox"/> SSL	4	177
<input type="checkbox"/> Acave Directory	0	1
<input type="checkbox"/> Tivoli	0	1
<input type="checkbox"/> 1000008	0	1

You can also use the Application Protocol Breakdown section to view details on servers using the detected protocols.

To view servers that use a listed application protocol:

ACCESS: Admin/Any Security Analyst

- ▶ Click the name of the application protocol you want to view.

The first page of the default servers workflow appears, constrained by the application protocol you picked. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

For information on working with servers, see [Working with Servers](#) on page 1418.

OS Breakdown

LICENSE: FireSIGHT

The OS Breakdown section lists the operating systems currently running on the monitored network, along with their vendors and the total number of hosts running each operating system.

A value of **unknown** for the operating system name or version means that the operating system or its version does not match any of the system's fingerprints. A value of **pending** means that the system has not yet gathered enough information to identify the operating system or its version.

OS Breakdown

OS Name	OS Vendor	Total
Linux	Linux	539
pending	pending	261
AIX	IBM	131
Mac OS X	Apple	130
Windows	Microsoft	120
unknown	unknown	105
FreeBSD	FreeBSD	11
OpenBSD	OpenBSD	2
Linux	Debian	1

You can use the OS Breakdown section to view details on the detected operating systems.

To view hosts by operating system or vendor:

ACCESS: Admin/Any Security Analyst

► You have two options:

- To view all hosts running a specific operating system, under **OS Name**, click the operating system name.
- To view all hosts running any operating system from a specific vendor, under **OS Vendor**, click the vendor name.

The first page of the default hosts workflow appears, constrained by the operating system or vendor you picked. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

For information on working with hosts, see [Working with Hosts](#) on page 1401.

Viewing Discovery Performance Graphs

LICENSE: FireSIGHT

You can generate graphs that display performance statistics for managed devices with discovery events.

IMPORTANT! New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs.

Descriptions of the available graph types follow.

Processed Events/Sec

Displays a graph that represents the number of events that the Data Correlator processes per second

Processed Connections/Sec

Displays a graph that represents the number of connections that the Data Correlator processes per second

Generated Events/Sec

Displays a graph that represents the number of events that the system generates per second

Mbits/Sec

Displays a graph that represents the number of megabits of traffic that are analyzed by the discovery process per second

Avg Bytes/Packet

Displays a graph that represents the average number of bytes included in each packet analyzed by the discovery process

K Packets/Sec

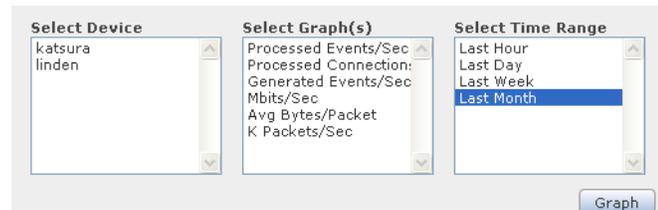
Displays a graph that represents the number of packets analyzed by the discovery process per second, in thousands

To generate discovery performance graphs:

ACCESS: Admin/Maint

1. Select **Overview > Summary > Discovery Performance**.

The Discovery Performance page appears.



2. From the **Select Device** list, select the Defense Center or managed devices you want to include.

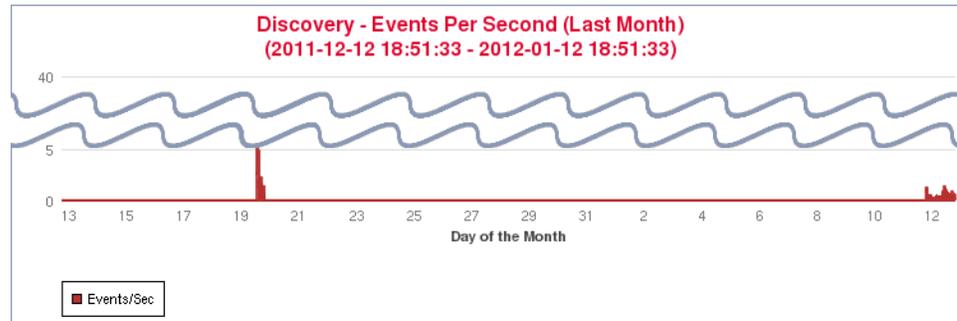
Depending on which appliance you select, the **Select Graph(s)** list adjusts to display the available graphs.

3. From the **Select Graph(s)** list, select the type of graph you want to create.

TIP! You can select multiple graphs by holding down the Ctrl or Shift keys while clicking on the graph type.

4. From the **Select Time Range** list, select the time range you would like to use for the graph. You can choose from last hour, last day, last week, or last month.

5. Click **Graph** to graph the selected statistics.
The selected graph appears.



Understanding Discovery Event Workflows

LICENSE: FireSIGHT

The Defense Center provides a set of workflows that you can use to analyze the discovery events that are generated for your network. The workflows are, along with the network map, a key source of information about your network assets. These workflows contain tables that are populated with discovery data generated by the system.

Access network discovery workflows from the **Analysis > Hosts** menu. The Defense Center provides predefined workflows for discovery events, as well as for detected hosts and their host attributes, servers, applications, application details, vulnerabilities, user activities, and users. You can also create custom workflows. For more information on workflows, see [Understanding and Using Workflows](#) on page 1791.

TIP! Select **Analysis > Custom > Custom Tables** to access workflows based on custom tables.

When you are using a network discovery workflow, you can perform many common actions, whatever the type of event. These common functions are described in the [Common Discovery Event Actions](#) table.

Common Discovery Event Actions

To...	YOU CAN...
view the host profile for an IP address	click the host profile icon () that appears next to the IP address.
view user profile information	click the user icon () that appears next to the user identity. For more information, see Understanding User Details and Host History on page 1449.
sort data	click the column title. Click the column title again to reverse the sort order.
drill down to the next page in the workflow	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To drill down to the next workflow page constraining on a specific value, click a value within a row. Note that this only works on drill-down pages. Clicking a value within a row in a table view only constrains the table view and does not drill down to the next page. • To drill down to the next workflow page constraining on some events, select the check boxes next to the events you want to view on the next workflow page, then click View. • To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include “Table View” in the page name.</p> <p>For more information, see Constraining Events on page 1830.</p>
constrain the columns that appear	<p>click the close icon () in the column heading that you want to hide. In the pop-up window that appears, click Apply.</p> <p>TIP! To hide or show other columns, Select or clear the appropriate check boxes before you click Apply. To add a disabled column back to the view, click the expand arrow to expand the search constraints, then click the column name under Disabled Columns.</p>
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1836.

Common Discovery Event Actions (Continued)

To...	YOU CAN...
navigate between pages in the current workflow, keeping the current constraints	click the appropriate page link at the top left of the workflow page. For more information, see Using Workflow Pages on page 1812.
delete items from the system, including: <ul style="list-style-type: none"> • discovery and host input events from discovery event workflows • hosts and network devices from host workflows • host attributes from host attribute workflows • servers from server workflows • applications from application workflows • third-party vulnerabilities from third-party vulnerability workflows • users from user workflows 	use one of the following methods: <ul style="list-style-type: none"> • To delete some items, select the check boxes next to items you want to delete, then click Delete. • To delete all items in the current constrained view, click Delete All, then confirm you want to delete all the items. These items remain deleted until the system's discovery function is restarted, when they may be detected again. <p>TIP! See Purging Discovery Data from the Database on page 2236 for information on deleting all discovery events from the database and also for information on how to restart discovery.</p> Note that you cannot delete Sourcefire (as opposed to third-party) vulnerabilities; you can, however, mark them reviewed. For more information, see Working with Sourcefire Vulnerabilities on page 1434.
navigate to other event views to view associated events	find more information in Navigating Between Workflows on page 1836.

Working with Discovery and Host Input Events

LICENSE: FireSIGHT

The system generates discovery events that communicate the details of changes in your monitored network segments. *New* events are generated for newly discovered network features, and change events are generated for any change in previously identified network assets.

During its initial network discovery phase, the system generates new events for each host and any TCP or UDP servers discovered running on each host. Optionally, you can configure the system to use data exported by NetFlow-enabled devices to generate these new host and server events.

In addition, the system generates new events for each network, transport, and application protocol running on each discovered host. When you create a discovery rule configured to include NetFlow-enabled devices, you can disable detection of application protocols. However, you cannot disable application detection in discovery rules that do not use a configured NetFlow-enabled device. If you enable host or user discovery in non-NetFlow discovery rules, applications are automatically discovered.

After the initial network mapping is complete, the system continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered asset changes.

When a discovery event is generated, it is logged to the database. You can use the Defense Center web interface to view, search, and delete discovery events. You can also use discovery events in correlation rules. Based on the type of discovery event generated as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria.

You can add data to the network map using the host input feature. You can add, modify, or delete operating system information, which causes the system to stop updating that information for that host. You can also manually add, modify, or delete application protocols, clients, servers, and host attributes or modify vulnerability information. When you do this, the system generates host input events.

See the following sections for more information:

- [Understanding Discovery Event Types](#) on page 1389
- [Understanding Host Input Event Types](#) on page 1394
- [Viewing Discovery and Host Input Events](#) on page 1396
- [Understanding the Discovery Events Table](#) on page 1397
- [Searching for Discovery Events](#) on page 1399

Understanding Discovery Event Types

LICENSE: FireSIGHT

There are many types of discovery events. For example, the system generates and logs a New Host event when it detects a new host on your monitored network segment. When you view a table of discovery events, the event type is listed in the **Event** column. For more information, see [Viewing Discovery and Host Input Events](#) on page 1396.

Contrast discovery events, which are generated when the system detects a change in your monitored network (such as detecting traffic from a previously undetected host), with host input events, which are generated when a user takes a specific action (such as manually adding a host). For more information on host input events, see [Understanding Host Input Event Types](#) on page 1394.

You can configure the types of discovery events the system logs by modifying your network discovery policy. By default, the system logs all types of discovery events. For more information, see [Configuring Database Event Limits](#) on page 1978.

If you understand the information the different types of discovery events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in correlation policies. In addition, knowing the

names of the event types can help you craft more effective event searches. Descriptions of the different types of discovery events follow.

Additional MAC Detected for Host

This event is generated when the system detects a new MAC address for a previously discovered host.

This event is often generated when the system detects hosts passing traffic through a router. While each host has a different IP address, they all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an “ARP/DHCP detected” message within the event description in the event view.

Client Timeout

This event is generated when the system drops a client from the database due to inactivity.

Client Update

This event is generated when the system detects a payload (that is, a specific type of content, such as audio, video, or webmail) in HTTP traffic.

DHCP: IP Address Changed

This event is generated when the system detects that a host IP address has changed due to DHCP address assignment.

DHCP: IP Address Reassigned

This event is generated when a host is reusing an IP address; that is, when a host obtains an IP address formerly used by another physical host due to DHCP IP address assignment.

Hops Change

This event is generated when the system detects a change in the number of network hops between a host and the device that detects the host.

This may happen if the device sees host traffic through different routers and is able to make a better determination of the host’s location. This may also happen if the device detects an ARP transmission from the host, indicating that the host is on a local segment.

Host Deleted: Host Limit Reached

This event is generated when the host limit on the Defense Center is exceeded and a monitored host is deleted from the Defense Center’s network map.

Host Dropped: Host Limit Reached

This event is generated when the host limit on the Defense Center is reached and a new host is dropped. Compare this with the previous event where old hosts are deleted from the network map when the host limit is reached.

To drop new hosts when the host limit is reached, go to **Policies > Network Discovery > Advanced** and set **When Host Limit Reached** to **Drop hosts**. See [Configuring Data Storage](#) on page 1293 for more information.

Host Timeout

This event is generated when a host is dropped from the network map because the host has not produced traffic within the interval defined in the network discovery policy. Note that individual host IP addresses and MAC addresses time out individually; a host does not disappear from the network map unless all of its associated addresses have timed out. See [Configuring Data Storage](#) on page 1293 for information about configuring the host timeout value.

If you change the networks you want to monitor in your network discovery policy, you may want to manually delete old hosts from the network map so that they do not count against your FireSIGHT license. For more information, see [Working with the Hosts Network Map](#) on page 1316.

Host Type Changed to Network Device

This event is generated when the system detects that a detected host is actually a network device.

Identity Conflict

This event is generated when the system detects a new server or operating system identity that conflicts with a current active identity for that server or operating system.

If you want to resolve identity conflicts by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation. For more information, see [Configuring Nmap Remediations](#) on page 1623.

For more information, see [Understanding Identity Conflicts](#) on page 1648 and [Configuring Identity Conflict Resolution](#) on page 1289. For information on manually resolving conflicts, see [Resolving Operating System Identity Conflicts](#) on page 1345 and [Resolving Server Identity Conflicts](#) on page 1353.

Identity Timeout

This event is generated when identity data that was added to the network map through an active source times out.

If you want to refresh identity data by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation. For more information, see [Configuring Nmap Remediations](#) on page 1623.

For more information, see [Resolving Server Identity Conflicts](#) on page 1353.

MAC Information Change

This event is generated when the system detects a change in the information associated with a specific MAC address or TTL value.

This event often occurs when the system detects hosts passing traffic through a router. While each host has a different IP address, they will all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an “ARP/DHCP detected” message within the event description in the event view. The TTL may change because the traffic may pass through different routers or if the system detects the actual MAC address of the host.

NETBIOS Name Change

This event is generated when the system detects a change to a host's NetBIOS name. This event will only be generated for hosts using the NetBIOS protocol.

New Client

This event is generated when the system detects a new client.

IMPORTANT! To collect and store client data for analysis, make sure that you enable application detection in your discovery rules in the network discovery policy. For more information, see [Understanding Application Detection](#) on page 1263.

New Host

This event is generated when the system detects a new host running on the network.

If you select the **Discover** option and select **Hosts** in a network discovery rule where a NetFlow device is selected, this event is also generated when a device processes NetFlow data that involves a new host.

New Network Protocol

This event is generated when the system detects that a host is communicating with a new network protocol (IP, ARP, and so on).

New OS

This event is generated when the system either detects a new operating system for a host, or a change in a host's operating system.

New TCP Port

This event is generated when the system detects a new TCP server port (for example, a port used by SMTP or web services) active on a host. Note that this event is not used to identify the application protocol or the server associated with it; that information is transmitted in the TCP Server Information Update event.

If you select the **Discover** option and select **Applications** in a network discovery rule for NetFlow data, this event is also generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map.

New Transport Protocol

This event is generated when the system detects that a host is communicating with a new transport protocol, such as TCP or UDP.

New UDP Port

This event is generated when the system detects a new UDP server port running on a host.

If you select the **Discover** option and select **Applications** in a network discovery rule for NetFlow data, this event is also generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map.

TCP Port Closed

This event is generated when the system detects that a TCP port has closed on a host.

TCP Port Timeout

This event is generated when the system has not detected activity from a TCP port within the interval defined in the system's network discovery policy. See [Configuring Data Storage](#) on page 1293 for information about configuring the server timeout value.

TCP Server Information Update

This event is generated when the system detects a change in a discovered TCP server running on a host.

This event may be generated if a TCP server is upgraded.

UDP Port Closed

This event is generated when the system detects that a UDP port has closed on a host.

UDP Port Timeout

This event is generated when the system has not detected activity from a UDP port within the interval defined in the network discovery policy. See [Configuring Data Storage](#) on page 1293 for information about configuring the server timeout value.

UDP Server Information Update

This event is generated when the system detects a change in a discovered UDP server running on a host.

This event may be generated if a UDP server is upgraded.

VLAN Tag Information Update

This event is generated when the system detects a change in the VLAN tag attributed to a host. For more information about VLAN tags, see [Working with VLAN Tags in the Host Profile](#) on page 1357.

Understanding Host Input Event Types

LICENSE: FireSIGHT

There are many types of host input events. For example, the system generates and logs an Add Host event when a user adds a host using the host import feature. When you view a table of discovery events, the event type is listed in the **Event** column. For more information, see [Viewing Discovery and Host Input Events](#) on page 1396.

Contrast host input events, which are generated when a user takes a specific action (such as manually adding a host), with discovery events, which are generated when the system itself detects a change in your monitored network (such as detecting traffic from a previously undetected host). For more information on host input events, see [Understanding Discovery Event Types](#) on page 1389.

You can configure the types of host input events that the system logs by modifying your network discovery policy. By default, the system logs all types of host input events. For more information, see [Configuring Database Event Limits](#) on page 1978.

If you understand the information the different types of host input events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in correlation policies. In addition, knowing the names of the event types can help you craft more effective event searches. Descriptions of the different types of host input events follow.

Add Client

This event is generated when a user adds a client.

Add Host

This event is generated when a user adds a host.

Add Protocol

This event is generated when a user adds a protocol.

Add Scan Result

This event is generated when the system adds the results of an Nmap scan to a host.

Add Port

This event is generated when a user adds a server port.

Delete Client

This event is generated when a user deletes a client from the system.

Delete Host/Network

This event is generated when a user deletes an IP address or subnet from the system.

Delete Protocol

This event is generated when a user deletes a protocol from the system.

Delete Port

This event is generated when a user deletes a server port or group of server ports from the system.

Host Attribute Add

This event is generated when a user creates a new host attribute.

Host Attribute Delete

This event is generated when a user deletes a user-defined host attribute.

Host Attribute Delete Value

This event is generated when a user deletes a value assigned to a host attribute.

Host Attribute Set Value

This event is generated when a user sets a host attribute value for a host.

Host Attribute Update

This event is generated when a user changes the definition of a user-defined host attribute.

Set Host Criticality

This event is generated when a user sets or modifies the host criticality value for a host.

Set Operating System Definition

This event is generated when a user sets the operating system for a host.

Set Server Definition

This event is generated when a user sets the vendor and version definitions for a server.

Set Vulnerability Impact Qualification

This event is generated when a vulnerability impact qualification is set.

When a vulnerability is disabled at a global level from being used for impact qualifications, or when a vulnerability is enabled at a global level, this event is generated.

Vulnerability Set Invalid

This event is generated when a user invalidates (or reviews) a vulnerability or vulnerabilities.

Vulnerability Set Valid

This event is generated when a user validates a vulnerability that was previously marked as invalid.

Viewing Discovery and Host Input Events

LICENSE: FireSIGHT

Both discovery events and host input events can be viewed using Discovery Events workflows. Discovery events record the detection of network discovery data based on the configured network discovery policy for an appliance. Host input events record the input of host data into the network map through the host input feature. For more information, see [Understanding Discovery Event Types](#) on page 1389 and [Understanding Host Input Event Types](#) on page 1394.

You can use the Defense Center to view a table of discovery or host input events. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of discovery events and a terminating host view page. You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The [Discovery Event Actions](#) table below describes some of the specific actions you can perform on a discovery events workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Discovery Event Actions

To...	You CAN...
modify the time and date range for displayed events	find more information in Setting Event Time Constraints on page 1820. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
learn more about the contents of the columns in the table	find more information in Understanding the Discovery Events Table on page 1397.

To view discovery events:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Hosts > Discovery Events**.

The first page of the default discovery events workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

Understanding the Discovery Events Table

LICENSE: FireSIGHT

The system generates discovery events that communicate the details of changes in your monitored network segments. *New* events are generated for newly discovered network features, and change events are generated for any change in previously identified network assets.

During its initial network discovery phase, the system generates new events for each host and any TCP or UDP servers it discovers on each host. In addition, the system generates new events for each network, transport, or application protocol running on each discovered host. For NetFlow-related traffic, you can control whether the system generates new events when it detects application protocols running on a host. After the initial network mapping is complete, the system continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered host, server, or client changes.

Descriptions of the fields in the discovery events table follow.

Time

The time that the system generated the event.

Event

The event type. See [Understanding Discovery Event Types](#) on page 1389 and [Understanding Host Input Event Types](#) on page 1394 for a description of each available event.

IP Address

The IP address associated with the host involved in the event.

User

The last user to log into the host involved in the event before the event was generated. If only non-authoritative users log in after an authoritative user, the authoritative user remains the current user for the host unless another authoritative user logs in.

MAC Address

The MAC address of the NIC used by the network traffic that triggered the discovery event. This MAC address can be either the actual MAC address of the host involved in the event, or the MAC address of a network device that the traffic passed through.

MAC Vendor

The MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.

Port

The port used by the traffic that triggered the event, if applicable.

Description

The text description of the event.

Device

The name of the device that generated the event. For new host and new server events based on NetFlow data, this is the device that processed the NetFlow data.

Searching for Discovery Events

LICENSE: FireSIGHT

You can search for specific discovery events. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify `n/a` or `blank` in the field to identify events where information is not available for that field; use `!n/a` or `!blank` to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Special Search Syntax for Discovery Events

The following table notes search information specific to particular discovery event fields. For more information on discovery event fields, see [Understanding the Hosts Table](#) on page 1403.

Discovery Event Search Criteria Notes

FIELD	SEARCH CRITERIA NOTES
Event	The range of event names is listed in Understanding Discovery Event Types on page 1389 and Understanding Host Input Event Types on page 1394
MAC Vendor	To search for virtual MAC vendors, that is, for events that involve virtual machines, type <code>virtual_mac_vendor</code> . To search for a vendor whose name includes a comma, enclose the entire search term in quotes. Otherwise, the Defense Center treats the term as two searches and returns events that match each search term.
Port	Note that you cannot: <ul style="list-style-type: none">• enter a port/protocol combination as you can when searching for other kinds of events• use spaces when specifying port numbers or ranges.

To search for discovery events:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Discovery Events**.
The page reloads with the appropriate constraints.



3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.

4. Enter your search criteria in the appropriate fields, as described in [General Search Syntax](#) on page 1399 and [Special Search Syntax for Discovery Events](#) on page 1400.

If you enter multiple criteria, the search returns only the records that match all the criteria. Click the add icon (+) that appears next to a search field to use an object as a search criterion.

5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:

- Click **Search** to start the search.

Your search results appear in the default discovery events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click (**switch workflow**). For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

- Click **Save** if you are modifying an existing search and want to save your changes.
- Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**) so that you can run it at a later time.

Working with Hosts

LICENSE: FireSIGHT

The system generates an event when it detects a host and collects information about it to build the host profile. You can use the Defense Center web interface to view, search, and delete hosts.

While viewing hosts, you can create traffic profiles and compliance white lists based on selected hosts. You can also assign host attributes, including host criticality values (which designate business criticality) to groups of hosts. You can then use these criticality values, white lists, and traffic profiles within correlation rules and policies.

Although you can configure the network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

For more information, see the following sections:

- [Viewing Hosts](#) on page 1402
- [Understanding the Hosts Table](#) on page 1403
- [Creating a Traffic Profile for Selected Hosts](#) on page 1407
- [Creating a Compliance White List Based on Selected Hosts](#) on page 1408
- [Searching for Hosts](#) on page 1408
- [Setting Host Attributes for Selected Hosts](#) on page 1415

Viewing Hosts

LICENSE: FireSIGHT

You can use the Defense Center to view a table of hosts that the system has detected. Then, you can manipulate the view depending on the information you are looking for.

The page you see when you access hosts differs depending on the workflow you use. Both predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs. For more information, see [Creating Custom Workflows](#) on page 1841.

The [Host Actions](#) table below describes some of the specific actions you can perform on an hosts workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Host Actions

To...	YOU CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Hosts Table on page 1403.
assign a host attribute to selected hosts	find more information in Setting Host Attributes for Selected Hosts on page 1415.
create traffic profiles for selected hosts	find more information in Creating a Traffic Profile for Selected Hosts on page 1407.
create a compliance white list based on selected hosts	find more information in Creating a Compliance White List Based on Selected Hosts on page 1408.

To view hosts:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Hosts > Hosts**.

The first page of the default hosts workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of hosts, click **(switch workflow)**, then select **Hosts**.

Understanding the Hosts Table

LICENSE: FireSIGHT

When the system discovers a host, it collects data about that host. That data can include the host's IP addresses, the operating system it is running, and more. You can view some of that information in the table view of hosts. For more information on the data that the system collects about detected hosts, see [Using Host Profiles](#) on page 1334.

Descriptions of the fields in the hosts table follow below.

Although you can configure the network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Last Seen

The date and time any of the host's IP addresses was last detected by the system. The Last Seen value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system generates a new host event for any of the host's IP addresses.

For hosts with operating system data updated using the host input feature, the Last Seen value indicates the date and time when the data was originally added.

IP Address

The IP addresses associated with the host.

MAC Address

The host's detected MAC address of the NIC.

The MAC Address field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Address field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

MAC Vendor

The host's detected MAC hardware vendor of the NIC.

The MAC Vendor field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Vendor field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-specified criticality value assigned to the host. See the description of the Host Criticality column in [Understanding the Host Attributes Table](#) on page 1413 for more information about this field.

NetBIOS Name

The NetBIOS name of the host. Only hosts running the NetBIOS protocol will have a NetBIOS name.

VLAN ID

VLAN ID used by the host. For more detailed information about VLAN IDs, see [Working with VLAN Tags in the Host Profile](#) on page 1357.

Hops

The number of network hops from the device that detected the host to the host.

Host Type

The type of host (host, mobile device, jailbroken mobile device, router, bridge, NAT device, or load balancer). The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a device is not identified as a network device, it is categorized as a host.

Hardware

The hardware platform for a mobile device.

OS

The detected operating system (name, vendor, and version) running on the host, or updated using Nmap or the host input feature. This field appears when you invoke the hosts event view from the Custom Analysis widget on the dashboard. It is also a field option in custom tables based on the Hosts table.

Note if the system detects multiple identities, it displays those identities in a comma-separated list.

In this field, a value of **unknown** means that the operating system does not match any of the known fingerprints. A value of **pending** means that the system has not yet gathered enough information to identify the operating system.

OS Vendor

The vendor of the operating system detected on the host or updated using Nmap or the host input feature.

Note if the system detects multiple vendors, it displays those vendors in a comma-separated list.

In this field, a value of **unknown** means that the operating system does not match any of the known fingerprints. A value of **pending** means that the system has not yet gathered enough information to identify the operating system.

OS Name

The detected operating system running on the host or updated using Nmap or the host input feature.

Note if the system detects multiple names, it displays those names in a comma-separated list.

In this field, a value of **unknown** means that the operating system does not match any of the known fingerprints. A value of **pending** means that the system has not yet gathered enough information to identify the operating system.

OS Version

The version of the operating system detected on the host or updated using Nmap or the host input feature.

Note if the system detects multiple versions, it displays those versions in a comma-separated list.

In this field, a value of **unknown** means that the operating system does not match any of the known fingerprints. A value of **pending** means that the system has not yet gathered enough information to identify the operating system.

Source Type

One of the following values for the source of the host's operating system identity:

- User: *user_name*
- Application: *app_name*
- Scanner: *scanner_type* (Nmap or scanner added through network discovery configuration)
- FireSIGHT, for operating systems detected by the system

The system may reconcile data from multiple sources to determine the identity of an operating system; see [Understanding Current Identities](#) on page 1646.

Confidence

One of:

- the percentage of confidence that the system has in the identity of the operating system running on the host, for hosts detected by the system
- 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner
- **unknown**, for hosts for which the system cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data

Notes

The user-defined content of the Notes host attribute.

Device

Either:

- the managed device that detected the traffic or
- the device that processed the NetFlow or host input data that added the host to the network map

If this field is blank, either:

- the host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy, or
- the host was added using the host input feature and has not also been detected by the system

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Creating a Traffic Profile for Selected Hosts

LICENSE: FireSIGHT

A traffic profile is a profile of the traffic on your network, based on connection data collected over a timespan that you specify. After you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

You can use the Hosts page to create a traffic profile for a group of hosts that you specify. The traffic profile will be based on connections detected where one of the hosts you specify is the initiating host. Use the sort and search features to isolate the hosts for which you want to create a profile.

To create a traffic profile for selected hosts:

ACCESS: Admin

1. On a table view in the hosts workflow, select the check boxes next to the hosts for which you want to create a traffic profile.
2. At the bottom of the page, click **Create Traffic Profile**.

The Create Profile page appears, populated with the IP addresses of the hosts you specified as the hosts to be monitored.

3. Modify and save the traffic profile according to your specific needs.
For more information on creating traffic profiles, see [Creating Traffic Profiles](#) on page 1585.

Creating a Compliance White List Based on Selected Hosts

LICENSE: FireSIGHT

Compliance white lists allow you to specify which operating systems, clients, and network, transport, or application protocols are allowed on your network.

You can use the Hosts page to create a compliance white list based on the host profiles of a group of hosts that you specify. Use the sort and search features to isolate the hosts that you want to use to create a white list.

To create a compliance white list based on selected hosts:

ACCESS: Admin

1. On a table view in the hosts workflow, select the check boxes next to the hosts for which you want to create a white list.
2. At the bottom of the page, click **Create White List**.
The Create White List page appears, populated with the information in the host profiles of the hosts you specified.
3. Modify and save the white list according to your specific needs.
For more information on creating compliance white lists, see [Creating Compliance White Lists](#) on page 1541.

Searching for Hosts

LICENSE: FireSIGHT

You can search for specific hosts by using one of the predefined searches or by using your own search criteria.

When searching for hosts, you should keep in mind that although you can configure the network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

You can search for specific discovery events. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify `n/a` or `blank` in the field to identify events where information is not available for that field; use `!n/a` or `!blank` to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IP addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

IMPORTANT! When you search for hosts by IP address, the results include all hosts for which at least one IP address matches your search conditions, that is, a search for an IPv6 address may return hosts whose primary address is IPv4.

- When you search hosts by IP address,
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Special Search Syntax for Hosts

The following table notes search information specific to particular host fields. For more information on host fields, see [Understanding the Hosts Table](#) on page 1403.

Host Search Criteria

FIELD	SEARCH CRITERIA NOTES
Host Type	To search for all network devices, type <code>!host</code> .
MAC Vendor	To search for virtual MAC vendors, that is, for events that involve virtual machines, type <code>virtual_mac_vendor</code> . To search for a vendor whose name includes a comma, enclose the entire search term in quotes. Otherwise, the Defense Center treats the term as two searches and returns events that match each search term.
OS Vendor/Name/Version	Type <code>unknown</code> to search for hosts where the operating system is unknown. Type <code>n/a</code> to search for hosts where the operating system has not yet been identified.
Confidence	You can precede the confidence with greater than (<code>></code>), greater than or equal to (<code>>=</code>), less than (<code><</code>), less than or equal to (<code><=</code>), or equal to (<code>=</code>) operators. Matches to an <code>n/a</code> search include hosts added to the network map based on NetFlow data.
OS Conflict	Note that the OS Conflict column does not appear in search results. To determine whether you are viewing hosts with or without operating system conflicts, expand the search constraints on the workflow page. For more information on resolving operating system conflicts, see Resolving Operating System Identity Conflicts on page 1345.

For more information on searching, including how to load and delete saved searches, see [Searching for Events](#) on page 1770.

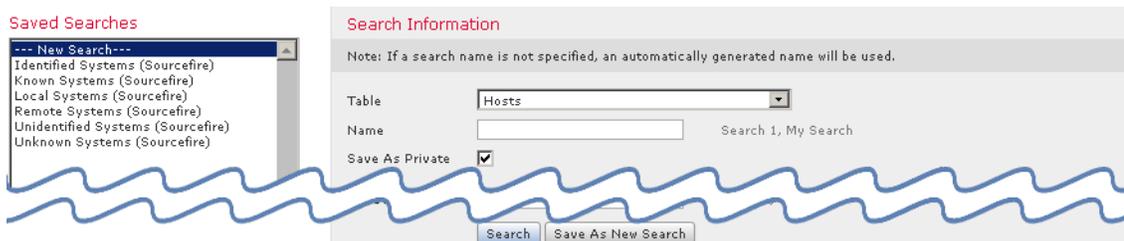
To search for hosts:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.

2. From the **Table** drop-down list, select **Hosts**.

The page reloads with the appropriate constraints.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.

If you do not enter a name, the Defense Center automatically creates one when you save the search.

4. Enter your search criteria in the appropriate fields, as described in the [Host Search Criteria](#) table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria. Click the add icon (+) that appears next to a search field to use an object as a search criterion.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default hosts workflow. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Host Attributes

LICENSE: FireSIGHT

The Sourcefire 3D System collects information about the hosts it detects and uses that information to build host profiles. However, there may be additional information about the hosts on your network that you want to provide to your analysts. You can add notes to a host profile, set the business criticality of a host, or provide any other information that you choose. Each piece of information is called a host attribute.

You can use host attributes in host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also set attribute values in response to a correlation rule.

For more information, see:

- [Viewing Host Attributes](#) on page 1412
- [Understanding the Host Attributes Table](#) on page 1413
- [Setting Host Attributes for Selected Hosts](#) on page 1415
- [Searching for Host Attributes](#) on page 1416
- [Configuring Set Attribute Remediations](#) on page 1629

Viewing Host Attributes

LICENSE: FireSIGHT

You can use the Defense Center to view a table of hosts detected by the system, along with their host attributes. Then, you can manipulate the view depending on the information you are looking for.

The page you see when you access host attributes differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of host attributes that lists all detected hosts and their attributes, and terminates in a host view page, which contains a host profile for every host that meets your constraints.

You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The [Host Attribute Actions](#) table below describes some of the specific actions you can perform on a host attributes workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Host Attribute Actions

To...	YOU CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Host Attributes Table on page 1413.
assign a host attribute to selected hosts	find more information in Setting Host Attributes for Selected Hosts on page 1415

To view host attributes:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Hosts > Host Attributes**.

The first page of the default host attributes workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of host attributes, click **(switch workflow)**, then select **Attributes**.

Understanding the Host Attributes Table

LICENSE: FireSIGHT

The Sourcefire 3D System collects information about the hosts it detects and uses that information to build host profiles. However, there may be additional information about the hosts on your network that you want to provide to your analysts. You can add notes to a host profile, set the business criticality, or provide any other information that you choose. Each piece of information is called a host attribute.

You can use host attributes in host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule.

Note that the host attributes table does not display hosts identified only by MAC addresses.

For more information on host attributes, see [Working with the Predefined Host Attributes](#) on page 1369 and [Working with User-Defined Host Attributes](#) on page 1370.

Descriptions of the fields in the host attributes table follow.

IP Address

The IP addresses associated with a host.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-assigned importance of a host to your enterprise. You can use the host criticality in correlation rules and policies to tailor policy violations and their responses to the importance of a host involved in an event. You can assign a host criticality of low, medium, high, or none.

For information on setting a host's criticality, see [Working with the Predefined Host Attributes](#) on page 1369 and [Setting Host Attributes for Selected Hosts](#) on page 1415.

Notes

Information about the host that you want other analysts to view. For information on how to add a note, see [Working with the Predefined Host Attributes](#) on page 1369.

Any user-defined host attribute, including those for compliance white lists

The value of the user-defined host attribute.

The host attributes table contains a field for each user-defined host attribute. For more information, see [Working with User-Defined Host Attributes](#) on page 1370.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Setting Host Attributes for Selected Hosts

LICENSE: FireSIGHT

There are two predefined host attributes that you can assign to each host: host criticality and host-specific notes.

Use the host criticality to designate the business criticality of a given host. You can tailor correlation policies and alerts based on host criticality. For example, your organization's mail servers are more critical to your business than a typical user workstation. You can assign a high host criticality value to your mail servers and other business-critical servers and medium or low values to other hosts. You could then create a correlation policy that launches different alerts based on the criticality of an affected host.

Use notes to record information about a host that you want other analysts to view. For example, if you have a computer on the network that has an older, unpatched version of an operating system that you use for testing, you can use the notes feature to indicate that the system is intentionally unpatched.

You can also create user-defined host attributes. For example, you could create a host attribute that assigns physical location identifiers to hosts, such as a facility code, city, or room number. For more information on created user-defined host attributes, see [Creating User-Defined Host Attributes](#) on page 1372.

You can also set the host criticality of selected hosts in a host workflow, and from within a host profile, or set it through a remediation. For more information, see [Working with the Predefined Host Attributes](#) on page 1369 or [Configuring Set Attribute Remediations](#) on page 1629.

To set host attributes for selected hosts:

ACCESS: Admin/Any Security Analyst

1. Select the check boxes next to the hosts to which you want to add a host attribute.

TIP! Use the sort and search features to isolate the hosts to which you want to assign particular attributes.

2. At the bottom of the page, click **Set Attributes**.
The Host Attributes pop-up window appears.



The screenshot shows a pop-up window titled "Host Attributes:". Inside the window, there is a "Host Criticality" label followed by a dropdown menu. Below that is a "Notes" label followed by a large text area. At the bottom of the window is a "Save" button.

3. Optionally, set the host criticality for the hosts you selected.
You can select **None**, **Low**, **Medium**, or **High**.
4. Optionally, add notes to the host profiles of the hosts you selected by entering up to 255 alphanumeric characters, special characters, and spaces in the text box.
5. Optionally, set any user-defined host attributes you have configured.
6. Click **Save**.
The host attributes you specified are assigned to the selected hosts.

Searching for Host Attributes

LICENSE: FireSIGHT

You can search for hosts that have specific host attributes. For example, if your company has several regional offices, you could configure a host attribute that tells you which city any one host resides in. You could then search for hosts in specific regions. For more information on host attributes, see [Working with User-Defined Host Attributes](#) on page 1370.

You may want to create searches customized for your network environment, then save them to reuse later. For more information on the host attribute fields, see [Understanding the Host Attributes Table](#) on page 1413.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify **n/a** or **blank** in the field to identify events where information is not available for that field; use **!n/a** or **!blank** to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

To search for host attributes:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Host Attributes**.
The page reloads with the appropriate constraints.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, the Defense Center automatically creates one when you save the search.
4. Enter your search criteria in the appropriate fields, as described in [Understanding the Host Attributes Table](#) on page 1413. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria. Click the add icon (+) that appears next to a search field to use an object as a search criterion.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default host attributes workflow. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

- Click **Save** if you are modifying an existing search and want to save your changes.
- Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Servers

LICENSE: FireSIGHT

The Sourcefire 3D System collects information about all servers running on hosts on monitored network segments. The information that the system collects includes the name of the server, the application and network protocols used by the server, the vendor and version of the server, the IP address associated with the host running a server, and the port on which the server communicates.

When the system detects a server, it generates a discovery event unless the associated host has already reached its maximum number of servers. For more information, see [Host Limits and Discovery Event Logging](#) on page 1268. You can use the Defense Center web interface to view, search, and delete server events.

You can also base correlation rules on server events. For example, you could trigger a correlation rule when the system detects a chat server, such as ircd, running on one of your hosts.

Although you can configure the network discovery policy to add servers to the network map based on application data exported by NetFlow-enabled devices, the available information about these servers is limited. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

See the following sections for more information:

- [Viewing Servers](#) on page 1418
- [Understanding the Servers Table](#) on page 1419
- [Searching for Servers](#) on page 1422
- [Editing Server Identities](#) on page 1352

Viewing Servers

LICENSE: FireSIGHT

You can use the Defense Center to view a table of detected servers. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access servers differs depending on the workflow you use. All the predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs. For more information, see [Creating Custom Workflows](#) on page 1841.

The [Server Actions](#) table below describes some of the specific actions you can perform on an servers workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Server Actions

To...	You CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Servers Table on page 1419.
edit server identities	select the check boxes next to the events for servers you want to edit, then click Set Server Identity . For more information, see Editing Server Identities on page 1352.

To view servers:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Hosts > Servers**.

The first page of the default servers workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of servers, click **(switch workflow)**, then select **Servers**.

Understanding the Servers Table

LICENSE: FireSIGHT

The Sourcefire 3D System collects information about servers running on hosts on monitored network segments.

Descriptions of the fields in the servers table follow below.

Although you can configure the network discovery policy to add servers to the network map based on data exported by NetFlow-enabled devices, the available information about these servers is limited. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Last Used

The date and time the server was last used on the network or the date and time that the server was originally updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects a server information update. For information on setting the update interval, see [Configuring Data Storage](#) on page 1293.

IP Address

The IP address associated with the host running the server.

Port

The port where the server is running.

Protocol

The network or transport protocol used by the server.

Application Protocol

The application protocol, as indicated by one of the following:

- the name of the application protocol for the server
- **pending**, if the system cannot positively or negatively identify the server for one of several reasons
- **unknown**, if the system cannot identify the server based on known server fingerprints or if the server was added through host input and did not include the application protocol

Category, Tags, Risk, or Business Relevance for Application Protocols

The categories, tags, risk level, and business relevance assigned to the application protocol. These filters can be used to focus on a specific set of data.

For more information, see the [Application Characteristics table](#) on page 1264.

Vendor

One of:

- the server vendor as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its vendor based on known server fingerprints, or if the server was added to the network map using NetFlow data

Version

One of:

- the server version as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its version based on known server fingerprints, or if the server was added to the network map using NetFlow data

Web Application

The web application based on the payload content detected by the system in the http traffic. Note that if the system detects an application protocol of **HTTP** but cannot detect a specific web application, the system supplies a generic web browsing designation.

Category, Tags, Risk, or Business Relevance for Web Applications

The categories, tags, risk level, and business relevance assigned to the web application. These filters can be used to focus on a specific set of data.

For more information, see the [Application Characteristics table](#) on page 1264.

Hits

The number of times the server was accessed. For servers added using the host input feature, this value is always 0.

Source Type

One of the following values:

- User: *user_name*
- Application: *app_name*
- Scanner: *scanner_type* (Nmap or scanner added through network discovery configuration)
- FireSIGHT, FireSIGHT Port Match, or FireSIGHT Pattern Match, for servers detected by the Sourcefire 3D System
- NetFlow, for servers added to the network map based on NetFlow data

The system may reconcile data from multiple sources to determine the identity of a server; see [Understanding Current Identities](#) on page 1646.

Device

The name of the device that either detected the server or processed the NetFlow or host input data that added the server to the network map.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for Servers

LICENSE: FireSIGHT

You can search for specific servers that are running on monitored hosts by using one of the predefined searches or by using your own search criteria. The predefined searches serve as examples and can provide quick access to important information about your network.

You may want to modify specific fields within the default searches to customize them for your network environment, then save them to reuse later. The fields you can use to retrieve data are described in the [Understanding the Servers Table](#) on page 1419.

When searching for servers, you should keep in mind that although you can configure the network discovery policy to add applications, including servers, to the network map based on data exported by NetFlow-enabled devices, the available information about these servers is limited. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify `n/a` or `blank` in the field to identify events where information is not available for that field; use `!n/a` or `!blank` to identify the events where that field is populated.
- Most fields are case-insensitive.

- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
- Click the add object icon (⊕) that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

To search for servers:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Servers**.
The page reloads with the appropriate constraints.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, the Defense Center automatically creates one when you save the search.
4. Enter your search criteria in the appropriate fields. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria. Click the add icon (⊕) that appears next to a search field to use an object as a search criterion.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default servers workflow. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Applications

LICENSE: FireSIGHT

When a monitored host connects to another host, the system can, in many cases, determine what application was used. The Sourcefire 3D System detects the use of many email, instant messaging, peer to peer, web applications, as well as other types of applications.

For each detected application, the system logs the IP address that used the application, the product, the version, and the number of times its use was detected. You can use the web interface to view, search, and delete application events. You can also update application data on a host or hosts using the host input feature.

If you know which applications are running on which hosts, you can use that knowledge to create host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also base correlation rules on the detection of application. For example, if you want your employees to use a specific mail client, you could trigger a correlation rule when the system detects a different mail client running on one of your hosts.

You should carefully read the release notes for each Sourcefire 3D System update as well as the advisories for each VDB update for information on updated detectors.

To collect and store application data for analysis, make sure that you enable application detection in your network discovery policy. For more information, see [Understanding Discovery Data Collection](#) on page 1251.

See the following sections for more information:

- [Viewing Application Details](#) on page 1429
- [Understanding the Application Detail Table](#) on page 1430
- [Searching for Application Details](#) on page 1432

Viewing Applications

LICENSE: FireSIGHT

You can use the Defense Center to view a table of detected applications. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access applications differs depending on the workflow you use. You can also create a custom workflow that displays only the information that matches your specific needs. For more information, see [Creating Custom Workflows](#) on page 1841.

The [Application Actions](#) table below describes some of the specific actions you can perform on an application workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Application Actions

To...	You CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Applications Table on page 1425.
open the Application Detail View for a specific application	click the application detail view icon () next to a client, application protocol, or web application.

To view applications:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Hosts > Application Details**.

The first page of the default application details workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then select **Clients**.

Understanding the Applications Table

LICENSE: FireSIGHT

When a monitored host connects to another host, the Sourcefire 3D System can, in many cases, determine what application was used. The system detects various web browsers or servers, email clients or servers, instant messengers,

peer-to-peer applications, and so on. When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

The Sourcefire 3D System classifies application data into three types: client, web application, and application protocol. The applications table provides a list combining all three types of detected applications on the appliance.

Descriptions of the fields in the applications table follow.

Application

The name of the detected application.

IP Address

The IP address associated with the host using the application.

Category

A general classification for the application that describes its most essential function. Each application belongs to at least one category.

Tag

Additional information about the application. Applications can have any number of tags, including none.

Risk

How likely the application is to be used for purposes that might be against your organization's security policy. An application's risk can range from **Very Low** to **Very High**.

Of Application Protocol Risk, Client Risk, and Web Application Risk, the highest of the three detected, when available, in the traffic that triggered the intrusion event.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from **Very Low** to **Very High**.

Of Application Protocol Business Relevance, Client Business Relevance, and Web Application Business Relevance, the lowest of the three detected, when available, in the traffic that triggered the intrusion event.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Type

The type of application:

- **Application Protocols** represent communications between hosts.
- **Client Applications** represent software running on a host.
- **Web Applications** represent the content or requested URL for HTTP traffic.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for Applications

LICENSE: FireSIGHT

You can search for hosts that are running specific clients, application protocols, or web applications. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify **n/a** or **blank** in the field to identify events where information is not available for that field; use **!n/a** or **!blank** to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

To search for applications:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Applications**.
The page reloads with the appropriate constraints.



3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, the Defense Center automatically creates one when you save the search.
4. Enter your search criteria in the appropriate fields. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria. Click the add icon () that appears next to a search field to use an object as a search criterion.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default clients workflow. To use a different workflow, including a custom workflow, click (**switch workflow**). For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

- Click **Save** if you are modifying an existing search and want to save your changes.
- Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Application Details

LICENSE: FireSIGHT

When a monitored host connects to another host, the system can, in many cases, determine what application was used. The Sourcefire 3D System detects the use of many email, instant messaging, peer to peer, web applications, as well as other types of applications.

For each detected application, the system logs the IP address that used the application, the product, the version, and the number of times its use was detected. You can use the web interface to view, search, and delete application events. You can also update application data on a host or hosts using the host input feature.

If you know which applications are running on which hosts, you can use that knowledge to create host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also base correlation rules on the detection of application. For example, if you want your employees to use a specific mail client, you could trigger a correlation rule when the system detects a different mail client running on one of your hosts.

You should carefully read the release notes for each Sourcefire 3D System update as well as the advisories for each VDB update for information on updated detectors.

To collect and store application data for analysis, make sure that you enable application detection in your network discovery policy. For more information, see [Understanding Application Detection](#) on page 1263.

See the following sections for more information:

- [Viewing Application Details](#) on page 1429
- [Understanding the Application Detail Table](#) on page 1430
- [Searching for Application Details](#) on page 1432

Viewing Application Details

LICENSE: FireSIGHT

You can use the Defense Center to view a table of detected application details. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access application details differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs. For more information, see [Creating Custom Workflows](#) on page 1841.

The [Application Details Actions](#) table below describes some of the specific actions you can perform on an application details workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Application Details Actions

To...	You CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Application Detail Table on page 1430.
open the Application Detail View for a specific application	click the application detail view icon () next to a client.

To view application details:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Hosts > Application Details**.

The first page of the default application details workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then select **Clients**.

Understanding the Application Detail Table

LICENSE: FireSIGHT

When a monitored host connects to another host, the Sourcefire 3D System can, in many cases, determine what application was used. The system detects various web browsers, email clients, instant messengers, peer-to-peer applications, and so on.

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it. Descriptions of the fields in the application details table follow.

Last Used

The time that the application was last used or the time that the application data was updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects an application information update. For information on setting the update interval, see [Configuring Data Storage](#) on page 1293.

IP Address

The IP address associated with the host using the application.

Client

The name of the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Version

The version of the application.

Category, Tags, Risk, or Business Relevance for Clients

The categories, tags, risk level, and business relevance assigned to the client. These filters can be used to focus on a specific set of data.

For more information, see [Understanding Application Detection](#) on page 1263.

Application Protocol

The application protocol used by the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Category, Tags, Risk, or Business Relevance for Application Protocols

The categories, tags, risk level, and business relevance assigned to the application protocol. These filters can be used to focus on a specific set of data.

For more information, see [Understanding Application Detection](#) on page 1263.

Web Application

The web application based on the payload content or URL detected by the system in the http traffic. Note that if the system detects an application protocol of `HTTP` but cannot detect a specific web application, the system supplies a generic web browsing designation here.

Category, Tags, Risk, or Business Relevance for Web Applications

The categories, tags, risk level, and business relevance assigned to the web application. These filters can be used to focus on a specific set of data.

For more information, see [Understanding Application Detection](#) on page 1263.

Hits

The number of times the system detected the application in use. For applications added using the host input feature, this value is always 0.

Device

The device that generated the discovery event containing the application detail.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for Application Details

LICENSE: FireSIGHT

You can search for hosts that are running specific clients, application protocols, or web applications. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.

- For some fields, you can specify n/a or blank in the field to identify events where information is not available for that field; use !n/a or !blank to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

To search for application details:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Application Details**.
The page reloads with the appropriate constraints.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, the Defense Center automatically creates one when you save the search.
4. Enter your search criteria in the appropriate fields. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria. Click the add icon () that appears next to a search field to use an object as a search criterion.

5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default clients workflow. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Sourcefire Vulnerabilities

LICENSE: FireSIGHT

The Sourcefire 3D System includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network.

The operating systems, servers, and clients running on your hosts have different sets of associated vulnerabilities. You can deactivate vulnerabilities for a host after you patch the host or otherwise judge it immune to a vulnerability. You can use the Defense Center to track and review the vulnerabilities for each host.

Note that vulnerabilities for vendorless and versionless servers are not mapped unless the applications protocols used by the servers are mapped in the system policy. Vulnerabilities for vendorless and versionless clients cannot be mapped. For more information, see [Mapping Vulnerabilities for Servers](#) on page 1996.

For more information, see:

- [Viewing Sourcefire Vulnerabilities](#) on page 1435
- [Understanding the Sourcefire Vulnerabilities Table](#) on page 1436
- [Deactivating Sourcefire Vulnerabilities](#) on page 1438
- [Searching for Sourcefire Vulnerabilities](#) on page 1439

Viewing Sourcefire Vulnerabilities

LICENSE: FireSIGHT

You can use the Defense Center to view a table of vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access vulnerabilities differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of vulnerabilities. The table view contains a row for each vulnerability in the database, regardless of whether any of your detected hosts exhibit the vulnerabilities. The second page of the predefined workflow contains a row for each vulnerability (that you have not deactivated) that applies to detected hosts on your network. The predefined workflow terminates in a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.

TIP! If you want to see the vulnerabilities that apply to a single host or set of hosts, perform a search for vulnerabilities, specifying an IP address or range of IP addresses for the hosts. For more information on searching for vulnerabilities, see [Searching for Sourcefire Vulnerabilities](#) on page 1439.

You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The following table describes some of the specific actions you can perform on an vulnerabilities workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Vulnerability Actions

To...	You CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Sourcefire Vulnerabilities Table on page 1436.
view the vulnerability details for a vulnerability	click the view icon () in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page. For more information, see Viewing Vulnerability Details on page 1365.

Vulnerability Actions (Continued)

To...	You CAN...
deactivate selected vulnerabilities so they are no longer used for intrusion impact correlation for currently vulnerable hosts	find more information in Deactivating Sourcefire Vulnerabilities on page 1438.
view the full text of a vulnerability title	right-click the title and select Show Full Text .

To view vulnerabilities:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Vulnerabilities > Vulnerabilities**.

The first page of the default vulnerabilities workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of vulnerabilities, click **(switch workflow)**, then select **Vulnerabilities**.

Understanding the Sourcefire Vulnerabilities Table

LICENSE: FireSIGHT

The Sourcefire 3D System includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network.

The operating systems, servers, and clients running on your hosts have different sets of associated vulnerabilities. You can deactivate vulnerabilities for a host after you patch the host or otherwise judge it immune to a vulnerability. You can use the Defense Center to track and review the vulnerabilities for each host.

For more information on vulnerabilities, see [Working with the Vulnerabilities Network Map](#) on page 1323 and [Working with Vulnerabilities in the Host Profile](#) on page 1363.

Descriptions of the fields in the vulnerabilities table follow.

SVID

The Sourcefire vulnerability identification number that the system uses to track vulnerabilities.

Click the view icon (🔍) to access the vulnerability details for the SVID. See [Viewing Vulnerability Details](#) on page 1365 for more information.

Bugtraq ID

The identification number associated with the vulnerability in the Bugtraq database. (<http://www.securityfocus.com/bid/>)

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

Title

The title of the vulnerability.

IP Address

The IP address associated with the host affected by the vulnerability.

Date Published

The date the vulnerability was published.

Vulnerability Impact

Displays the severity assigned to the vulnerability in the Bugtraq database on a scale of 0 to 10, with 10 being the most severe. The vulnerability impact is determined by the writer of the Bugtraq entry based on his or her best judgment and guided by SANS Critical Vulnerability Analysis (CVA) criteria.

Remote

Indicates whether the vulnerability is remotely exploitable.

Available Exploits

Indicates whether there are known exploits for the vulnerability.

Description

A brief description of the vulnerability.

Technical Description

A detailed technical description of the vulnerability.

Solution

Information about repairing the vulnerability.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Deactivating Sourcefire Vulnerabilities

LICENSE: FireSIGHT

Deactivate a vulnerability after you patch the hosts on your network or otherwise judge them immune. Deactivated vulnerabilities are not used for intrusion impact correlation. Note that if the system discovers a new host that is affected by that vulnerability, the vulnerability is considered valid (and is not automatically deactivated) for that host.

You can deactivate vulnerabilities within the vulnerabilities workflow **only** on a workflow page that shows vulnerabilities for specific hosts on your network, that is:

- on the second page of the default vulnerabilities workflow, **Vulnerabilities on the Network**, which shows only the vulnerabilities that apply to the hosts on your network
- on any page in a vulnerabilities workflow, custom or predefined, that you constrained based on IP address using a search.

Deactivating a vulnerability within a vulnerabilities workflow that is not constrained on IP addresses deactivates the vulnerability for *all* detected hosts on your network. To deactivate a vulnerability for a single host, you have three options:

- Use the network map.
For more information, see [Working with the Vulnerabilities Network Map](#) on page 1323.
- Use the host's host profile.
For more information, see [Setting Vulnerabilities for Individual Hosts](#) on page 1368.
- Constrain the vulnerabilities workflow based on the IP addresses of the host or hosts for which you want to deactivate vulnerabilities. For hosts with multiple associated IP addresses, this function applies only to the single, selected IP address of that host.

To constrain the view based on IP address, perform a search for vulnerabilities, specifying an IP address or range of IP addresses for the hosts for which you want to deactivate vulnerabilities. For more information on searching for vulnerabilities, see [Searching for Sourcefire Vulnerabilities](#) on page 1439.

To deactivate vulnerabilities:

ACCESS: Admin/Any Security Analyst

- ▶ On the Vulnerabilities on the Network page, select the check boxes next to vulnerabilities you want to deactivate, then click **Review**.

Searching for Sourcefire Vulnerabilities

LICENSE: FireSIGHT

You can search for vulnerabilities that affect the hosts on your network. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify `n/a` or `blank` in the field to identify events where information is not available for that field; use `!n/a` or `!blank` to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Specific Search Criteria for Vulnerabilities

Note the following information specific to searching for vulnerabilities:

- Find Bugtraq ID numbers at <http://www.securityfocus.com/bid>.
- Enter `TRUE` to search for vulnerabilities that are exploited, or `FALSE` to exclude such vulnerabilities.

To search for vulnerabilities:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Vulnerabilities**.
The page reloads with the appropriate constraints.



3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields.
If you enter multiple criteria, the search returns only the records that match all the criteria. Click the add icon (+) that appears next to a search field to use an object as a search criterion.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.
6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default vulnerabilities workflow. To use a different workflow, including a custom workflow, click (**switch workflow**). For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Third-Party Vulnerabilities

LICENSE: FireSIGHT

The Sourcefire 3D System includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network.

If your organization can write scripts or create command line import files to import network map data from third-party applications, you can import third-party vulnerability data to augment the system's vulnerability data. For more information, see the *Sourcefire 3D System Host Input API Guide*.

To include imported data in impact correlations, you must map third-party vulnerability information to the operating system and application definitions in the database. You cannot map third-party vulnerability information to client definitions.

For more information, see:

- [Viewing Third-Party Vulnerabilities](#) on page 1441
- [Understanding the Third-Party Vulnerabilities Table](#) on page 1442
- [Searching for Third-Party Vulnerabilities](#) on page 1444

Viewing Third-Party Vulnerabilities

LICENSE: FireSIGHT

After you use the host input feature to import third-party vulnerability data, you can use the Defense Center to view a table of third-party vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access third-party vulnerabilities differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs. For more information, see [Creating Custom Workflows](#) on page 1841.

The following table describes some of the specific actions you can perform on a third-party vulnerabilities workflow page. You can also perform the tasks described in the [Common Discovery Event Actions table](#) on page 1387.

Third-Party Vulnerability Actions

To...	You CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Third-Party Vulnerabilities Table on page 1442.
view the vulnerability details for a third-party vulnerability	click the view icon (🔍) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page. For more information, see Viewing Vulnerability Details on page 1365.

To view third-party vulnerabilities:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Vulnerabilities > Third-Party Vulnerabilities**.

The first page of the default third-party vulnerabilities workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of third-party vulnerabilities, click **(switch workflow)**, then select **Vulnerabilities by Source** or **Vulnerabilities by IP Address**.

Understanding the Third-Party Vulnerabilities Table

LICENSE: FireSIGHT

When you import third-party vulnerability information using the host input feature, the system stores that information in its database. The fields in the third-party vulnerabilities table are described in the following table.

Vulnerability Source

The source of the third-party vulnerabilities, for example, QualysGuard or NeXpose.

Vulnerability ID

The ID number associated with the vulnerability for its source.

IP Address

The IP address associated with the host affected by the vulnerability.

Port

A port number, if the vulnerability is associated with a server running on a specific port.

Bugtraq ID

The identification number associated with the vulnerability in the Bugtraq database. (<http://www.securityfocus.com/bid/>)

CVE ID

The identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (<http://www.cve.mitre.org/>).

SVID

The Sourcefire Vulnerability identification number that the system uses to track vulnerabilities

Click the view icon () to access the vulnerability details for the SVID. See [Viewing Vulnerability Details](#) on page 1365 for more information.

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

Title

The title of the vulnerability.

Description

A brief description of the vulnerability.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for Third-Party Vulnerabilities

LICENSE: FireSIGHT

You can search for third-party vulnerabilities that affect the hosts on your network. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify **n/a** or **b1ank** in the field to identify events where information is not available for that field; use **!n/a** or **!b1ank** to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Specific Search Criteria for Vulnerabilities

Note the following information specific to searching for vulnerabilities:

- Find Bugtraq ID numbers at <http://www.securityfocus.com/bid>.
- Enter **TRUE** to search for vulnerabilities that are exploited, or **FALSE** to exclude such vulnerabilities.

To search for third-party vulnerabilities:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.

- From the **Table** drop-down list, select **Third-Party Vulnerabilities**.
The page reloads with the appropriate constraints.



- Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
- Enter your search criteria in the appropriate fields.
If you enter multiple criteria, the search returns only the records that match all the criteria. Click the add icon (+) that appears next to a search field to use an object as a search criterion.
- If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.
- You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default third-party vulnerabilities workflow. To use a different workflow, including a custom workflow, click (**switch workflow**). For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with Users

LICENSE: FireSIGHT

When either an Active Directory Agent or a managed device detects a user login for a user who is not already in the database, the user is added to the database, unless you have specifically restricted that login type (see [Restricting User](#)

[Logging](#) on page 1285).

IMPORTANT! Although the system detects SMTP logins, the system does not record them unless there is already a user with a matching email address in the database; users are **not** added to the database based on SMTP logins.

The type of login that the system detected determines what information is stored about the new user, as described in the following table.

Login Types and User Data Stored

LOGIN TYPE	USER DATA STORED
LDAP AIM Oracle SIP	<ul style="list-style-type: none">• username• current IP address• login type (<code>aim</code>, <code>ldap</code>, <code>oracle</code>, or <code>sip</code>)
POP3 IMAP	<ul style="list-style-type: none">• username• current IP address• email address• login type (<code>pop3</code> or <code>imap</code>)

If you configured Defense Center-LDAP server connections, the Defense Center queries the LDAP servers every five minutes and obtains metadata for the new users in the user database. At the same time, the Defense Center also queries the LDAP servers for updated information on users whose records in the Defense Center database are more than 12 hours old. It may take five to ten minutes for the Defense Center database to update with user metadata after the system detects a new user login. From the LDAP servers, the Defense Center obtains the following information and metadata about each user:

- LDAP username
- first and last names
- email address
- department
- telephone number

The number of users the Defense Center can store in its database depends on your FireSIGHT license. Note that AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that the system obtains from LDAP servers. To prevent overuse of user count because of duplicate user records from these protocols, disable logging of the protocols in the network discovery policy. For more information, see [Restricting User Logging](#) on page 1285.

You can search, view, and delete users from the database; you can also purge all users from the database. For more information, see the following sections:

- [Viewing Users](#) on page 1447
- [Understanding the Users Table](#) on page 1447
- [Understanding User Details and Host History](#) on page 1449
- [Searching for Users](#) on page 1451

Viewing Users

LICENSE: FireSIGHT

You can view a table of users, and then manipulate the event view depending on the information you are looking for.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

For more information about the contents of the columns in the table, see [Understanding the Users Table](#) on page 1447. The following table, see describes some of the specific actions you can perform on an users workflow page. You can also perform the actions in the [Common Discovery Event Actions table](#) on page 1387.

To view users:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Users > Users**.

The first page of the default users workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of users, click **(switch workflow)**, then select **Users**.

Understanding the Users Table

LICENSE: FireSIGHT

When the system discovers a user, it collects data about that user and stores it in the database. Descriptions of the fields in the users table follow.

User

One of:

- the first name, last name, and username of the user as collected via the optional Defense Center-LDAP server connections
- the username only, if you have not configured Defense Center-LDAP server connections, or for users that the Defense Center cannot correlate with an LDAP record

The Defense Center also displays the protocol used to detect the user.

Note that because unsuccessful AIM login attempts are recorded, the Defense Center can store invalid AIM users (for example, if a user misspelled his or her username).

Current IP

The IP address associated with the host that the user is logged into. This field is blank if another authoritative user logs into the host with the same IP address after the user's login, unless the user is an authoritative user and the new user is a non-authoritative user. (The system associates the IP address with the last authoritative user that logged in with the host.) For more information on authoritative vs. non-authoritative users, see [Users Database](#) on page 1258.

First Name

The user's first name, as obtained from the optional Defense Center-LDAP server connections. This field is blank if:

- you have not configured a Defense Center-LDAP server connection
- the Defense Center cannot correlate the user in the Defense Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login)
- there is no first name associated with the user on your LDAP servers

Last Name

The user's last name, as obtained from the optional Defense Center-LDAP server connections. This field is blank if:

- you have not configured a Defense Center-LDAP server connection
- the Defense Center cannot correlate the user in the Defense Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login)
- there is no last name associated with the user on your LDAP servers

E-Mail

The user's email address. This field is blank if:

- the user was added to the database via an AIM login
- the user was added to the database via an LDAP login and there is no email address associated with the user on your LDAP servers

Department

The user's department, as obtained from the optional Defense Center-LDAP server connections. If there is no department explicitly associated with the user on your LDAP servers, the department is listed as whatever default group the server assigns. For example, on Active Directory, this is `users (ad)`. This field is blank if:

- you have not configured a Defense Center-LDAP server connection
- the Defense Center cannot correlate the user in the Defense Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login)

Phone

The user's telephone number, as obtained from the optional Defense Center-LDAP server connections. This field is blank if:

- you have not configured a Defense Center-LDAP server connection
- the Defense Center cannot correlate the user in the Defense Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login)
- there is no telephone number associated with the user on your LDAP servers

User Type

The protocol used to detect the user. For example, for users added to the database when detects a POP3 login, the user type is `pop3`.

Count

The number of users that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Understanding User Details and Host History

LICENSE: FireSIGHT

From any event view that associates user identity data with other kinds of events, as well as from a table view of users, you can display the User Identity pop-up

window to learn more about a specific user. User information also appears in the terminating page for users workflows.

The screenshot shows a 'User Identity' section with the following fields: Username (leonard), Authentication Protocol (LDAP), First Name, Last Name, Email, Department, and Phone. Below this is a 'Host History' section with a table showing login events. The table has two columns: 'Hosts' and 'Time'. The first row shows '10.5.41.44' and '2011-11-06 14:42:22'. The second row shows '2011-11-07 14:42:22' with a bar graph below it representing the session duration.

Hosts	Time
10.5.41.44	2011-11-06 14:42:22
	2011-11-07 14:42:22

The user data you see is the same as you would see in the table view of users; for more information, see [Understanding the Users Table](#) on page 1447.

The host history provides a graphic representation of the last twenty-four hours of the user's activity. A list of IP addresses of the hosts that the user logged into and logged off of approximates login and logout times with bar graphs. A typical user might log on to and off of multiple hosts in the course of a day. For example, periodic automated logins to a mail server would display as multiple short sessions, while longer logins (such as during working hours) display longer sessions.

Note that when a non-authoritative user login to a host is detected, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user login is detected for that host, only another authoritative user login changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control. If you configure capture of failed logins in the network discovery policy, the host history also includes hosts where the user failed to log in.

The data used to generate the host history is stored in the user history database, which by default stores 10 million user login events. If you do not see any data in the host history for a particular user, either that user is inactive, or you may need to increase the database limit. For more information, see [Configuring Database Event Limits](#) on page 1978.

To view user details and host history:

ACCESS: Admin/Any Security Analyst

- ▶ You have two options:
 - In any event view that lists users, click the user icon () that appears next to a user identity.
 - In any users workflow, click the Users terminating page.
- User details appear.

Searching for Users

LICENSE: FireSIGHT

You can search for specific users. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify `n/a` or `blank` in the field to identify events where information is not available for that field; use `!n/a` or `!blank` to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

Specific User Search Criteria

For user type, valid search criteria are `ldap`, `pop3`, `imap`, and `aim`; because users are not added to the database based on SMTP logins, entering `smtp` will not return any results.

To search for users:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.

2. From the **Table** drop-down list, select **Users**.

The Users search page appears.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields. If you enter multiple criteria, the search returns only the records that match all the criteria.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default users workflow. To use a different workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with User Activity

LICENSE: FireSIGHT

The Sourcefire 3D System generates events that communicate the details of user activity on your network. Descriptions of the four types of user activity follow.

New User Identity

This event is generated when the system detects a user login for a user that is not in the database.

User Login

This event is generated when any of the following occur:

- an Active Directory Agent that you installed on an Active Directory server detects an LDAP login
- a managed device detects an LDAP, POP3, IMAP, SMTP, AIM, Oracle or SIP login

There are several points to keep in mind regarding user login events:

- SMTP logins are not recorded unless there is already a user with a matching email address in the database.
- Failed logins are only for LDAP, IMAP, and POP3, and only when detected in traffic. Users are not added to the detected users database as a result of a failed login, but the activity is optionally recorded in the user activity database, based on the user logging configuration in the network discovery policy.
- A user login is not recorded if you have specifically restricted its login type; see [Restricting User Logging](#) on page 1285.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Delete User Identity

This event is generated when you manually delete a user from the database.

User Identity Dropped: User Limit Reached

This event is generated when the system detects a user that is not in the database, but cannot add the user because you have reached the maximum number of users in the database as determined by your FireSIGHT license.

The total number of detected users the Defense Center can store depends on your FireSIGHT license. After you reach the licensed limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative users. If you have reached the limit and the system detects a login for a previously undetected authoritative user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new authoritative user.

When the system detects user activity, it is logged to the database. You can view, search, and delete user activity; you can also purge all user activity from the database.

Whenever possible the Sourcefire 3D System correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event. This can tell you who owns the host that was targeted by an attack, or who initiated an internal attack or portscan.

You can also use user activity in correlation rules. Based on the type of user activity as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and alert responses when network traffic meets your criteria. For more information on user activity, see [Understanding User Data Collection](#) on page 1252.

For more information, see the following sections:

- [Viewing User Activity Events](#) on page 1454
- [Understanding the User Activity Table](#) on page 1455
- [Searching for User Activity](#) on page 1456

Viewing User Activity Events

LICENSE: FireSIGHT

You can view a table of user activity, and then manipulate the event view depending on the information you are looking for.

The page you see when you access user activity differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of user activity and terminates in a user details page, which contains user details for every user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

For more information about the contents of the columns in the table, see [Understanding the User Activity Table](#) on page 1455. The following table, see describes some of the specific actions you can perform on an user activity workflow page. You can also perform the actions in the [Common Discovery Event](#)

[Actions table](#) on page 1387.

To view user activity:

ACCESS: Admin/Any Security Analyst

► Select **Analysis > Users > User Activity**.

The first page of the default user activity workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

TIP! If you are using a custom workflow that does not include the table view of user activity, click **(switch workflow)**, then select **User Activity**.

Understanding the User Activity Table

LICENSE: FireSIGHT

When the system detects user activity, it is logged to the database. Descriptions of the fields in the users table follow.

Time

The time that the system detected the user activity.

Event

The user activity type. For more information, see [Working with User Activity](#) on page 1453.

User

The user associated with the activity. At a minimum, this field contains a username and the protocol used to detect the user. If there is LDAP metadata on the user, this field may also contain the first name and last name of the user.

User Type

The protocol used to detect the user. For example, for users added to the database when the system detects a POP3 login, the user type is **pop3**.

IP Address

For User Login activity, the IP address involved in the login, which can be an IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for SMTP and Oracle logins), or the session originator (for SIP logins).

Note that an associated IP address does not mean the user is the current user for that IP address; when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user.

For other types of user activity, this field is blank.

Description

For Delete User Identity and User Identity Dropped activity, the username of the user who was deleted from the database or failed to be added to the database. For logins to network resources, **network login** is displayed. For other types of user activity, this field is blank.

Device

For user activity detected by a managed device, the name of the device. For other types of user activity, the managing Defense Center.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for User Activity

LICENSE: FireSIGHT

You can search for specific user activity. You may want to create searches customized for your network environment, then save them to reuse later.

General Search Syntax

The system displays examples of valid syntax next to each search field. When entering search criteria, keep the following points in mind:

- All fields accept negation (!).
- All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
- Many fields accept one or more asterisks (*) as wild cards.
- For some fields, you can specify **n/a** or **blank** in the field to identify events where information is not available for that field; use **!n/a** or **!blank** to identify the events where that field is populated.
- Most fields are case-insensitive.
- IP addresses may be specified using CIDR notation. For information on entering IPv4 and IPv6 addresses in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

- Click the add object icon () that appears next to a search field to use an object as a search criterion.

For detailed information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

To search for user activity:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Tables** drop-down menu, select **User Activity**.
The User Activity search page appears.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields. If you enter multiple criteria, the search returns only the records that match all the criteria. Click the add icon () that appears next to a search field to use an object as a search criterion.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges, you **must** save it as a private search.

6. You have the following options:

- Click **Search** to start the search.

Your search results appear in the default user activity workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

- Click **Save** if you are modifying an existing search and want to save your changes.
- Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

CHAPTER 36

CONFIGURING CORRELATION POLICIES AND RULES

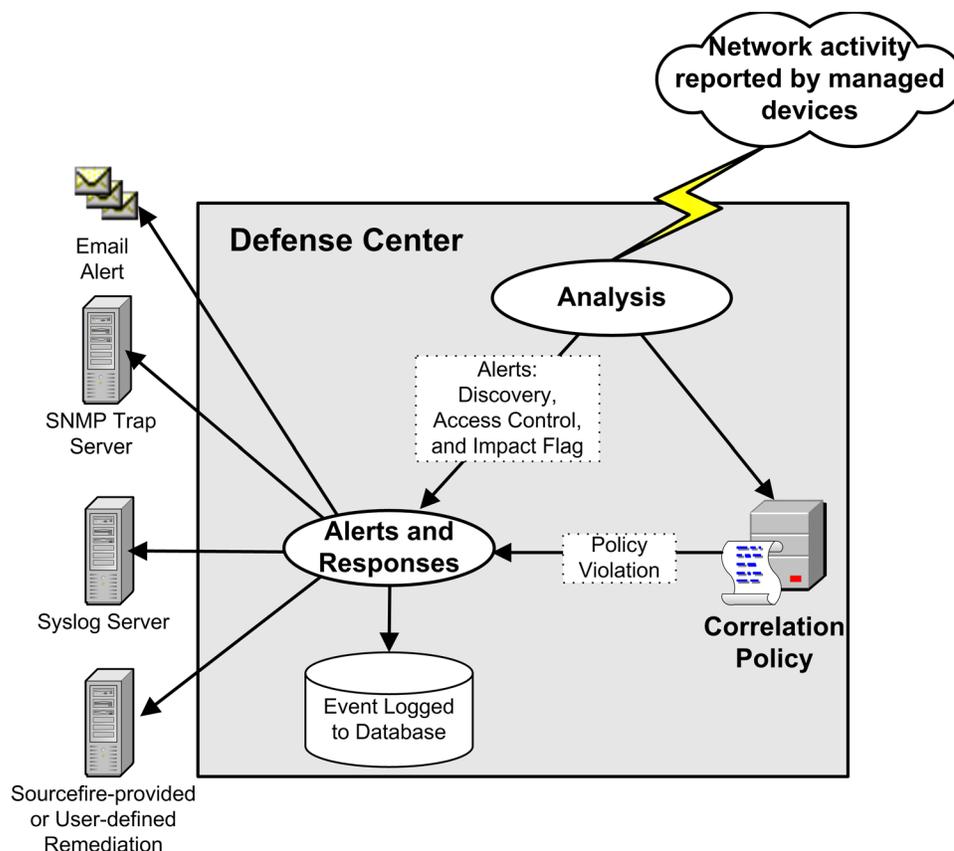
You can use the Sourcefire 3D System's *correlation* feature to build *correlation policies*, which are populated with *correlation rules* and *compliance white lists*, and that let you respond in real time to threats to your network. A *correlation policy violation* occurs when the activity on your network triggers either a correlation rule or white list.

A correlation rule triggers when a specific event generated by the Sourcefire 3D System either meets criteria that you specify, or when your network traffic deviates from your normal network traffic pattern as characterized in an existing traffic profile.

Compliance white lists, on the other hand, trigger when the system determines that a host on your network is running a prohibited operating system, client application (or client), application protocol, or protocol.

You can configure the Sourcefire 3D System to initiate responses to policy violations. Responses include simple alerts as well as various remediations (such as scanning a host). You can group responses so that the system launches multiple responses for each policy violation.

The following graphic illustrates the event notification and correlation process:



This chapter focuses on creating correlation rules, using those rules in policies, associating responses and response groups with those rules, and analyzing correlation events. For more information, see:

- [Creating Rules for Correlation Policies](#) on page 1461
- [Managing Rules for Correlation Policies](#) on page 1508
- [Grouping Correlation Responses](#) on page 1510
- [Creating Correlation Policies](#) on page 1513
- [Managing Correlation Policies](#) on page 1519
- [Working with Correlation Events](#) on page 1521

For information on creating compliance white lists and correlation responses (alerts and remediations), see:

- [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530
- [Working with Alert Responses](#) on page 526.
- [Configuring Remediations](#) on page 1606.

Creating Rules for Correlation Policies

LICENSE: FireSIGHT, Protection, URL Filtering, or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

Before you create a correlation policy, you should create correlation rules or compliance white lists (or both) to populate it.

IMPORTANT! This section describes how to create correlation rules. For information on creating compliance white lists, see [Creating Compliance White Lists](#) on page 1541.

A correlation rule triggers (and generates a correlation event) when your network traffic meets criteria that you specify. When you create correlation rules, you can use simple conditions or you can create more elaborate constructs by combining and nesting conditions and constraints.

You can further add to correlation rules in the following ways:

- Add a *host profile qualification* to constrain the rule using information from the host profile of a host involved in the triggering event.
- Add a *connection tracker* to a correlation rule so that after the rule's initial criteria are met, the system begins tracking certain connections. Then, a correlation event is generated only if the tracked connections meet additional criteria.
- Add a *user qualification* to a correlation rule to track certain users or groups of users. For example, you could constrain a correlation rule so that it triggers only when the identity of the source or destination user is a certain user or, for example, one from the marketing department.
- Add *snooze periods* and *inactive periods*. When a correlation rule triggers once, a snooze period causes that rule not to trigger again for a specified interval, even if the rule is violated again during the interval. After the snooze period has elapsed, the rule can trigger again (and start a new snooze period). During inactive periods, the correlation rule does not trigger.

WARNING! Evaluating complex correlation rules that trigger on frequently occurring events can degrade Defense Center performance. For example, a multi-condition rule that the Defense Center must evaluate against every connection logged by the system can cause resource overload.

The following table explains the licenses you must have to build effective correlation rules. If you do not have the appropriate licenses, correlation rules that use an unlicensed aspect of the Sourcefire 3D System do not trigger.

License Requirements for Building Correlation Rules

To...	YOU NEED THIS LICENSE...
trigger a correlation rule on an intrusion event	Protection
trigger a correlation rule on a discovery event, host input event, or user activity, or to add a host profile or user qualification to a correlation rule	FireSIGHT
trigger a correlation rule on a connection event or endpoint-based malware event, or to add a connection tracker to a rule	Any
trigger a correlation rule on a connection event with URL data, or build a connection tracker using URL data Note that neither Series 2 devices nor the DC500 Defense Center support URL filtering by category or reputation, and Series 2 devices do not support URL filtering by literal URL or URL group.	URL Filtering
trigger a correlation rule on a malware event based on network-based malware data or retrospective network-based malware data Note that neither Series 2 devices nor the DC500 Defense Center support network-based malware protection.	Malware

When you create either correlation rule trigger criteria, host profile qualifications, user qualifications, or connection trackers, the syntax varies but the mechanics remain consistent. For more information, See [Understanding Rule Building Mechanics](#) on page 1499.

To create a correlation rule:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then select the **Rule Management** tab.
The Rule Management page appears.

2. Click **Create Rule**.

The Create Rule page appears.

The screenshot shows the 'Create Rule' configuration page. It has a light gray background and is organized into several sections. At the top is the 'Rule Information' section, which includes three input fields: 'Rule Name', 'Rule Description', and 'Rule Group' (which is currently set to 'Ungrouped'). Below this is a section titled 'Select the type of event for this rule' with a dropdown menu. Underneath that is a blue bar with the text 'If [dropdown] and it meets the following conditions:'. The 'Rule Options' section follows, containing a 'Snooze' field with a value of '0' and a unit of 'hours', and an 'Add Inactive Period' button. At the bottom of the form are 'Save' and 'Cancel' buttons.

3. Provide basic rule information, such as the rule name, description, and group. See [Providing Basic Rule Information](#) on page 1464.

4. Specify the basic criteria on which you want the rule to trigger. See [Specifying Correlation Rule Trigger Criteria](#) on page 1464.

5. Optionally, add a host profile qualification to the rule. See [Adding a Host Profile Qualification](#) on page 1481.

6. Optionally, add a connection tracker to the rule. See [Adding a Connection Tracker](#) on page 1486.

7. Optionally, add a user qualification to the rule. See [Adding a User Qualification](#) on page 1496.

8. Optionally, add an inactive period or snooze period (or both) to the rule. See [Adding Snooze and Inactive Periods](#) on page 1498.

9. Click **Save Rule**.

The rule is saved. You can now use the rule within correlation policies or within other correlation rules that trigger on the same event type.

Providing Basic Rule Information

LICENSE: Any

You must give each correlation rule a name and, optionally, a short description. You can also place the rule in a rule group.



Rule Information

Rule Name: Bagle Worm

Rule Description: Detects Bagle worm activity

Rule Group: Worms

To provide basic rule information:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then select the **Rule Management** tab.
The Rule Management page appears.
2. Click **Create Rule**.
The Create Rule page appears.
3. On the Create Rule page, in the **Rule Name** field, type a name for the rule.
4. In the **Rule Description** field, type a description for the rule.
5. Optionally, select a group for the rule from the **Rule Group** drop-down list.
For more information on rule groups, see [Managing Rules for Correlation Policies](#) on page 1508.
6. Continue with the procedure in the next section, [Specifying Correlation Rule Trigger Criteria](#).

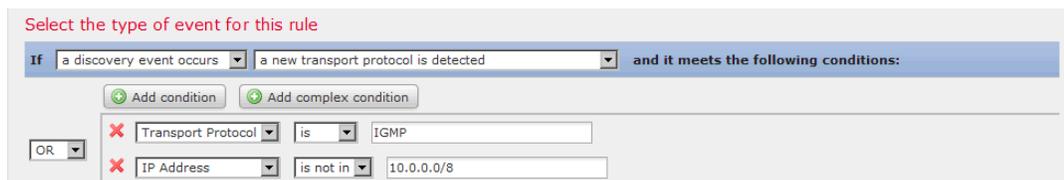
Specifying Correlation Rule Trigger Criteria

LICENSE: FireSIGHT, Protection, URL Filtering, or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

A simple correlation rule requires only that an event of a certain type occurs; you do not need to provide more specific conditions. For example, correlation rules based on traffic profile changes do not require any conditions at all. In contrast, correlation rules may be complex, with multiple nested conditions. For example, the rule shown in the following graphic comprises criteria that direct the rule to trigger if an IP address that is not in the 10.x.x.x subnet transmits an IGMP message.



Select the type of event for this rule

If a discovery event occurs a new transport protocol is detected and it meets the following conditions:

OR

- ✗ Transport Protocol is IGMP
- ✗ IP Address is not in 10.0.0.0/8

To specify correlation rule trigger criteria:

ACCESS: Admin/Discovery Admin

1. Select the type of event on which you want to base your rule.

When you build a correlation rule, you must first select the type of event on which you want to base your rule. You have a few options under **Select the type of event for this rule:**

- Select **an intrusion event occurs** to trigger your rule when a specific intrusion event occurs.
- Select **a Malware event occurs** to trigger the rule when a specific malware event occurs.

Note that because neither Series 2 devices nor the DC500 Defense Center support network-based malware protection, these appliances do not support triggering a correlation rule on a malware event based on network-based malware data or retrospective network-based malware data.

- Select **a discovery event occurs** to trigger the rule when a specific discovery event occurs. When triggering a correlation rule on a discovery event, you must also choose the type of event you want to use. You can choose from a subset of the discovery events described in [Understanding Discovery Event Types](#) on page 1389; you cannot, for example, trigger a correlation rule on a hops change. You can, however, choose **there is any type of event** to trigger the rule when any kind of discovery event occurs.
- Select **user activity is detected** to trigger the rule when a new user is detected or a user logs in to a host.
- Select **a host input event occurs** to trigger the rule when a specific host input event occurs. When triggering a correlation rule on a host input event, you must also choose the type of event you want to use. You can choose from a subset of the events described in [Understanding Host Input Event Types](#) on page 1394.
- Select **a connection event occurs** to trigger the rule when connection data meets specific criteria. When triggering a correlation rule on a connection event, you must also choose whether you want to use connection events that represent the beginning or the ending of the connection, or either.

Note that because neither Series 2 devices nor the DC500 Defense Center support URL filtering by category or reputation, these appliances do not support triggering a correlation rule on a connection event with URL data, or building a connection tracker using URL data.
- Select **a traffic profile changes** to trigger the correlation rule when network traffic deviates from your normal network traffic pattern as characterized in an existing traffic profile.

2. Specify the rule's conditions.

The syntax you can use within correlation rule trigger criteria conditions varies depending on the base event you chose in step 1, but the mechanics are the same. For more information, see [Understanding Rule Building Mechanics](#) on page 1499.

The syntax you can use to build conditions is described in the following sections:

- [Syntax for Intrusion Events](#) on page 1467
- [Syntax for Malware Events](#) on page 1469
- [Syntax for Discovery Events](#) on page 1470
- [Syntax for User Activity Events](#) on page 1474
- [Syntax for Host Input Events](#) on page 1475
- [Syntax for Connection Events](#) on page 1476
- [Syntax for Traffic Profile Changes](#) on page 1479

TIP! You can nest rules that share the base event type you specified in step 1. For example, if you create a new rule based on the detection of an open TCP port, the trigger criteria for the new rule could include **rule "MyDoom Worm" is true** and **rule "Kazaa (TCP) P2P" is true**.

3. Optionally, continue with the procedures in the following sections:

- [Adding a Host Profile Qualification](#) on page 1481
- [Adding a Connection Tracker](#) on page 1486
- [Adding a User Qualification](#) on page 1496
- [Adding Snooze and Inactive Periods](#) on page 1498

If you are finished building the correlation rule, continue with step 9 of the procedure in [Creating Rules for Correlation Policies](#) on page 1461 to save the rule.

Syntax for Intrusion Events

LICENSE: Protection

The [Syntax for Intrusion Events](#) table describes how to build a correlation rule condition when you choose an intrusion event as the base event.

Syntax for Intrusion Events

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Access Control Policy	Select one or more access control policies that use the intrusion policy that generated the intrusion event.
Access Control Rule Name	Type all or part of the name of the access control rule that uses the intrusion policy that generated the intrusion event.
Application Protocol	Select one or more application protocols associated with the intrusion event.
Classification	Select one or more classifications.
Client	Select one or more clients associated with the intrusion event.
Destination IP, Source IP, or Source/Destination IP	Specify a single IP address, an address block, or a comma-separated list comprised of any of these. For information on using IP address notation and prefix lengths in the Sourcefire 3D System, see IP Address Conventions on page 57.
Destination Port/ICMP Code or Source Port/ICMP Type	Type the port number or ICMP type for source traffic or the port number or ICMP type for destination traffic.
Device	Select one or more devices that may have generated the event.
Egress Interface or Ingress Interface	Select one or more interfaces.
Egress Security Zone or Ingress Security Zone	Select one or more security zones.
Generator ID	Select one or more preprocessors. See Using Advanced Settings in an Intrusion Policy on page 773 for more information about available preprocessors.

Syntax for Intrusion Events (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Impact Flag	<p>Select the impact flag assigned to the intrusion event. You select any of the following along with operators that specify is, is not, is greater than, and so on:</p> <ul style="list-style-type: none"> • 0 — gray (Unknown) • 1 — red (Vulnerable) • 2 — orange (Potentially Vulnerable) • 3 — yellow (Currently Not Vulnerable) • 4 — blue (Unknown Target) <p>IMPORTANT! Because there is no operating system information available for hosts added to the network map based on NetFlow data, the Defense Center cannot assign Vulnerable (level 1: red) impact flags for intrusion events involving those hosts, unless you use the host input feature to manually set the host operating system identity.</p> <p>For more information, see Using Impact Flags to Evaluate Events on page 641.</p>
Inline Result	<p>Select either:</p> <ul style="list-style-type: none"> • dropped, to specify whether the packet was dropped in an inline, switched, or routed deployment • would have dropped, to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline, switched, or routed deployment <p>Note that the system does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of the rule state or the drop behavior of the intrusion policy. For more information, see Setting Rule States on page 743, Setting Drop Behavior in an Inline Deployment on page 689, Configuring Passive Interfaces on page 267, and Tap Mode on page 275.</p>
Intrusion Policy	<p>Select one or more intrusion policies that generated the intrusion event.</p>
Priority	<p>Select the rule priority: low, medium, or high.</p> <p>For rule-based intrusion events, the priority corresponds to either the value of the priority keyword or the value for the classtype keyword. For other intrusion events, the priority is determined by the decoder or preprocessor.</p>
Protocol	<p>Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.</p>
Rule Message	<p>Type all or part of the rule message.</p>

Syntax for Intrusion Events (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Rule SID	Type a single Short ID number (SID) or multiple SIDs separated by commas. IMPORTANT! If you choose is in or is not in as the operator, you cannot use the multi-selection pop-up window. You must type a comma-separated list of SIDs.
Rule Type	Specify whether the rule is or is not local. Local rules include custom standard text intrusion rules, standard text rules that you modified, and any new instances of shared object rules created when you saved the rule with modified header information. For more information, see Modifying Existing Rules on page 1184.
Username	Type the username of the user logged into the source host in the intrusion event.
VLAN ID	Type the innermost VLAN ID associated with the packet that triggered the intrusion event
Web Application	Select one or more web applications associated with the intrusion event.

Syntax for Malware Events

LICENSE: Any or Malware

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

The syntax for correlation rule conditions based on malware events depends on whether the event is reported by an endpoint-based malware agent, detected by a managed device, or detected by a managed device and retrospectively identified as malware.

Note that because neither Series 2 devices nor the DC500 Defense Center support network-based malware protection, these appliances do not support triggering a correlation rule on a malware event based on network-based malware data or retrospective network-based malware data.

The [Syntax for Malware Events](#) table describes how to build a correlation rule condition when you choose a malware event as the base event.

Syntax for Malware Events

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Application Protocol	Select one or more application protocols associated with the malware event.
Client	Select one or more clients associated with the malware event.
Destination IP, Host IP, or Source IP	Specify a single IP address or address block. For information on using IP address notation in the Sourcefire 3D System, see IP Address Conventions on page 57.
Destination Port/ICMP Code	Type the port number or ICMP code for destination traffic.
Event Type	Select one or more endpoint-based event types associated with the malware event. For more information, see Malware Event Types on page 1238.
File Name	Type the name of the file.
File Type	Select the type of file, for example, PDF or MSEXEX.
File Type Category	Select the general category of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.
SHA-256	Type or paste the SHA-256 hash value of the file.
Source Port/ICMP Type	Type the port number or ICMP type for source traffic.
Web Application	Select one or more web applications associated with the malware event.

Syntax for Discovery Events

LICENSE: FireSIGHT

If you base your correlation rule on a discovery event, you must first choose the type of event you want to use from a drop-down list. The following table lists the events you can choose as trigger criteria from the drop-down list, cross-referenced with their corresponding event types. For detailed descriptions

of discovery event types, see [Understanding Discovery Event Types](#) on page 1389.

Correlation Rule Trigger Criteria vs. Discovery Event Types

SELECT THIS OPTION...	TO TRIGGER THE RULE ON THIS EVENT TYPE...
a client has changed	Client Update
a client timed out	Client Timeout
a host ip address is reused	DHCP: IP Address Reassigned
a host is deleted because the host limit was reached	Host Deleted: Host Limit Reached
a host is identified as a network device	Host Type Changed to Network Device
a host timed out	Host Timeout
a host's IP address has changed	DHCP: IP Address Changed
a NETBIOS name change is detected	NETBIOS Name Change
a new client is detected	New Client
a new IP host is detected	New Host
a new MAC address is detected	Additional MAC Detected for Host
a new MAC host is detected	New Host
a new network protocol is detected	New Network Protocol
a new transport protocol is detected	New Transport Protocol
a TCP port closed	TCP Port Closed
a TCP port timed out	TCP Port Timeout
a UDP port closed	UDP Port Closed
a UDP port timed out	UDP Port Timeout
a VLAN tag was updated	VLAN Tag Information Update

Correlation Rule Trigger Criteria vs. Discovery Event Types (Continued)

SELECT THIS OPTION...	TO TRIGGER THE RULE ON THIS EVENT TYPE...
an open TCP port is detected	New TCP Port
an open UDP port is detected	New UDP Port
the OS information for a host has changed	New OS
the OS or server identity for a host has a conflict	Identity Conflict
the OS or server identity for a host has timed out	Identity Timeout
there is any kind of event	any event type
there is new information about a MAC address	MAC Information Change
there is new information about a TCP server	TCP Server Information Update
there is new information about a UDP server	UDP Server Information Update

Note that you cannot trigger a correlation rule on hops changes, or when the system drops a new host due to reaching the licensed host limit. You can, however, choose **there is any type of event** to trigger the rule when any type of discovery event occurs.

After you choose the discovery event type, you can build correlation rule conditions as described in the [Syntax for Discovery Events](#) table below. Depending on the type of event you choose, you can build conditions using subsets of the criteria in the following table. For example, if you trigger your correlation rule when a new client is detected, you can build conditions based on

the IP or MAC address of the host, the client name, type, or version, and the device that detected the event.

Syntax for Discovery Events

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Application Protocol	Select one or more application protocols.
Application Port	Type the application protocol port number.
Client	Select one or more clients.
Client Version	Type the version number of the client.
Device	Select one or more devices that may have generated the discovery event.
Hardware	Type the hardware model for the mobile device. For example, to match all Apple iPhones, type iPhone .
Host Type	Select one or more host types from the drop-down list. You can choose between a host or one of several types of network device.
IP Address or New IP Address	Type a single IP address or address block. For information on using IP address notation in the Sourcefire 3D System, see IP Address Conventions on page 57.
Jailbroken	Select Yes to indicate that the host in the event is a jailbroken mobile device or No to indicate that it is not.
MAC Address	Type all or part of the MAC address of the host. For example, if you know that devices from a certain hardware manufacturer have MAC addresses that begin with 0A:12:34, you could choose begins with as the operator, then type 0A:12:34 as the value.
MAC Type	Select whether the MAC address was ARP/DHCP Detected . That is, select whether the system positively identified the MAC address as belonging to the host (is ARP/DHCP Detected), or whether the system is seeing many hosts with that MAC address because, for example, there is a router between the managed device and the host (is not ARP/DHCP Detected).
MAC Vendor	Type all or part of the name of the MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.
Mobile	Select Yes to indicate that the host in the event is a mobile device or No to indicate that it is not.
NETBIOS Name	Type the NetBIOS name of the host.

Syntax for Discovery Events (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Network Protocol	Type the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers .
OS Name	Select one or more operating system names.
OS Vendor	Select one or more operating system vendors.
OS Version	Select one or more operating system versions.
Protocol or Transport Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers .
Source	Select the source of the host input data (for operating system and server identity changes and timeouts).
Source Type	Select the type of the source for the host input data (for operating system and server identity changes and timeouts).
VLAN ID	Type the VLAN ID of the host involved in the event.
Web Application	Select a web application.

Syntax for User Activity Events

LICENSE: FireSIGHT

If you base your correlation rule on user activity, you must first choose the type of user activity you want to use from a drop-down list, either:

- a user logged into a host, or
- a new user identity was detected

After you choose the user activity type, you can build correlation rule conditions as described in the [Syntax for User Activity](#) table. Depending on the type of user activity you choose, you can build conditions using subsets of the criteria in the

following table; for correlation rules triggered on new user identity, you cannot specify an IP address.

Syntax for User Activity

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Device	Select one or more devices that may have detected the user activity.
IP Address	Type a single IP address or address block. For information on using IP address notation in the Sourcefire 3D System, see IP Address Conventions on page 57.
Username	Type a username.

Syntax for Host Input Events

LICENSE: FireSIGHT

If you base your correlation rule on a host input event, you must first choose the type of host input event you want to use from a drop-down list. The following table lists the events you can choose as trigger criteria from the drop-down list, cross-referenced with their corresponding host input event types. For detailed descriptions of host input event types, see [Understanding Host Input Event Types](#) on page 1394.

Correlation Rule Trigger Criteria vs. Host Input Event Types

SELECT THIS OPTION...	TO TRIGGER THE RULE ON THIS EVENT TYPE...
a client is added	Add Client
a client is deleted	Delete Client
a host is added	Add Host
a protocol is added	Add Protocol
a protocol is deleted	Delete Protocol
a scan result is added	Add Scan Result
a server definition is set	Set Server Definition
a server is added	Add Port
a server is deleted	Delete Port

Correlation Rule Trigger Criteria vs. Host Input Event Types (Continued)

SELECT THIS OPTION...	TO TRIGGER THE RULE ON THIS EVENT TYPE...
a vulnerability is marked invalid	Vulnerability Set Invalid
a vulnerability is marked valid	Vulnerability Set Valid
an address is deleted	Delete Host/Network
an attribute value is deleted	Host Attribute Delete Value
an attribute value is set	Host Attribute Set Value
an OS definition is set	Set Operating System Definition
host criticality is set	Set Host Criticality

You cannot trigger a correlation rule when you add, delete, or change the definition of a user-defined host attribute, or set a vulnerability impact qualification.

After you choose the host input event type, you can build correlation rule conditions as described in the [Syntax for Host Input Events](#) table below. Depending on the type of host input event you choose, you can build conditions using subsets of the criteria in the following table. For example, if you trigger your correlation rule when a client is deleted, you can build conditions based on the IP address of the host involved in the event, the source type of the deletion (manual, third-party application, or scanner), and the source itself (a specific scanner type or user).

Syntax for Host Input Events

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
IP Address	Type a single IP address or address block. For information on using IP address notation in the Sourcefire 3D System, see IP Address Conventions on page 57.
Source	Select the source for the host input data.
Source Type	Select the type of the source for the host input data.

Syntax for Connection Events

LICENSE: Any

If you base your correlation rule on a connection event, you must first choose whether you want to evaluate events that represent the beginning or ending of

the connection, or either the beginning or the end. After you choose the connection event type, you can build correlation rule conditions as described in [Syntax for Connection Events](#) on page 1477.

When you build rule conditions, you should make sure that your network traffic can trigger the rules. The information available for any individual connection or connection summary event depends on several factors, including the detection method, the logging method, and event type. For more information, see [Information Available in Connection Events](#) on page 550.

Syntax for Connection Events

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Access Control Policy	Select one or more access control policies that logged the connection.
Access Control Rule Action	Select one or more actions associated with the access control rule that logged the connection. IMPORTANT! Select Monitor to trigger correlation events when network traffic matches the conditions of any Monitor rule, regardless of the rule or default action that later handles the connection
Access Control Rule Name	Type all or part of the name of the access control rule that logged the connection. IMPORTANT! You can type the name of any Monitor rule whose conditions were matched by a connection, regardless of the rule or default action that later handled the connection.
Application Protocol	Select one or more application protocols associated with the connection.
Client	Select one or more clients.
Client Version	Type the version number of the client.
Connection Duration	Type the duration of the connection event, in seconds.
Connection Type	Select whether you want to trigger the correlation rule based on whether the connection was detected by a Sourcefire managed device (FireSIGHT) or was exported by a NetFlow-enabled device (NetFlow).
Device	Select one or more devices that either detected the connection, or that processed the connection (for connection data exported by a NetFlow-enabled device).
Egress Interface or Ingress Interface	Select one or more interfaces.
Egress Security Zone or Ingress Security Zone	Select one or more security zones.

Syntax for Connection Events (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Initiator Bytes, Responder Bytes, or Total Bytes	Type one of: <ul style="list-style-type: none"> • the number of bytes transmitted (Initiator Bytes). • the number of bytes received (Responder Bytes). • the number of bytes both transmitted and received (Total Bytes).
Initiator IP, Responder IP, or Initiator/Responder IP	Specify a single IP address, an address block, or a comma-separated list comprised of any of these. For information on using IP address notation and prefix lengths in the Sourcefire 3D System, see IP Address Conventions on page 57.
Initiator Packets, Responder Packets, or Total Packets	Type one of: <ul style="list-style-type: none"> • the number of packets transmitted (Initiator Packets). • the number of packets received (Responder Packets). • the number of packets both transmitted and received (Total Packets).
Initiator Port/ICMP Type or Responder Port/ICMP Code	Type the port number or ICMP type for initiator traffic or the port number or ICMP code for responder traffic.
NETBIOS Name	Type the NetBIOS name of the monitored host in the connection.
NetFlow Device	Select the IP address of the NetFlow-enabled device that exported the connection data you want to use to trigger the correlation rule. If you did not add any NetFlow-enabled devices to your deployment, the NetFlow Device drop-down list is blank.
Reason	Select one or more reasons associated with the connection event.
TCP Flags	Select a TCP flag that a connection event must contain in order to trigger the correlation rule. IMPORTANT! Only connection data exported by NetFlow-enabled devices contain TCP flags.
Transport Protocol	Type the transport protocol used by the connection: TCP or UDP .
URL	Type all or part of the URL visited in the connection.
URL Category	Select one or more URL categories for the URL visited in the connection.
URL Reputation	Select one or more URL reputation values for the URL visited in the connection.
Username	Type the username of the user logged into either host in the connection.
Web Application	Select one or more web applications associated with the connection.

Syntax for Traffic Profile Changes

LICENSE: Any

If you base your correlation rule on a traffic profile change, the rule triggers when network traffic deviates from your normal network traffic pattern as characterized in an existing traffic profile. For information on how to build a traffic profile, see [Creating Traffic Profiles](#) on page 1585.

You can trigger the rule based on either raw data or on the statistics calculated from the data. For example, you could write a rule that triggers if the amount of data traversing your network (measured in bytes) suddenly spikes, which could indicate an attack or other security policy violation. You could specify that the rule trigger if either:

- the number of bytes traversing your network spikes above a certain number of standard deviations above or below the mean amount of traffic

Note that to create a rule that triggers when the number of bytes traversing your network falls outside a certain number of standard deviations (whether above or below), you must specify upper and lower bounds, as shown in the following graphic.

Select the type of event for this rule

If a traffic profile changes and the profile is Sample Traffic Profile and it meets the following conditions:

OR

- Responder Bytes data are greater than 3 standard deviation(s) use velocity data
- Responder Bytes data are less than 3 standard deviation(s) use velocity data

To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations **above** the mean, use only the first condition shown in the graphic.

To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations **below** the mean, use only the second condition.

- the number of bytes traversing your network spikes above a certain number of bytes

You can select the **use velocity data** check box (see [Changing the Graph Type](#) on page 557), to trigger the correlation rule based on rates of change between data points. If you wanted to use velocity data in the above example, you could specify that the rule triggers if either:

- the change in the number of bytes traversing your network spikes above or below a certain number of standard deviations above the mean rate of change
- the change in the number of bytes traversing your network spikes above a certain number of bytes

The following table describes how to build a condition in a correlation rule when you choose a traffic profile change as the base event. If your traffic profile uses

connection data exported by NetFlow-enabled devices, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272 to learn about how the detection method can affect the data used to create your traffic profile.

Syntax for Traffic Profile Changes

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN TYPE...	AND THEN CHOOSE ONE OF THE FOLLOWING...
Number of Connections	the total number of connections detected or the number of standard deviations either above or below the mean that the number of connections detected must be in to trigger the rule	connections standard deviation(s)
Total Bytes, Initiator Bytes, or Responder Bytes	one of: <ul style="list-style-type: none"> • the total bytes transmitted (Total Bytes) • the number of bytes transmitted (Initiator Bytes) • the number of bytes received (Responder Bytes) or the number of standard deviations either above or below the mean that one of the above criteria must be in to trigger the rule	bytes standard deviation(s)

Syntax for Traffic Profile Changes (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN TYPE...	AND THEN CHOOSE ONE OF THE FOLLOWING...
Total Packets, Initiator Packets, or Responder Packets	one of: <ul style="list-style-type: none"> the total packets transmitted (Total Packets) the number of packets transmitted (Initiator Packets) the number of packets received (Responder Packets) or the number of standard deviations either above or below the mean that one of the above criteria must be in trigger the rule	packets standard deviation(s)
Unique Initiators	the number of unique hosts that initiated sessions or the number of standard deviations either above or below the mean that the number of unique initiators detected must be to trigger the rule	initiators standard deviation(s)
Unique Responders	the number of unique hosts that responded to sessions or the number of standard deviations either above or below the mean that the number of unique responders detected must be to trigger the rule	responders standard deviation(s)

Adding a Host Profile Qualification

LICENSE: FireSIGHT

If you are using a connection, intrusion, discovery, user activity, or host input event to trigger your correlation rule, you can constrain the rule based on the host profile of a host involved in the event. This constraint is called a *host profile qualification*.

IMPORTANT! You **cannot** add a host profile qualification to a correlation rule that triggers on a malware event, traffic profile change, or on the detection of a new IP host.

For example, you could constrain a correlation rule so that it triggers only when a Microsoft Windows host is the target of the offending traffic, because only Microsoft Windows computers are vulnerable to the vulnerability the rule is written for. As another example, you could constrain a correlation rule so that it triggers only when the host is out of compliance with a white list.

To match against *implied* or generic clients, create a host profile qualification based on the application protocol used by the server responding to the client. When the client list on a host that acts as the initiator or source of a connection includes an application protocol name followed by **client**, that client may actually be an implied client. In other words, the system reports that client based on server response traffic that uses the application protocol for that client, not on detected client traffic.

For example, if the system reports **HTTPS client** as a client on a host, create a host profile qualification for **Responder Host** or **Destination Host** where **Application Protocol** is set to **HTTPS**, because **HTTPS client** is reported as a generic client based on the HTTPS server response traffic sent by the responder or destination host.

Note that to use a host profile qualification, the host must exist in the network map and the host profile property you want to use as a qualification must already be included in the host profile. For example, if you configure a correlation rule to trigger when an intrusion event is generated for a host running Windows, the rule only triggers if the host is already identified as Windows when the intrusion event is generated.

To add a host profile qualification:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then select the **Rule Management** tab.
The Rule Management page appears.
2. Click **Create Rule**.
The Create Rule page appears.
3. On the Create Rule page, click **Add Host Profile Qualification**.
The Host Profile Qualification section appears.



TIP! To remove a host profile qualification, click **Remove Host Profile Qualification**.

4. Build the host profile qualification's conditions.
You can create a single, simple condition, or you can create more elaborate constructs by combining and nesting conditions. See [Understanding Rule Building Mechanics](#) on page 1499 for information on how to use the web interface to build conditions.
The syntax you can use to build conditions is described in [Syntax for Host Profile Qualifications](#) on page 1483.

5. Optionally, continue with the procedures in the following sections:

- [Adding a Connection Tracker](#) on page 1486
- [Adding a User Qualification](#) on page 1496
- [Adding Snooze and Inactive Periods](#) on page 1498

If you are finished building the correlation rule, continue with step 9 of the procedure in [Creating Rules for Correlation Policies](#) on page 1461 to save the rule.

Syntax for Host Profile Qualifications

LICENSE: FireSIGHT

When you build a host profile qualification condition, you must first select the host you want to use to constrain your correlation rule. The host you can choose depends on the type of event you are using to trigger the rule, as follows:

- If you are using a connection event, select **Responder Host** or **Initiator Host**.
- If you are using an intrusion event, select **Destination Host** or **Source Host**.
- If you are using a discovery event, host input event, or user activity, select **Host**.

After you select the host type, you continue building your host profile qualification condition, as described in the [Syntax for Host Profile Qualifications](#) table.

Note that although you can configure the network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. In addition, if you use connection data exported by NetFlow-enabled devices, keep in mind that NetFlow records do not contain information about which host is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Syntax for Host Profile Qualifications

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Host Type	Select one or more host types. You can choose between a host or one of several types of network device.
NETBIOS Name	Type the NetBIOS name of the host.
Operating System > OS Name	Select one or more operating system names.

Syntax for Host Profile Qualifications (Continued)

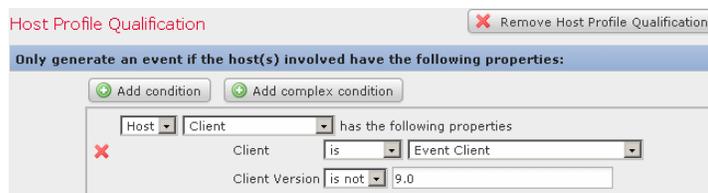
IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Operating System > OS Vendor	Select one or more operating system vendor names.
Operating System > OS Version	Select one or more operating system versions.
Hardware	Type the hardware model for the mobile device. For example, to match all Apple iPhones, type iPhone .
Jailbroken	Select Yes to indicate that the host in the event is a jailbroken mobile device or No to indicate that it is not.
Mobile	Select Yes to indicate that the host in the event is a mobile device or No to indicate that it is not.
Network Protocol	Type the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers .
Transport Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers .
Host Criticality	Select the host criticality: None , Low , Medium , or High . For more information on host criticality, see Working with the Predefined Host Attributes on page 1369.
VLAN ID	Type the VLAN ID associated with the host.
Application Protocol > Application Protocol	Select one or more application protocols.
Application Protocol > Application Port	Type the application protocol port number. If you are using an intrusion event to trigger the correlation rule, depending on the host you chose for the host profile qualification, this field is prepopulated with a port in the event: dst_port (for Destination Host) or src_port (for Source Host).
Application Protocol > Protocol	Select one or more protocols.
Client > Client	Select one or more clients. Select Event Client to choose the client from the event.
Client > Client Version	Type the client version.
Web Application	Select a web application.

Syntax for Host Profile Qualifications (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
MAC Address > MAC Address	Type all or part of the MAC address of the host. For example, if you know that devices from a certain hardware have MAC addresses that begin with 0A:12:34, you could choose begins with as the operator, then type 0A:12:34 as the value.
MAC Address > MAC Type	Select whether the MAC type is ARP/DHCP Detected . That is, select whether the system positively identified the MAC address as belonging to the host (is ARP/DHCP Detected), whether the system is seeing many hosts with that MAC address because, for example, there is a router between the managed device and the host (is not ARP/DHCP Detected), or whether the MAC type is irrelevant (is any).
MAC Vendor > MAC Vendor	Type all or part of the name of the MAC hardware vendor of the host.
any available host attribute, including the default compliance white list host attribute	Specify the appropriate value, which depends on the type of host attribute you select: <ul style="list-style-type: none"> • If the host attribute type is Integer, enter an integer value in the range defined for the attribute. • If the host attribute type is Text, enter a text value. • If the host attribute type is List, select a valid list string. • If the host attribute type is URL, enter a URL value. For more information on host attributes, see Working with User-Defined Host Attributes on page 1370.

Note that you can often use event data when constructing a host profile qualification. For example, assume your correlation rule triggers when the system detects the use of Internet Explorer on one of your monitored hosts. Further assume that when you detect this use, you want to generate an event if the version of the browser is not the latest (for this example, assume the latest version is 9.0).

You could add a host profile qualification to this correlation rule so that the rule triggers only if the **Client** is the **Event Client** (that is, Internet Explorer), but the **Client Version** is not 9.0.



Adding a Connection Tracker

LICENSE: Any

If you are using a connection, intrusion, discovery, user activity, or host input event to trigger your correlation rule, you can add a *connection tracker* to the rule. A connection tracker constrains a correlation rule so that after its initial criteria are met, the system begins tracking certain connections. The Defense Center generates a correlation event for the rule only if the tracked connections meet additional criteria. You **cannot** add a connection tracker to a rule that triggers on a malware event or traffic profile change.

Unlike traffic profiles, which typically monitor a broad range of network traffic and run persistently, connection trackers typically monitor very specific traffic and, when triggered, run only for a finite, specified time. For more information on traffic profiles, see [Creating Traffic Profiles](#) on page 1585.

When you configure a connection tracker, you must specify:

- which connections you want to track
- the conditions that the connections you are tracking must meet for the Defense Center to generate a correlation event
- the amount of time that can elapse before the connection tracker expires, that is, the amount of time after which the system stops tracking connections if the connection tracker conditions have not yet been met

When you add a connection tracker to a correlation rule, after your network traffic satisfies the basic conditions of the rule as well as of any host profile and user qualifications you added, the system begins tracking connections according to the criteria you specified.

If your network traffic meets the connection tracker's conditions before time expires, the Defense Center generates a correlation event and immediately stops tracking connections for that rule instance. On the other hand, if time expires before your network traffic meets the conditions in the connection tracker, the Defense Center does not generate a correlation event. After time expires, the system stops tracking connections for that rule instance.

TIP! You can add a connection tracker to a simple correlation rule that requires only that any connection, intrusion, discovery, user identity, or host input event occurs.

For example, consider a scenario where you archive sensitive files on network 10.1.0.0/16, and where hosts outside this network typically do not initiate connections to hosts inside the network. An occasional connection initiated from outside the network might occur, but you have determined that when four or more connections are initiated within two minutes, there is cause for concern.

The rule shown in the following graphic specifies that when a connection occurs from outside the 10.1.0.0/16 network to inside the network, the system begins tracking connections that meet that criterion. The Defense Center then generates

a correlation event if the system detects four connections (including the original connection) within two minutes that match that signature.

Rule Information Add User Qualification Add Host Profile Qualification

Rule Name:

Rule Description:

Rule Group:

Select the type of event for this rule

If at either the beginning or the end of the connection and it meets the following conditions:

AND is not in

is in

Connection Tracker Remove Connection Tracker

... start tracking connections that meet the following conditions:

AND is not in (switch to event fields)

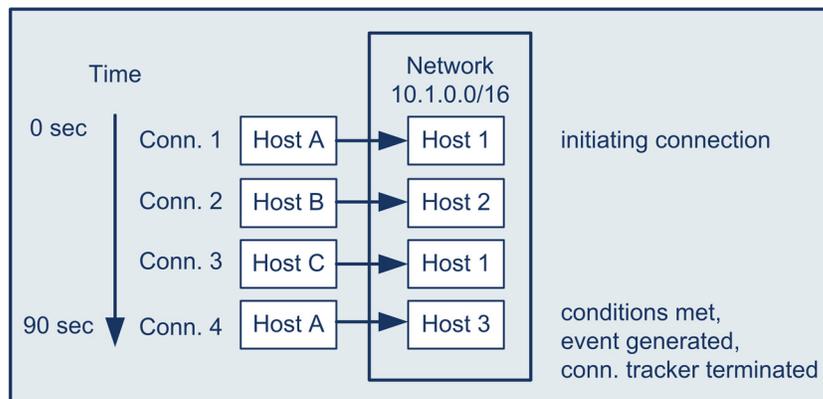
is in (switch to event fields)

... and generate an event if:

Number of Connections are greater than or equal to

in the next minutes

The following graphic shows how network traffic can trigger the above correlation rule.



In this example, the system detected a connection that met the basic conditions of the correlation rule, that is, the system detected a connection from a host outside the 10.1.0.0/16 network to a host inside the network. This created a connection tracker.

The connection tracker is processed in the following stages:

1. The system starts tracking connections when it detects a connection from Host A outside the network to Host 1 inside the network.
2. The system detects two more connections that match the connection tracker signature: Host B to Host 2 and Host C to Host 1.
3. The system detects a fourth qualifying connection when Host A connects to Host 3 within the two-minute time limit. The rule conditions are met.
4. The Defense Center generates a correlation event and the system stops tracking connections.

Note that when you add a connection tracker to a correlation rule, depending on the type of event you use to trigger the rule, the connection tracker is populated with default constraints. For example, the following graphic shows a connection tracker added to a rule that triggers when the system detects the BitTorrent application protocol on your monitored network.

The connection tracker constrains the rule so that it does not trigger until more than 5MB of data is transmitted by the host, using the application protocol and protocol in the discovery event that initially violated the correlation rule, in the five minutes following the initial policy violation.

Select the type of event for this rule

If a discovery event occurs there is new information about a TCP application and it meets the following conditions:

Application Protocol is BitTorrent

Connection Tracker Remove Connection Tracker

... start tracking connections that meet the following conditions:

Responder IP is Event IP Address (switch to manual entry)

AND Application Protocol is BitTorrent

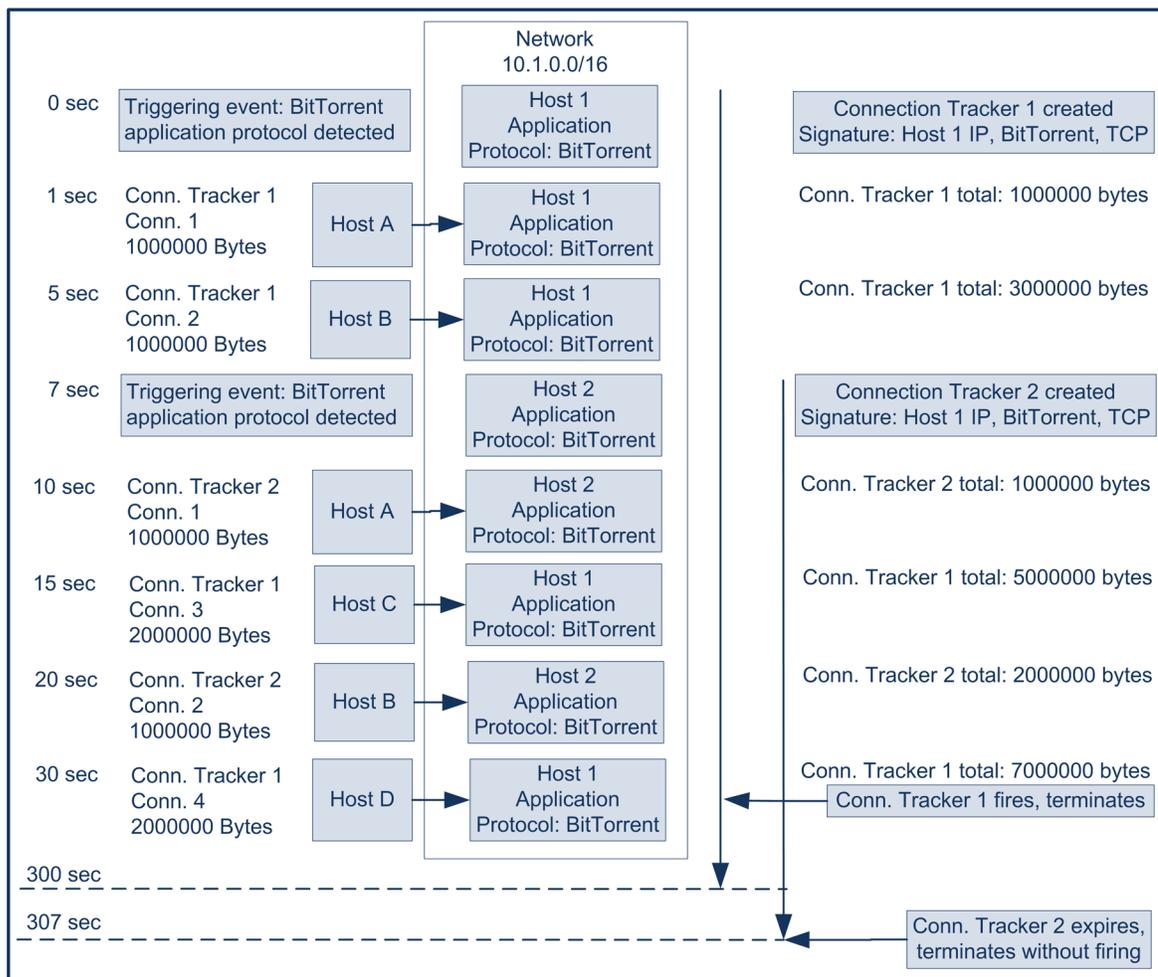
Protocol is tcp

... and generate an event if:

total Total Bytes are 5000000

in the next 5 minutes

The following graphic shows how network traffic can trigger the above correlation rule.



In this example, the system detected two connections that meet the basic conditions of the correlation rule, that is, the system detected the BitTorrent TCP application protocol on two different hosts: Host 1 and Host 2. This created two connection trackers, Connection Tracker 1 and Connection Tracker 2, each with its own signature.

The signatures represent unique combinations of protocol, IP address, and application protocol. That is, the system created two connection trackers because the IP addresses of Host 1 and Host 2 differ. If the system had detected the BitTorrent TCP application protocol on a different host, it would have created a third connection tracker. However, it would **not** have created an additional connection tracker if it detected the BitTorrent UDP application protocol on either Host 1 or Host 2. This is because the basic conditions of the correlation rule require the TCP protocol.

Connection Tracker 1 is processed in the following stages:

1. The system starts tracking connections at the 0-second marker when the system detects the BitTorrent application protocol on Host 1. Note that Connection Tracker 1 will expire if Host 1 does not transmit 5MB of BitTorrent TCP data in the next 5 minutes (by the 300-second marker).
2. The system detects that Host 1 transmitted data that matches the signature for Connection Tracker 1:
 - 1MB to Host A, at the 1-second marker
 - 2MB to Host B, at the 5-second marker
 - 2MB to Host C, at the 15-second marker

Although Host 1 has now transmitted 5MB of data, the rule does not trigger because the total number of bytes transmitted must be **more** than 5MB (**Responder Bytes are greater than 5000000**).

3. The system detects that Host 1 has transmitted 2MB via BitTorrent to Host D using the TCP protocol at the 30-second marker, for a total of 7MB transmitted by Host 1. The rule conditions are met.
4. The Defense Center generates a correlation event and stops tracking connections.

Note that the Defense Center generates the correlation event *after* Host 1 transmits the entire 2MB to Host D, because it does not tally connection data until the session terminates.

Also, note that the connection tracker terminates after the Defense Center generates the correlation event, even though the 5-minute period has not expired. If the system detects the BitTorrent TCP application protocol on Host 1 at any time after the connection tracker terminates, it will create a new connection tracker.

Connection Tracker 2 is processed in the following stages:

1. The system starts tracking connections at the 7-second marker when the system detects the BitTorrent application protocol on Host 2. Note that Connection Tracker 2 will expire if Host 2 does not transmit 5MB of BitTorrent TCP data over in the next 5 minutes (by the 307-second marker).
2. The system detects that Host 2 transmitted data, totaling 2MB, that matches the signature for Connection Tracker 2:
 - 1MB to Host A, at the 10-second marker
 - 1MB to Host B, at the 20-second marker
3. Connection Tracker 2 expires after the 5 minutes, at the 307-second marker. the system has not detected any more data transmissions that match the appropriate signature.
4. The system stops tracking connections and the Defense Center does not generate a correlation event.

To add a connection tracker:

ACCESS: Admin/Discovery Admin

1. On the Create Rule page, click **Add Connection Tracker**.
The Connection Tracker section appears.

Connection Tracker Remove Connection Tracker

... start tracking connections that meet the following conditions:

Add condition Add complex condition

X [dropdown]

... and generate an event if:

Add condition Add complex condition

X total [dropdown]

in the next 5 [minutes]

TIP! To remove a connection tracker, click **Remove Connection Tracker**.

2. Specify which connections you want to track by setting connection tracker criteria.

You can set connection tracker criteria by creating a single, simple condition, or you can create more elaborate constructs by combining and nesting conditions.

See [Understanding Rule Building Mechanics](#) on page 1499 for information on how to use the web interface to build conditions. The syntax you can use to build connection tracker conditions is described in [Syntax for Connection Trackers](#) on page 1492.

3. Based on the connections you decided to track in step 2, describe when you want to generate a correlation event.

You can create a single, simple condition that describes when you want to generate an event, or you can create more elaborate constructs by combining and nesting conditions.

You must also specify the interval (in seconds, minutes, or hours) during which the conditions you specify must be met to generate a correlation event.

See [Understanding Rule Building Mechanics](#) on page 1499 for information on how to use the web interface to build conditions. The syntax you can use to build connection tracker conditions is described in [Syntax for Connection Tracker Events](#) on page 1495.

4. Optionally, continue with the procedures in the following sections:
 - [Adding a User Qualification](#) on page 1496
 - [Adding Snooze and Inactive Periods](#) on page 1498

If you are finished building the correlation rule, continue with step 9 of the procedure in [Creating Rules for Correlation Policies](#) on page 1461 to save the rule.

Syntax for Connection Trackers

LICENSE: Any

The [Syntax for Connection Trackers](#) table describes how to build a connection tracker condition that specifies the kind of connections you want to track.

You should keep in mind that connections detected by Sourcefire managed devices and connection data exported by NetFlow-enabled devices contain different information. For example, connections detected by managed devices do not contain TCP flag information. Therefore, if you want to specify that a connection event have a certain TCP flag to trigger a correlation rule, none of the connections detected by managed devices will trigger the rule. As another example, NetFlow records do not contain information about which host in the connection is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Syntax for Connection Trackers

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Ingress Security Zone or Egress Security Zone	Select one or more security zones.
Device	Select one or more devices whose detected connections you want to track. If you want to track NetFlow connections, select the devices that process the connection data exported by your NetFlow-enabled devices.
Ingress Interface or Egress Interface	Select one or more interfaces.
Initiator IP, Responder IP, or Initiator/Responder IP	Type a single IP address or address block. For information on using IP address notation in the Sourcefire 3D System, see IP Address Conventions on page 57.
Initiator Port/ICMP Type or Responder Port/ICMP Code	Type the port number or ICMP type for initiator traffic or the port number or ICMP code for responder traffic.

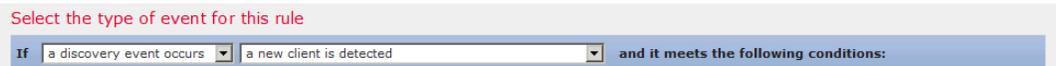
Syntax for Connection Trackers (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Transport Protocol	Type the transport protocol used by the connection: TCP or UDP .
NETBIOS Name	Type the NetBIOS name of the monitored host in the connection.
Connection Type	Select whether you want to track connections based on how they were detected: by a Sourcefire managed device (FireSIGHT) or exported by a NetFlow-enabled device (NetFlow).
Connection Duration	Type the connection duration, in seconds.
NetFlow Device	Select the IP address of the NetFlow-enabled device that exported the connections you want to track. If you did not add any NetFlow-enabled devices to your deployment, the NetFlow Device drop-down list is blank.
TCP Flags	Select the TCP flag that connections must contain in order to track them. IMPORTANT! Only connections exported by NetFlow-enabled devices contain TCP flag data.
Client	Select one or more clients.
Client Version	Type the version of the client.
Initiator Bytes, Responder Bytes, or Total Bytes	Type one of: <ul style="list-style-type: none"> the number of bytes transmitted (Initiator Bytes) the number of bytes received (Responder Bytes) the number of bytes both transmitted and received (Total Bytes)
Initiator Packets, Responder Packets, or Total Packets	Type one of: <ul style="list-style-type: none"> the number of packets transmitted (Initiator Packets) the number of packets received (Responder Packets) the number of packets both transmitted and received (Total Packets)
Username	Type the username of the user logged into either host in the connections you want to track.
Application Protocol	Select one or more application protocols.
Web Application	Select one or more web applications.
Access Control Policy	Select one or more access control policies that logged the connections you want to track.

Syntax for Connection Trackers (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Access Control Rule Name	Type all or part of the name of the access control rule that logged the connections you want to track. IMPORTANT! To track connections that match Monitor rules, type the name of the Monitor rule. The system tracks the connections, regardless of the rule or default action that later handles them.
Access Control Rule Action	Select one or more access control rule actions associated with the access control rule that logged the connections you want to track. IMPORTANT! Select Monitor to track connections that match the conditions of any Monitor rule, regardless of the rule or default action that later handles the connections.
Reason	Select one or more reasons associated with the connections you want to track.
URL Category	Select one or more URL categories for the URL visited in the connections you want to track.
URL Reputation	Select one or more URL reputation values for the URL visited in the connections you want to track
URL	Type all or part of the URL visited in the connections you want to track.

Note that you can often use event data when constructing a connection tracker. For example, assume your correlation rule triggers when the system detects a new client on one of your monitored hosts; that is, the rule triggers when a system event whose base event type is **a new client is detected** is generated.



Further assume that when you detect this new client, you want to track connections involving the new client on the host where it was detected. Because the system knows the IP address of the host and the client name, you can build a simple connection tracker that tracks those connections.

In fact, when you add a connection tracker to this type of correlation rule, the connection tracker is populated with those default constraints; that is, the

Initiator/Responder IP is set to the **Event IP Address** and the **Client** is set to the **Event Client**.



TIP! To specify that the connection tracker track connections for a specific IP address or range of IP addresses, click **switch to manual entry** to manually specify the IP. Click **switch to event fields** to go back to using the IP address in the event.

Syntax for Connection Tracker Events

LICENSE: Any

The [Syntax for Connection Tracker Events](#) table describes how to build a connection tracker condition that specifies when you want to generate a correlation event based on the connections you are tracking.

Syntax for Connection Tracker Events

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Number of Connections	Type the total number of connections detected.
Total Bytes, Initiator Bytes, or Responder Bytes	Type one of: <ul style="list-style-type: none"> the total bytes transmitted (Total Bytes) the number of bytes transmitted (Initiator Bytes) the number of bytes received (Responder Bytes)
Total Packets, Initiator Packets, or Responder Packets	Type one of: <ul style="list-style-type: none"> the total packets transmitted (Total Packets) the number of packets transmitted (Initiator Packets) the number of packets received (Responder Packets)
Unique Initiators or Unique Responders	Type one of: <ul style="list-style-type: none"> the number of unique hosts that initiated sessions that were detected (Unique Initiators) the number of unique hosts that responded to connections that were detected (Unique Responders)

Adding a User Qualification

LICENSE: FireSIGHT

If you are using a connection, intrusion, discovery, or host input event to trigger your correlation rule, you can constrain the rule based on the identity of a user involved in the event. This constraint is called a *user qualification*. You **cannot** add a user qualification to a correlation rule that triggers on a traffic profile change or on the detection of user activity.

For example, you could constrain a correlation rule so that it triggers only when the identity of the source or destination user is one from the sales department.

To add a user identity qualification:

ACCESS: Admin/Discovery Admin

1. On the Create Rule page, click **Add User Qualification**.

The User Identity Qualification section appears.



The screenshot shows a web interface for configuring a user identity qualification. At the top, it says "User Identity Qualification" with a "Remove User Qualification" button. Below that, a blue bar contains the text "Only generate an event if the user(s) involved have the following properties:". Underneath, there are two buttons: "Add condition" and "Add complex condition". Below these buttons is a dropdown menu with a red 'X' icon on the left and a downward arrow on the right.

TIP! To remove a user qualification, click **Remove User Qualification**.

2. Build the user qualification's conditions.

You can create a single, simple condition, or you can create more elaborate constructs by combining and nesting conditions. See [Understanding Rule Building Mechanics](#) on page 1499 for information on how to use the web interface to build conditions.

The syntax you can use to build conditions is described in [Syntax for User Qualifications](#) on page 1497.

3. Optionally, continue with [Adding Snooze and Inactive Periods](#) on page 1498.

If you are finished building the correlation rule, continue with step 9 of the procedure in [Creating Rules for Correlation Policies](#) on page 1461 to save the rule.

Syntax for User Qualifications

LICENSE: FireSIGHT

When you build a user qualification condition, you must first select the identity you want to use to constrain your correlation rule. The identity you can choose depends on the type of event you are using to trigger the rule, as follows:

- If you are using a connection event, select **Identity on Initiator** or **Identity on Responder**.
- If you are using an intrusion event, select **Identity on Destination** or **Identity on Source**.
- If you are using a discovery event, select **Identity on Host**.
- If you are using a host input event, select **Identity on Host**.

After you select the user type, you continue building your user qualification condition, as described in the [Syntax for User Qualifications](#) table.

The Defense Center obtains certain information about users, including first and last names, department, telephone number, and email address, from an optional Defense Center-LDAP server connection; see [Creating LDAP Connections with the Defense Center](#) on page 1298. This information may not be available for all users in the database.

Syntax for User Qualifications

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Username	Type the username of the user you want to use to constrain the correlation rule.
Authentication Protocol	Select an authentication protocol (or user type) protocol. This is the protocol that was used to detect the user.
First Name	Type the first name of the user you want to use to constrain the correlation rule.
Last Name	Type the last name of the user you want to use to constrain the correlation rule.
Department	Type the department of the user you want to use to constrain the correlation rule.
Phone	Type the telephone number of the user you want to use to constrain the correlation rule.
Email	Type the email address of the user you want to use to constrain the correlation rule.

Adding Snooze and Inactive Periods

LICENSE: Any

You can configure *snooze periods* in correlation rules. When a correlation rule triggers, a snooze period instructs the Defense Center to stop firing that rule for a specified interval, even if the rule is violated again during the interval. When the snooze period has elapsed, the rule can trigger again (and start a new snooze period).

For example, you may have a host on your network that should never generate traffic. A simple correlation rule that triggers whenever the system detects a connection involving that host may create multiple correlation events in a short period of time, depending on the network traffic to and from the host. To limit the number of correlation events exposing your policy violation, you can add a snooze period so that the Defense Center generates a correlation event only for the first connection (within a time period that you specify) that the system detects involving that host.

You can also set up inactive periods in correlation rules. During inactive periods, the correlation rule will not trigger. You can set up inactive periods to recur daily, weekly, or monthly. For example, you might perform a nightly Nmap scan on your internal network to look for host operating system changes. In that case, you could set a daily inactive period on the affected correlation rules for the time and duration of your scan so that those rules do not trigger erroneously.

The following graphic shows a portion of a correlation rule that is configured with a snooze period and an inactive period.

The screenshot shows a configuration panel titled "Rule Options". In the top right corner, there is a button labeled "Add Inactive Period". Under the "Snooze" section, the text reads "If this rule generates an event, snooze for" followed by a text input field containing "10" and a dropdown menu set to "minutes". Under the "Inactive Periods" section, there is a red "X" icon on the left. The configuration shows "Daily" in a dropdown, "at 12:00 AM" with time and AM/PM dropdowns, and "for 10 minutes" with a text input field.

To add a snooze period:

ACCESS: Admin/Discovery Admin

- ▶ On the Create Profile page, under **Rule Options**, specify the interval that the Defense Center should wait to trigger a rule again, after the rule triggers.

TIP! To remove a snooze period, specify an interval of 0 (seconds, minutes, or hours).

To add an inactive period:

ACCESS: Admin/Discovery Admin

1. On the Create Profile page, under **Rule Options**, click **Add Inactive Period**.
2. Using the drop-down lists and text field, specify when and how often you want the Defense Center to refrain from evaluating network traffic against the correlation rule.

TIP! To delete an inactive period, click the delete icon (✖) next to the inactive period you want to delete.

When you are finished adding snooze and inactive periods, continue with step 9 of the procedure in [Creating Rules for Correlation Policies](#) on page 1461 to save the rule.

Understanding Rule Building Mechanics

LICENSE: Any

You build correlation rules, connection trackers, user qualifications, and host profile qualifications by specifying the conditions under which they trigger. You can create simple conditions, or you can create more elaborate constructs by combining and nesting conditions.

For example, if you want to generate a correlation event every time a new host is detected, you can create a very simple rule with no conditions, as shown in the following graphic.

Select the type of event for this rule

If a discovery event occurs a new IP host is detected

and it meets the following conditions:

Add condition Add complex condition

✖

Detailed description: This screenshot shows a web-based interface for creating a rule. At the top, it says 'Select the type of event for this rule'. Below that, there are two dropdown menus: 'If a discovery event occurs' and 'a new IP host is detected'. Underneath, a blue bar contains the text 'and it meets the following conditions:'. Below this bar are two buttons: 'Add condition' and 'Add complex condition'. At the bottom, there is a single dropdown menu with a red 'X' icon to its left, indicating it is empty or has been deleted.

If you wanted to further constrain the rule and generate an event only if that new host was detected on the 10.4.x.x network, you can add a single condition, as shown in the following graphic.

Select the type of event for this rule

If a discovery event occurs a new IP host is detected and it meets the following conditions:

Add condition Add complex condition

✖ IP Address is in 10.4.0.0/16

Detailed description: This screenshot shows the same rule configuration interface as the previous one, but with an additional condition added. The blue bar now contains the text 'and it meets the following conditions:'. Below this bar are two buttons: 'Add condition' and 'Add complex condition'. At the bottom, there is a single dropdown menu with a red 'X' icon to its left. The dropdown menu is expanded, showing 'IP Address' selected, followed by 'is in' and a text field containing '10.4.0.0/16'.

But the following rule, which detects SSH activity on a non-standard port on the 10.4.x.x network and the 192.168.x.x network, has four conditions, with the bottom two constituting a complex condition.

The screenshot shows a configuration window titled "Select the type of event for this rule". The main configuration area is a blue bar with the text "If a discovery event occurs there is new information about a TCP application and it meets the following conditions:". Below this bar are two rows of conditions. The first row contains two conditions: "Application Protocol is SSH" and "Application Port is not 22". The second row is preceded by an "AND" dropdown and contains two conditions: "IP Address is 10.4.0.0/16" and "IP Address is 192.168.0.0/16". The second row is preceded by an "OR" dropdown. There are "Add condition" and "Add complex condition" buttons between the rows.

The syntax you can use within conditions varies depending on the element you are creating, but the mechanics are the same.

WARNING! Evaluating complex correlation rules that trigger on frequently occurring events can degrade Defense Center performance. For example, a multi-condition rule that the Defense Center must evaluate against every connection logged by the system can cause resource overload.

For more information on condition building, see:

- [Building a Single Condition](#) on page 1500
- [Adding and Linking Conditions](#) on page 1503
- [Using Multiple Values in a Condition](#) on page 1506

Building a Single Condition

LICENSE: Any

Most conditions have three parts: a *category*, an *operator*, and a *value*; some conditions are more complex and contain several categories, each of which may have their own operators and values.

For example, the following correlation rule triggers if a new host is detected on the 10.4.x.x network. The category of the condition is **IP Address**, the operator is **is in**, and the value is **10.4.0.0/16**.

The screenshot shows a configuration window titled "Select the type of event for this rule". The main configuration area is a blue bar with the text "If a discovery event occurs a new IP host is detected and it meets the following conditions:". Below this bar are two rows of conditions. The first row contains one condition: "IP Address is in 10.4.0.0/16". There are "Add condition" and "Add complex condition" buttons above the row.

To build the correlation rule trigger criteria in the example above:

ACCESS: Admin/Discovery Admin

1. Begin building a correlation rule.
For more information, see [Creating Rules for Correlation Policies](#) on page 1461.
2. On the Create Rule page, under **Select the type of event for this rule**, select a **discovery event occurs**, then select a **new IP host is detected** from the drop-down list.
3. Start building the rule's single condition by selecting **IP Address** from the first (or *category*) drop-down list.
4. Select **is in** from the operator drop-down list that appears.

TIP! When the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a range of IP addresses, as expressed in special notation such as CIDR. For information on using IP address notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

5. Type **10.4.0.0/16** in the text field.

In contrast, the following host profile qualification is more complex; it constrains a correlation rule such that the rule triggers only if the host involved in the discovery event on which the rule is based is running a version of Microsoft Windows.

Property	Operator	Value
Destination Host		
Operating System		has the following properties
OS Vendor	is	Microsoft
OS Name	is	Windows
OS Version	is	any

To build the host profile qualification in the example above:

ACCESS: Admin/Discovery Admin

1. Build a correlation rule that triggers on an discovery event.
For more information, see [Creating Rules for Correlation Policies](#) on page 1461.

2. On the Create Rule page, click **Add Host Profile Qualification**.

The Host Profile Qualification section appears.



3. Under **Host Profile Qualification**, in the first condition, specify the host whose host profile you want to use to constrain your correlation rule.
Because this host profile qualification is part of a correlation rule based on an discovery event, the only available category is **Host**.
4. Begin specifying the details of the operating system of the host by choosing the **Operating System** category.
Three subcategories appear: **OS Vendor**, **OS Name**, and **OS Version**.
5. To specify that the host can be running any version of Microsoft Windows, use the same operator for all three subcategories: **is**.
6. Finally, specify the values for the subcategories.
Select **Microsoft** as the value for **OS Vendor**, **Windows** as the value for **OS Name**, and leave **any** as the value for **OS Version**.

Note that the categories you can choose from depend on whether you are building correlation rule triggers, a host profile qualification, a connection tracker, or a user qualification. Within correlation rule triggers, the categories further depend on what kind of event is the basis for your correlation rule.

In addition, a condition's available operators depend on the category you choose. Finally, the syntax you can use to specify a condition's value depends on the category and operator. Sometimes you must type the value in a text field. Other times, you can pick a value from a drop-down list.

IMPORTANT! Where the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list. For more information, see [Using Multiple Values in a Condition](#) on page 1506.

For more information on the syntax for building correlation rule trigger criteria, see:

- [Syntax for Intrusion Events](#) on page 1467
- [Syntax for Malware Events](#) on page 1469
- [Syntax for Discovery Events](#) on page 1470
- [Syntax for User Activity Events](#) on page 1474
- [Syntax for Host Input Events](#) on page 1475

- [Syntax for Connection Events](#) on page 1476
- [Syntax for Traffic Profile Changes](#) on page 1479

For more information on the syntax for building host profile qualifications, user qualifications, and connection trackers, see:

- [Syntax for Host Profile Qualifications](#) on page 1483
- [Syntax for Connection Trackers](#) on page 1492
- [Syntax for Connection Tracker Events](#) on page 1495
- [Syntax for User Qualifications](#) on page 1497

Adding and Linking Conditions

LICENSE: Any

You can create simple correlation rule triggers, connection trackers, host profile qualifications, and user qualifications, or you can create more elaborate constructs by combining and nesting conditions.

When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

For example, the following correlation rule trigger criteria contains two conditions, linked by **OR**. This means that the rule triggers if either condition is true, that is, if a host with an IP address is not in the 10.x.x.x subnet or if a host transmits an IGMP message.

Select the type of event for this rule

If a discovery event occurs a new transport protocol is detected and it meets the following conditions:

OR

✖	Transport Protocol	is	IGMP
✖	IP Address	is not in	10.0.0.0/8

In contrast, the following rule, which detects SSH activity on a non-standard port on the 10.4.x.x network and the 192.168.x.x network, has four conditions, with the bottom two constituting a complex condition.

The screenshot shows a rule configuration window titled "Select the type of event for this rule". The rule is defined as: "If a discovery event occurs there is new information about a TCP application and it meets the following conditions:". The conditions are:

- Application Protocol is SSH
- Application Port is not 22
- IP Address is 10.4.0.0/16
- IP Address is 192.168.0.0/16

The conditions are grouped with "AND" between the first two and "OR" between the last two.

This rule triggers if SSH is detected on a non-standard port; the first two conditions demand that the application protocol name is SSH and the port is not 22. The rule further requires that the IP address of the host involved in the event is in either the 10.4.x.x network or the 192.168.x.x network.

Logically, the rule is evaluated as follows:

(A and B and (C or D))

WHERE...	IS THE CONDITION THAT STATES...
A	Application Protocol is SSH
B	Application Port is not 22
C	IP Address is in 10.4.0.0/8
D	IP Address is in 196.168.0.0/16

To add a single condition:

ACCESS: Admin/Discovery Admin

- ▶ To add a single condition, click **Add condition** above the current condition. A new condition is added below the current set of conditions, on the same level as the current set of conditions. By default, it is linked to the conditions on the same level with the **OR** operator, though you can change the operator to **AND**.

For example, if you add a simple condition to the following rule:

Select the type of event for this rule

If

and it meets the following conditions:

The result is:

Select the type of event for this rule

If and it meets the following conditions:

To add a complex condition:

ACCESS: Admin/Discovery Admin

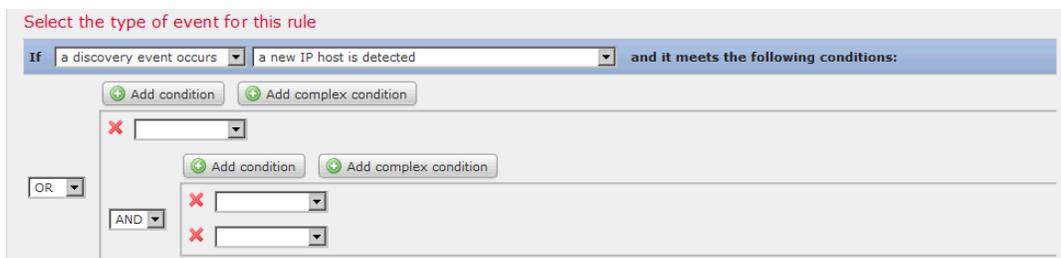
- ▶ Click **Add complex condition** above the current condition.

A complex condition is added below the current set of conditions. The complex condition comprises two subconditions, which are linked to each other with the opposite operator from the one used to link the conditions on the level above it.

For example, if you add a complex condition to the following rule:



The result is:



To link conditions:

ACCESS: Admin/Discovery Admin

- ▶ Use the drop-down list to the left of a set of conditions. Choose:
 - the **AND** operator to require that all conditions on the level it controls be met
 - the **OR** operator to require that only one of the conditions on the level it controls be met

Using Multiple Values in a Condition

LICENSE: Any

When you are building a condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list. For example, if you want to add a host profile qualification to a rule that requires that a host be running some flavor of UNIX, instead of constructing multiple conditions linked with the OR operator, use the following procedure.

To include multiple values in one condition:

ACCESS: Admin/Discovery Admin

1. Build a condition, choosing **is in** or **is not in** as the operator. The drop-down list changes to a text field.

Host Profile Qualification ✖ Remove Host Profile Qualification

Only generate an event if the host(s) involved have the following properties:

Responder Host
Operating System
 has the following properties

✖
 OS Vendor is in

OS Name is any

OS Version is any

2. Click anywhere in the text field or on the **Edit** link. A pop-up window appears.

Available Selected

12Planet
 2Wire
 A-V Tronics/NNTP-Services.
 Amazon
 Amino
 Apple
 Avaya

3. Under **Available**, use Ctrl or Shift while clicking to select multiple values. You can also click and drag to select multiple adjacent values.
4. Click the right arrow (>) to move the selected entries to **Selected**.
5. Click **OK**. The Create Rule page appears again. Your selections appear in the value field of your condition.

Host Profile Qualification ✖ Remove Host Profile Qualification

Only generate an event if the host(s) involved have the following properties:

Responder Host
Operating System
 has the following properties

✖
 OS Vendor is in "Sun", "OpenBSD", "NetBSD", "Linux"

OS Name is any

OS Version is any

Managing Rules for Correlation Policies

LICENSE: Any

Use the Rule Management page to manage correlation rules used within correlation policies.

Policy Management		Rule Management	White List	Traffic Profiles	
				Create Rule	Create Group
Rules					
Peer to Peer (11)					
Trojans (11)					
Worms					
Bagle Worm Detects Bagle worm activity					
Bugbear Worm Detects the Bugbear HTTP server backdoor					
Lovgate Worm Detects activity by the Lovgate worm backdoor component					
MyDoom Worm Detects activity by the backdoor component of MyDoom					
Netsky.S Detects the backdoor component of the NetSky.S worm.					
Sasser Worm Detects the command and transfer channels for the Sasser worm					

You can create, modify, and delete rules. You can also create rule groups to help you organize correlation rules. For more information on modifying rules, deleting rules, and creating rule groups, see:

- [Modifying a Rule](#) on page 1508
- [Deleting a Rule](#) on page 1509
- [Creating a Rule Group](#) on page 1509

For more information on creating rules, see [Creating Rules for Correlation Policies](#) on page 1461.

Modifying a Rule

LICENSE: Any

Use the following procedure to modify an existing correlation rule.

To modify an existing rule:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then select the **Rule Management** tab.
The Rule Management page appears.
2. If the rule is in a rule group, click the group name to expand the group.
3. Next to the rule you want to modify, click the edit icon ().
The Create Rule page appears.

4. Make modifications as necessary and click **Save**.
The rule is updated.

Deleting a Rule

LICENSE: Any

You cannot delete correlation rules that you are using in one or more correlation policies; you must first delete the rule from all policies in which it is included. For information on deleting a rule from a policy, see [Editing a Correlation Policy](#) on page 1520.

To delete an existing rule:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then select the **Rule Management** tab.
The Rule Management page appears.
2. If the rule is in a rule group, click the group name to expand the group.
3. Next to the rule you want to delete, click the delete icon ().
4. Confirm that you want to delete the rule.
The rule is deleted.

Creating a Rule Group

LICENSE: Any

Create rule groups to help you organize correlation rules. The Sourcefire 3D System ships with many default rules, which are grouped according to function. For example, the Worms rule group comprises rules that detect activity by common worms. Note that rule groups exist only to help you organize correlation rules; you cannot assign a group of rules to a correlation policy. Instead, add each rule individually.

You can add a rule to an existing group when you create the rule. You can also modify an existing rule to add it to a group. For more information, see the following sections:

- [Creating Rules for Correlation Policies](#) on page 1461
- [Modifying a Rule](#) on page 1508

TIP! To delete a rule group, click the delete icon () next to the group you want to delete. When you delete a rule group, rules that were in the group are **not** deleted. Rather, they merely become ungrouped

To create a rule group:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then select the **Rule Management** tab.
The Rule Management page appears.

2. Click **Create Group**.
The Create Group page appears.



3. In the **Group Name** field, type a name for the group.

4. Click **Add Group**.
The group is added.

Grouping Correlation Responses

LICENSE: Any

After you create alert responses and remediations, (see [Working with Alert Responses](#) on page 526 and [Creating Remediations](#) on page 1607), you can group them so that a policy violation triggers all of the responses within the group. Before you can assign response groups to correlation rules, you must create the groups on the Groups page.



The slider next to the group indicates whether the group is active. If you want to assign a response group to a rule within a correlation policy, you must activate it. You can sort response groups by state (active versus inactive) or alphabetically by name using the **Sort by** drop-down list.

See the following sections for more information:

- [Creating a Response Group](#) on page 1511
- [Modifying a Response Group](#) on page 1512
- [Deleting a Response Group](#) on page 1512
- [Activating and Deactivating Response Groups](#) on page 1512

Creating a Response Group

LICENSE: Any

You can place individual alerts and remediations in response groups, which can then be assigned to rules within correlation policies so that a group of alerts and remediations can be launched when a policy is violated. After a group has been assigned to rules in active policies, changes to the group and to alerts or remediations within the group are automatically applied to active policies.

To create a response group:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **Groups**.

The Groups page appears.

2. Click **Create Group**.

The Response Group page appears.

Response Group Information

Name

Active

Select Responses for Group

Available Responses		Responses in Group
Nmap_Scan_Remediation Sample Email Alert Response Sample SNMP Alert Response Sample Syslog Alert Response	> <	

Save Cancel

3. In the **Name** field, type a name for the new group.
4. Select **Active** to activate the group so that you can use it in response to a correlation policy violation.
5. From the **Available Responses** list, select the alerts and remediations you want to include in the group.

TIP! Hold down the Ctrl key while clicking to select multiple responses.

6. Click **>** to move alerts and remediations into the group.
Conversely, you can select alerts and remediations from the **Responses in Group** list and click **<** to move the alerts out of the response group.
7. Click **Save**.
The group is created.

Modifying a Response Group

LICENSE: Any

Use the following procedure to modify a response group.

To modify a response group:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **Groups**.
The Groups page appears.
2. Click the edit icon () next to the group you want to modify.
The Response Group page appears.
3. Make changes as needed and click **Save**.
If the group is active and in use, the changes you made take effect immediately.

Deleting a Response Group

LICENSE: Any

You can delete a response group if it is not used in a correlation policy. Deleting a response group does **not** delete the responses in the group, just their association with each other.

To delete a response group:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **Groups**.
The Groups page appears.
2. Click the delete icon () next to the group you want to delete.
3. Confirm that you want to delete the group.
The group is deleted.

Activating and Deactivating Response Groups

LICENSE: Any

You can temporarily deactivate a response group without deleting it. This leaves the group on the system but does not launch it when a policy to which the group is assigned is violated. Note that if a response group is used in a correlation policy when you deactivate it, it is still considered in use even though it is deactivated; you cannot delete in-use response groups.

To activate or deactivate a response group:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **Groups**.
The Groups page appears.
2. Next to the response group you want to activate or deactivate, click the slider.
If the group was activated, it is deactivated. If it was deactivated, it is activated.

Creating Correlation Policies

LICENSE: Any

After you create correlation rules or compliance white lists (or both), and, optionally, alert responses and remediations, you can use them to build correlation policies.

When your network traffic meets the criteria specified in a correlation rule or white list in an active policy, the Defense Center generates either a correlation event or white list event. It also launches any responses you assigned to the rule or white list. You can map each rule or white list to a single response or to a group of responses. If the network traffic triggers multiple rules or white lists, the Defense Center launches all the responses associated with each rule and white list.

For more information on creating the correlation rules, compliance white lists, and responses you can use to build a correlation policy, see the following sections:

- [Creating Rules for Correlation Policies](#) on page 1461
- [Creating Compliance White Lists](#) on page 1541
- [Configuring External Alerting](#) on page 524
- [Configuring Remediations](#) on page 1606

TIP! Optionally, create a skeleton policy and modify it later to add rules and responses.

To create a correlation policy:

ACCESS: Admin/Discovery Admin

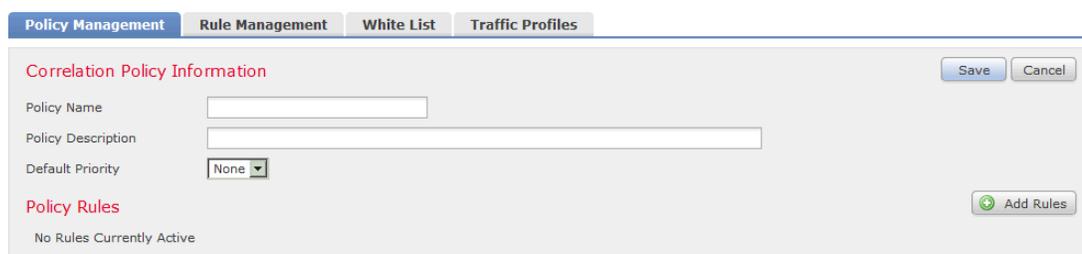
1. Select **Policies > Correlation**.

The Policy Management page appears.



2. Click **Create Policy**.

The Create Policy page appears.



3. Provide basic policy information, such as the name and description.
See [Providing Basic Policy Information](#) on page 1515.
4. Add one or more rules or white lists to the correlation policy.
See [Adding Rules and White Lists to a Correlation Policy](#) on page 1515.
5. Optionally, set rule and white list priorities.
See [Setting Rule and White List Priorities](#) on page 1516.
6. Optionally, add responses to the rules or white lists you added.
[Adding Responses to Rules and White Lists](#) on page 1517.
7. Click **Save**.
The policy is saved.

IMPORTANT! You must activate the policy before it can generate correlation and white list events and launch responses to policy violations. For more information, see [Managing Correlation Policies](#) on page 1519.

Providing Basic Policy Information

LICENSE: Any

You must give each policy an identifying name. Optionally, you can add a short description to the policy.



The screenshot shows a form titled "Correlation Policy Information" with three fields: "Policy Name" (a text input field), "Policy Description" (a larger text input field), and "Default Priority" (a dropdown menu currently set to "None").

You can also assign a user-defined priority to your policy. If your correlation policy is violated, the resultant correlation events display the priority value you assign to the policy (unless the rule that was triggered has its own priority).

IMPORTANT! Rule and white list priorities override policy priorities. For more information, see [Adding Rules and White Lists to a Correlation Policy](#) on page 1515.

To provide basic policy information:

ACCESS: Admin/Discovery Admin

1. On the Create Policy page, in the **Policy Name** field, type a name for the policy.
2. In the **Policy Description** field, type a description for the policy.
3. From the **Default Priority** drop-down list, select a priority for the policy.
You can select a priority value from 1 to 5, where 1 is highest and 5 is lowest. Or, you can select **None** to only use the priorities assigned to specific rules.
4. Continue with the procedure in the next section, [Adding Rules and White Lists to a Correlation Policy](#) on page 1515.

Adding Rules and White Lists to a Correlation Policy

LICENSE: Any

A correlation policy contains one or more correlation rules or white lists. When any rule or white list in a policy is violated, the system logs an event to the database. If you assigned one or more responses to the rule or white list, those responses are launched.

The following graphic shows a correlation policy composed of a compliance white list and a set of correlation rules, configured with a variety of responses.

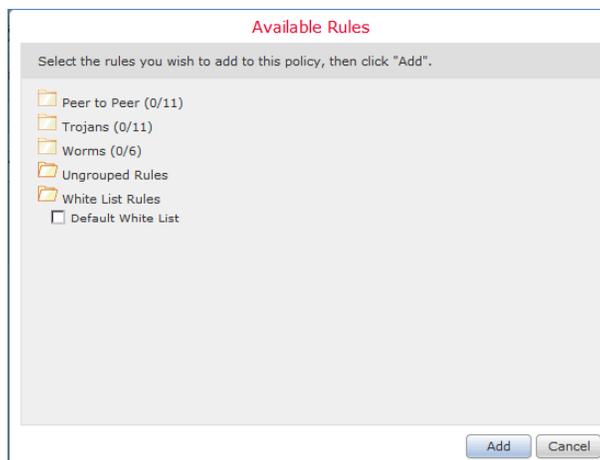
Rule	Responses	Priority	
Bugbear Worm Detects the Bugbear HTTP server backdoor	Sample Email Alert Response (Email)	Default	 
Default White List	Sample SNMP Alert Response (SNMP)	Default	 
Lovgate Worm Detects activity by the Lovgate worm backdoor component	Sample Syslog Alert Response (Syslog)	Default	 
MyDoom Worm Detects activity by the backdoor component of MyDoom	Sample Syslog Alert Response (Syslog) Sample SNMP Alert Response (SNMP) Sample Email Alert Response (Email)	Default	 
Netsky.S Detects the backdoor component of the NetSky.S worm.	This rule does not have any responses.	Default	 

To add rules or white lists to a correlation policy:

ACCESS: Admin/Discovery Admin

1. On the Create Policy page, click **Add Rules**.

The Available Rules pop-up appears.



2. Click the appropriate folder name to expand it.
3. Select the rules and white lists that you want to use in the policy and click **Add**.

The Create Policy page appears again. The rules and white lists you selected populate the policy.

4. Continue with the procedure in the next section, [Setting Rule and White List Priorities](#) on page 1516.

Setting Rule and White List Priorities

LICENSE: Any

You can assign a user-defined priority to each correlation rule or compliance white list in your correlation policy. If a rule or white list triggers, the resulting event

displays the priority you assign to the rule or white list. On the other hand, if you do not assign a priority value and the rule or white list triggers, the resulting event displays the priority value of the policy.

For example, consider a policy where the policy itself has a priority of 1 and its rules or white lists are set with the default priority, with the exception of one rule given a priority of 3. If the priority 3 rule triggers, the resulting correlation event shows 3 as its priority value. If other rules or white lists in the policy trigger, the resulting events show 1 as their priority values, retained from the policy's priority.

To set rule or white list priorities:

ACCESS: Admin/Discovery Admin

1. On the Create Policy page, from the **Priority** list for each rule or white list, select a default priority. You can select:
 - a priority value from 1 to 5, where 1 is highest and 5 is lowest
 - **None**
 - **Default** to use the policy's default priority
2. Continue with the procedure in the next section, [Adding Responses to Rules and White Lists](#) on page 1517.

Adding Responses to Rules and White Lists

LICENSE: Any

Within a correlation policy, you can map each rule or white list to a single response or to a group of responses. When any one of the rules or white lists in a policy is violated, the system logs an associated event to the database and launches the responses assigned to that rule or white list. If multiple rules or white lists within a policy trigger, the Defense Center launches the responses associated with each rule or white list.

For more information on creating responses and response groups, see:

- [Configuring External Alerting](#) on page 524
- [Configuring Remediations](#) on page 1606
- [Grouping Correlation Responses](#) on page 1510

IMPORTANT! Do **not** assign an Nmap remediation as a response to a correlation rule that triggers on a traffic profile change. The remediation will not launch.

The following graphic shows a correlation policy composed of a compliance white list and a set of correlation rules, configured with a variety of responses.

Rule	Responses	Priority	
Bugbear Worm Detects the Bugbear HTTP server backdoor	Sample Email Alert Response (Email)	Default	 
Default White List	Sample SNMP Alert Response (SNMP)	Default	 
Lovgate Worm Detects activity by the Lovgate worm backdoor component	Sample Syslog Alert Response (Syslog)	Default	 
MyDoom Worm Detects activity by the backdoor component of MyDoom	Sample Syslog Alert Response (Syslog) Sample SNMP Alert Response (SNMP) Sample Email Alert Response (Email)	Default	 
Netsky.S Detects the backdoor component of the Netsky.S worm.	This rule does not have any responses.	Default	 

To add responses to rules and white lists:

ACCESS: Admin/Discovery Admin

1. On the Create Policy page, next to a rule or white list where you want to add responses, click the responses icon .

A pop-up window appears.



2. Under **Unassigned Responses**, select the response, multiple responses, or response group you want to launch when the rule or white list triggers, and click the up arrow.

TIP! Hold down the Ctrl key while clicking to select multiple responses.

3. Click **Update**.

The Create Policy page appears again. The responses you specified are added to the rule or white list.

Managing Correlation Policies

LICENSE: Any

You manage correlation policies on the Policy Management page. You can create, modify, sort, activate, deactivate, and delete policies.



The slider next to the policy indicates whether the group is active. If you want the policy to generate correlation events and white list events, you must activate it. You can sort policies by state (active versus inactive) or alphabetically by name using the **Sort by** drop-down list.

If an active correlation policy contains a compliance white list, the following actions do **not** delete the host attribute associated with the white list, nor do they change that host attribute's values:

- deactivating the policy
- modifying the policy to remove the white list
- deleting the policy

That is, hosts that were compliant when you performed the action still appear as compliant on the host attributes network map, and so on. To delete the host attribute, you must delete its corresponding white list.

To update the white list compliance of the hosts on your network, you must either reactivate the correlation policy (if you deactivated it) or add the white list to another active correlation policy (if you deleted the white list from a correlation policy or deleted the policy itself). Note that the reevaluation of the white list that occurs when you do this does **not** generate white list events and therefore does not trigger any responses you associated with the white list. For more information on compliance white lists, see [Using the Sourcefire 3D System as a Compliance Tool](#) on page 1530.

For more information on managing correlation policies, see:

- [Activating and Deactivating Correlation Policies](#) on page 1520
- [Editing a Correlation Policy](#) on page 1520
- [Deleting a Correlation Policy](#) on page 1520

For information on creating new policies, see [Creating Correlation Policies](#) on page 1513.

Activating and Deactivating Correlation Policies

LICENSE: Any

Use the following procedure to either activate or deactivate a correlation policy.

To activate or deactivate a policy:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**.
The Policy Management page appears.
2. Next to the policy you want to activate or deactivate, click the slider.
If the policy was active, it is deactivated. If it was deactivated, it is activated.

Editing a Correlation Policy

LICENSE: Any

Use the following procedure to modify a correlation policy.

To edit a policy:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**.
The Policy Management page appears.
2. Click the edit icon () next to the policy.
The Create Policy page appears. See [Creating Correlation Policies](#) on page 1513 for information on the various configurations you can change. To remove a rule or white list from a correlation policy, on the Create Policy page, click the delete icon () next to the rule or white list you want to remove.
3. Make changes as needed and click **Save**.
The policy is changed. If the policy is active, the changes you made take effect immediately.

Deleting a Correlation Policy

LICENSE: Any

Use the following procedure to delete a correlation policy.

To delete a policy:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**.
The Policy Management page appears.

2. Click the delete icon () next to the policy you want to delete.
The policy is deleted.

Working with Correlation Events

LICENSE: Any

When a correlation rule within an active correlation policy triggers, the Defense Center generates a correlation event and logs it to the database. For information on configuring the number of correlation events saved in the database, see [Configuring Database Event Limits](#) on page 1978.

IMPORTANT! When a compliance white list within an active correlation policy triggers, the Defense Center generates a white list event. For more information, see [Working with White List Events](#) on page 1572.

For more information, see the following sections:

- [Viewing Correlation Events](#) on page 1521
- [Understanding the Correlation Events Table](#) on page 1524
- [Searching for Correlation Events](#) on page 1526

Viewing Correlation Events

LICENSE: Any

You can view a table of correlation events, then manipulate the event view depending on the information you are looking for.

The page you see when you access correlation events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of correlation events. You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The [Correlation Event Actions](#) table below describes some of the specific actions you can perform on an correlation events workflow page.

Correlation Event Actions

To...	YOU CAN...
view the host profile for an IP address	click the host profile icon that appears next to the IP address.
view user profile information	click the user icon () that appears next to the user identity. For more information, see Understanding User Details and Host History on page 1449.
sort and constrain events on the current workflow page	find more information in Sorting Drill-Down Workflow Pages on page 1835.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1836.
navigate between pages in the current workflow, keeping the current constraints	click the appropriate page link at the top left of the workflow page. For more information, see Using Workflow Pages on page 1812.
learn more about the columns that appear	find more information in Understanding the Correlation Events Table on page 1524.
modify the time and date range for displayed events	find more information in see Setting Event Time Constraints on page 1820. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

Correlation Event Actions (Continued)

To...	You CAN...
drill down to the next page in the workflow, constraining on a specific value	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • on a drill-down page that you created in a custom workflow, click a value within a row. Note that clicking a value within a row in a table view constrains the table view and does not drill down to the next page. • To drill down to the next workflow page constraining on some users, select the check boxes next to the users you want to view on the next workflow page, then click View. • To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include “Table View” in the page name.</p> <p>For more information, see Constraining Events on page 1830.</p>
delete correlation events from the system	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To delete some events, select the check boxes next to the events you want to delete, then click Delete. • To delete all events in the current constrained view, click Delete All, then confirm you want to delete all the events.
navigate to other event views to view associated events	find more information in Navigating Between Workflows on page 1836.

To view correlation events:

ACCESS: Admin/Any Security Analyst

- ▶ Select **Analysis > Correlation > Correlation Events**.

The first page of the default correlation events workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

TIP! If you are using a custom workflow that does not include the table view of correlation events, click **(switch workflow)**, then select **Correlation Events**.

Understanding the Correlation Events Table

LICENSE: Any

When a correlation rule triggers, the Defense Center generates a correlation event. The fields in the correlation events table are described in the following table.

Correlation Event Fields

FIELD	DESCRIPTION
Time	The date and time that the correlation event was generated.
Impact Flag	The impact flag assigned to the correlation event based on the correlation between intrusion data, discovery data, and vulnerability information. For more information, see Using Impact Flags to Evaluate Events on page 641.
Inline Result	<p>One of:</p> <ul style="list-style-type: none"> • a black down arrow, indicating that the system dropped the packet that triggered the intrusion rule • a gray down arrow, indicating that the system would have dropped the packet in an inline, switched, or routed deployment if you enabled the Drop when Inline intrusion policy option • blank, indicating that the triggered intrusion rule was not set to Drop and Generate Events <p>Note that the system does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of the rule state or the drop behavior of the intrusion policy. For more information, see Setting Rule States on page 743, Setting Drop Behavior in an Inline Deployment on page 689, Configuring Passive Interfaces on page 267, and Tap Mode on page 275.</p>
Source IP or Destination IP	The IP address of the source or destination host in the event that triggered the policy violation.
Source User or Destination User	The name of the user logged in to the source or destination host in the event that triggered the policy violation
Source Port/ICMP Type or Destination Port/ICMP Code	The source port or ICMP type for the source traffic or the destination port or ICMP code for destination traffic associated with the event that triggered the policy violation.
Description	<p>The description of the correlation event. The information in the description depends on how the rule was triggered.</p> <p>For example, if the rule was triggered by an operating system information update event, the new operating system name and confidence level appears.</p>
Policy	The name of the policy that was violated.

Correlation Event Fields (Continued)

FIELD	DESCRIPTION
Rule	The name of the rule that triggered the policy violation.
Priority	The priority specified by the policy or rule that triggered the policy violation.
Src Host Criticality or Dst Host Criticality	The user-assigned host criticality of the source or destination host involved in the correlation event: None , Low , Medium , or High . Note that only correlation events generated by rules based on discovery events, host input events, or connection events contain a source host criticality. For more information on host criticality, see Working with the Predefined Host Attributes on page 1369.
Ingress Security Zone or Egress Security Zone	The ingress or egress security zone in the intrusion or connection event that triggered the policy violation.
Device	The name of the device that generated the event that triggered the policy violation.
Ingress Interface or Egress Interface	The ingress or egress interface in the intrusion or connection event that triggered the policy violation.
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

For more information on displaying the correlation events table, see the following:

- [Viewing Correlation Events](#) on page 1521
- [Searching for Correlation Events](#) on page 1526

Searching for Correlation Events

LICENSE: Any

You can search for specific correlation events. You may want to create searches customized for your network environment, then save them to reuse later. The following table describes the search criteria you can use.

Correlation Event Search Criteria

FIELD	SEARCH CRITERIA RULES
Policy	Type the name of the correlation policy you want to search for.
Rule	Type the name of the correlation rule you want to search for.
Description	Type all or part of the correlation event description. The information in the description depends on the event that caused the rule to trigger.
Priority	Specify the priority of the correlation event, which is determined by the priority of either the triggered rule or the violated correlation policy. Enter none for no priority. For information on setting correlation rule and policy priorities, see Providing Basic Policy Information on page 1515 and Setting Rule and White List Priorities on page 1516.
Source IP Destination IP, or Source/Destination IP	Specify the IP address of the source, destination, or source or destination hosts in the event that triggered the policy violation. You can specify a single IP address or address block, or a comma-separated list of either or both. You can also use negation. See Specifying IP Addresses in Searches on page 1776 for more information.
Source User or Destination User	Specify the user logged in to the source or destination host in the event that triggered the policy violation.
Source Port/ICMP Type or Destination Port/ICMP Code	Specify the source port or ICMP type for source traffic or destination port or ICMP code for destination traffic associated with the event that triggered the policy violation.
Impact Flag	Specify the impact flag assigned to the correlation event. Valid case-insensitive values are Impact Level 0 , Impact Level 1 , Impact Level 2 , Impact Level 3 , and Impact Level 4 . Do not use impact flag colors or partial strings (for example, do not use blue , level 1 , or 0). For more information, see Using Impact Flags to Evaluate Events on page 641.

Correlation Event Search Criteria (Continued)

FIELD	SEARCH CRITERIA RULES
Inline Result	<p>For policy violations triggered by intrusion events, type either:</p> <ul style="list-style-type: none"> • dropped, to specify whether the packet was dropped in an inline, switched, or routed deployment • would have dropped, to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline, switched, or routed deployment <p>Note that the system does not drop packets in a passive deployment, including when an inline set is in tap mode, regardless of the rule state or the drop behavior of the intrusion policy. For more information, see Setting Rule States on page 743, Setting Drop Behavior in an Inline Deployment on page 689, and Tap Mode on page 275.</p>
Source Host Criticality or Destination Host Criticality	<p>Specify the host criticality of the source or destination host involved in the policy violation: None, Low, Medium, or High. Note that only correlation events generated by rules based on discovery events, host input events, or connection events contain a source host criticality. For more information on host criticality, see Working with the Predefined Host Attributes on page 1369.</p>
Ingress Security Zone, Egress Security Zone, or Ingress/Egress Security Zone	<p>Specify the ingress, egress, or ingress or egress security zone in the intrusion or connection event that triggered the policy violation.</p>
Device	<p>Type the name, group name, or IP address of the device that generated the event that triggered the policy violation. See Managing Devices on page 190, Editing Assigned Device Names on page 244, and Managing Device Groups on page 217.</p>
Ingress Interface or Egress Interface	<p>Specify the ingress or egress interface in the intrusion or connection event that triggered the policy violation.</p>

To search for correlation events:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.

- From the **Table** drop-down list, select **Correlation Events**.
The page reloads with the appropriate constraints.

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: Correlation Events

Name: Search 1, My Search

Save As Private:

Constraint

Policy: Policy 1, My Policy

Rule: My R

Egress Security Zone: My Security Zone

Ingress / Egress Security Zone: My Security Zone

Device: device1.example.com, *.example.com, 192.168.1.3

Ingress Interface: s1p1

Egress Interface: s1p1

Search Save As New Search

- Optionally, if you want to save the search, enter a name for the search in the **Name** field.

If you do not enter a name, one is created automatically when you save the search.

- Enter your search criteria in the appropriate fields, as described in the [Correlation Event Search Criteria](#) table:
 - All fields accept negation (!).
 - All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
 - Many fields accept one or more asterisks (*) as wild cards.
 - Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
 - Click the add object icon (⊕) that appears next to a search field to use an object as a search criterion.

For more information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

- If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default correlation events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**) so that you can run it at a later time.

CHAPTER 37

USING THE SOURCEFIRE 3D SYSTEM AS A COMPLIANCE TOOL

A *compliance white list* (or *white list*) is a set of criteria that allows you to specify the operating systems, applications, and protocols that are allowed to run on a specific subnet, and automatically generate an event if a host on the subnet violates the white list. For example, your security policy might state that while your web servers are allowed to run HTTP, none of the other hosts on your network are. You could create a white list that evaluates your entire network, excluding your web farm, to determine which hosts are running HTTP.

Note that you could create a correlation rule that performs this function by configuring the rule so that it triggers when:

- the system discovers new information about an application protocol
- the application protocol name is http
- the IP address of the host involved in the event is not in your web farm

The screenshot shows a rule configuration window. At the top, it says "If a discovery event occurs there is new information about a TCP application". Below this, it says "and it meets the following conditions:". There are two buttons: "Add condition" and "Add complex condition". The conditions are listed as follows:

AND	Application	is	HTTP
	IP Address	is in	10.0.0.0/8
	IP Address	is not in	10.1.0.0/16

However, correlation rules, which provide you with a more flexible way of alerting you and responding to policy violations on your network, are more complex to configure and maintain than white lists. Correlation rules are also wider in scope, allowing you to generate a correlation event when one of many types of event meets any criteria that you specify. On the other hand, white lists are specifically

meant to help you evaluate the operating systems, application protocols, clients, web applications, and protocols that are running on your network and whether that violates your organization's policies.

You can create custom white lists that meet your specific needs, or you can use the default white list created by the Sourcefire Vulnerability Research Team (VRT) that contains recommended settings for allowed operating systems, application protocols, clients, web applications, and protocols. You may also want to customize the default white list for your network environment.

If you add a white list to an active correlation policy, when the system detects that a host is violating the white list, the system logs a white list event — which is a special kind of correlation event — to the database. Further, you can configure the system to trigger responses (remediations and alerts) automatically when it detects a white list violation.

IMPORTANT! Although you can configure the network discovery policy to add hosts and application protocols to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts and application protocols is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. This may affect the way you build compliance white lists. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Because the system creates a host attribute for each host that indicates whether it is in compliance with any white lists you create, you can obtain an at-a-glance summary of the compliance of your network. In just a few seconds, you can determine exactly which hosts in your organization are running HTTP in violation of your policy, and take appropriate action.

Then, using the correlation feature, you can configure the system to alert you whenever a host that is not in your web farm starts running HTTP.

In addition, the system allows you to use host profiles to determine whether an individual host is violating any of the white lists you have configured, and in which way it is violating the white list. The Sourcefire 3D System also includes workflows that allow you to view each of the individual white list violations, as well as the number of violations per host.

Finally, you can use the dashboard to monitor recent system-wide compliance activity, including white list events and summary views of the overall white list compliance of your network.

For more information on creating and managing compliance white lists and on interpreting white list events and violations, see the following sections:

- [Understanding Compliance White Lists](#) on page 1532
- [Creating Compliance White Lists](#) on page 1541
- [Managing Compliance White Lists](#) on page 1563
- [Working with Shared Host Profiles](#) on page 1564

- [Working with White List Events](#) on page 1572
- [Working with White List Violations](#) on page 1579

In addition, see the following chapters and sections for more information:

- [Creating Correlation Policies](#) on page 1513 explains how to create and configure correlation policies that include compliance white lists, and explains how to assign responses and priorities to the white lists.
- [Using Host Profiles](#) on page 1334 explains how to use a host's profile to determine whether it is violating any white lists.
- [Using Dashboards](#) on page 66 explains how to obtain an at-a-glance view of your current system status, including white list compliance activity.

Understanding Compliance White Lists

LICENSE: FireSIGHT

A *compliance white list* is a set of criteria that specify the operating systems, clients, application protocols, web applications, and protocols that are allowed to run on your network. You can create custom white lists that meet your specific needs, or you can use the default white list created by the VRT that contains recommended settings.

Custom white list criteria can be simple; you can specify that only hosts running a certain operating system are allowed. Your criteria can also be complex; you can specify that while all operating systems are allowed, only hosts running a certain operating system are allowed to run a certain application protocol on a specific port.

White lists comprise two main parts: *targets* and *host profiles*. The targets are the specific hosts that are evaluated by the white list, while the host profiles specify the operating systems, clients, application protocols, web applications, and protocols are allowed to run on the targets.

After you create a white list and add it to an active correlation policy, the system evaluates the white list's targets against its host profiles to determine whether the targets are in compliance with the white list. After this initial evaluation, the system generates a *white list event* when it detects that a valid target is violating the white list.

For more information, see the following sections:

- [Understanding White List Targets](#) on page 1533 explains how white lists only target the hosts that you specify.
- [Understanding White List Host Profiles](#) on page 1534 explains the different kinds of profiles that describe which clients, application protocols, web applications, and protocols are allowed to run on your network.

- [Understanding White List Evaluations](#) on page 1538 explains how the system evaluates the hosts on your network against white lists, and how you can tell which hosts are in compliance and which are not.
- [Understanding White List Violations](#) on page 1539 explains how the system detects and notifies you of white list violations.

Understanding White List Targets

LICENSE: FireSIGHT

When you create a white list, you first specify the portions of your network it applies to. You can use a white list to evaluate all the hosts on your monitored network, or you can restrict the white list to evaluate only certain network segments or even individual hosts. You can further restrict the white list so that it evaluates only hosts that have a certain host attribute or that belong to a certain VLAN. A host that is eligible to be evaluated by a white list is called a *valid target* (or *target*). A valid target:

- must be in one of the IP address ranges you specify. You can also exclude ranges of IP addresses.
- must have at least one of the host attributes you specify.

For example, you could configure a white list to evaluate only hosts that have a high host criticality. For information on host attributes, including host criticality, see [Working with User-Defined Host Attributes](#) on page 1370 and [Working with the Predefined Host Attributes](#) on page 1369.

- must belong to one of the VLANs you specify.

If a host does not meet all of these criteria, it is not evaluated against the white list, regardless of whether its host profile is in violation of the white list.

If your white list contains more than one target, a host must meet the criteria specified in only one of them to be considered valid. For example, if you create a target that includes the 10.10.x.x network and one that excludes the 10.10.x.x network, a host in that network is considered a valid target. Note that if your white list does not contain any targets, none of the hosts on your network will be evaluated against the white list.

The targets for your white list are listed on the left of the Create White List page.



Note that the default white list uses targets of 0.0.0.0/0 and ::/0, which represent the entire monitored network. If you choose to use this white list, you can leave the target as-is or modify it to reflect your network environment.

For information on creating white list targets, see [Configuring Compliance White List Targets](#) on page 1545.

Understanding White List Host Profiles

LICENSE: FireSIGHT

After you specify which targets the white list evaluates, the next step is to configure *host profiles*. Host profiles in a white list specify which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the target hosts.

There are three kinds of host profiles you can configure in a white list: global host profiles, host profiles for specific operating systems, and shared host profiles. Each type of host profile appears differently when you are creating a white list.



The following table explains how to identify and access the different kinds of host profiles.

Accessing Compliance White List Host Profiles

TO VIEW...	UNDER ALLOWED HOST PROFILES, CLICK...
the global host profile for the white list	Any Operating System
a host profile for a specific operating system	a host profile name that is listed in plain text rather than italics
a shared host profile used by the white list	a host profile name that is listed in italics

For more information, see the following sections:

- [Understanding the Global Host Profile](#) on page 1535
- [Understanding Host Profiles for Specific Operating Systems](#) on page 1535
- [Understanding Shared Host Profiles](#) on page 1537

Understanding the Global Host Profile

LICENSE: FireSIGHT

Every white list contains a global host profile, which specifies the application protocols, clients, web applications, and protocols that are allowed to run on target hosts, regardless of the host's operating system.



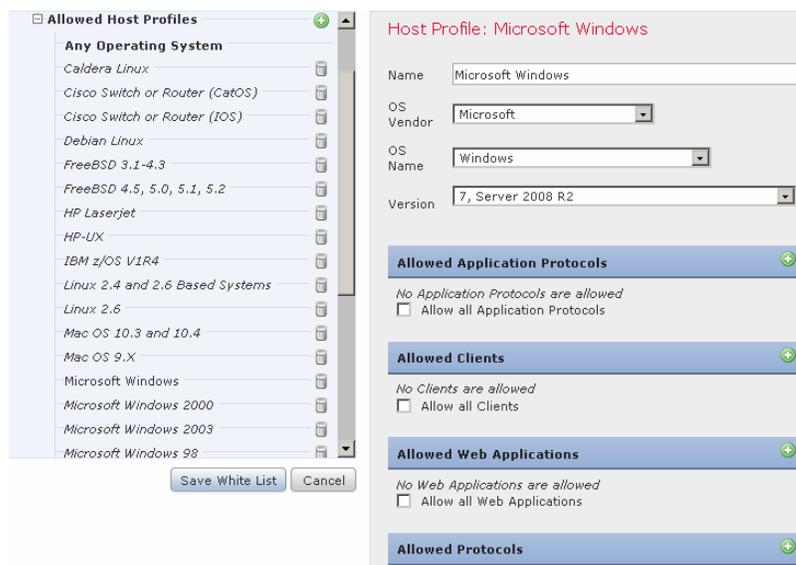
For example, instead of editing multiple Microsoft Windows and Linux host profiles to allow Internet Explorer, you can configure the global host profile to allow it regardless of the operating system on which it was detected. Note that the ARP, IP, TCP, and UDP protocols are always allowed to run on every host; you cannot disallow them. For more information, see [Configuring the Global Host Profile](#) on page 1549.

Understanding Host Profiles for Specific Operating Systems

LICENSE: FireSIGHT

You must create one host profile for each operating system you want to allow on your network. To disallow an operating system on your network, do not create a host profile for that operating system. For example, to make sure that all the

hosts on your network are running Microsoft Windows, configure the white list to only contain host profiles for that operating system.



When you create a host profile for an operating system, you can also require that it have a particular version. For example, you could require that compliant hosts run Windows 7 or Server 2008 R2.

After you create a host profile for a particular operating system, you can specify the application protocols, clients, web applications, and protocols that are allowed to run on target hosts running that operating system. For example, you could allow SSH to run on Linux hosts on port 22. You could also restrict the particular vendor and version to OpenSSH 4.2.

Note that unidentified hosts remain in compliance with all white lists until they are identified. You can, however, create a white list host profile for unknown hosts.

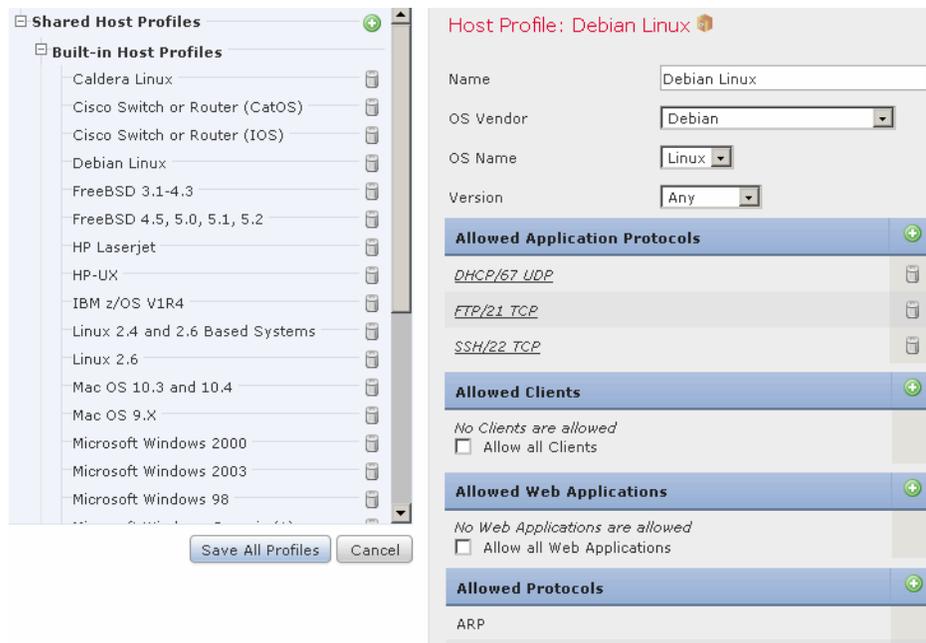
IMPORTANT! Unidentified hosts are not the same as unknown hosts. *Unidentified* hosts are hosts about which the system has not yet gathered enough information to identify their operating systems. *Unknown* hosts are hosts whose traffic has been analyzed by the system, but whose operating systems do not match any of the known fingerprints.

For more information, see [Creating Host Profiles for Specific Operating Systems](#) on page 1550.

Understanding Shared Host Profiles

LICENSE: FireSIGHT

Shared host profiles are tied to specific operating systems, but you can use each shared host profile in more than one white list. That is, if you create multiple white lists but want to use the same host profile to evaluate hosts running a particular operating system across the white lists, use a shared host profile.



For example, if you have offices worldwide and you want to create a separate white list for each location, but always want to use the same profile for all hosts running Apple Mac OS X, you can create a shared profile for that operating system and use it in all your white lists.

The default white list represents recommended “best practices” settings for allowed operating systems, clients, application protocols, web applications, and protocols. This white list uses a special category of shared host profiles, called *built-in host profiles*.



Note that built-in host profiles are marked with the built-in host profile icon (🏠).

Host Profile: Caldera Linux 🏠

Built-in host profiles use built-in application protocols, protocols, and clients. You can use these elements as-is in both the default white list and in any custom

white list that you create or you can modify them to suit your needs. They are displayed in italics within the built-in host profile and in any other host profile that uses them.



Keep in mind that like all shared host profiles, if you modify a built-in host profile, it affects every white list that uses it. Likewise, if you modify a built-in application protocol, protocol, or client, it affects every white list that uses it.

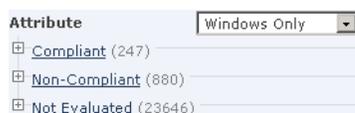
For more information on shared host profiles, [Working with Shared Host Profiles](#) on page 1564.

Understanding White List Evaluations

LICENSE: FireSIGHT

After you create white list host profiles and save the white list, you can add the white list to a correlation policy, just as you would a correlation rule. For more information, see [Configuring Correlation Policies and Rules](#) on page 1459.

After you activate the correlation policy, the system evaluates the targets of the white list against the white list criteria. You can then use the host attributes network map to gain an overall view of the white list compliance of the hosts on your network.



Every host on the network is assigned a host attribute that has the same name as the white list. This host attribute has one of the following values:

- **Compliant**, for valid targets that are compliant with the white list
- **Non-Compliant**, for valid targets that violate the white list
- **Not Evaluated**, for invalid targets and hosts that have not yet been evaluated for any reason

Note that if your network is large and the system is in the process of evaluating all the valid targets in the network map against the white list, targets that have not yet been evaluated are marked as **Not Evaluated**. As the system completes its processing, more hosts move from **Not Evaluated** to either **Compliant** or **Non-Compliant**. The system can evaluate approximately 100 hosts per second.

Additionally, a host may be marked as **Not Evaluated** if the system has insufficient information to determine whether the host is in compliance. For example, this may occur if the system has detected a new host but has not yet

gathered relevant information on the operating system, clients, application protocols, web applications, or protocols running on the host.

IMPORTANT! If you change or delete a host attribute from a host and that change or deletion means that the host is no longer a valid target, the host changes from either **Compliant** or **Non-Compliant** to **Not Evaluated**.

For more information on host attributes, see [Working with the Host Attributes Network Map](#) on page 1325.

Understanding White List Violations

LICENSE: FireSIGHT

After the initial white list evaluation, the system generates a *white list event* when it detects that a valid target is violating the white list. White list events are a special kind of correlation event, and are logged to the Defense Center correlation event database. You can view white list events in a workflow, or search for specific white list events. For more information, see [Working with White List Events](#) on page 1572.

White list violations occur when the system generates an event that indicates that a host is out of compliance. Similarly, discovery events may indicate that a previously non-compliant host is now compliant, although the system does **not** generate a white list event when this occurs.

The following events can change the compliance of a host:

- the system detects a change in a host's operating system
- the system detects an identity conflict for a host's operating system or an application protocol on the host
- the system detects a new TCP server port (for example, a port used by SMTP or web servers) active on a host, or a new UDP server running on a host
- the system detects a change in a discovered TCP or UDP server running on a host, for example, a version change due to an upgrade
- the system detects a new client running on a host
- the system drops a client from its database due to inactivity
- the system detects a new web application running on a host
- the system drops a web application from a host profile due to inactivity
- the system detects that a host is communicating with a new network protocol, such as Novell Netware or IPv6, or a new transport protocol, such as ICMP or EGP
- the system detects a new mobile device that is jailbroken
- the system detects that a TCP or UDP port has closed or timed out on a host

In addition, you can trigger a compliance change for a host by using the host input feature or the host profile to:

- add a client, protocol, or server to a host
- delete a client, protocol, or server from a host
- set the operating system definition for a host
- change a host attribute for a host so that the host is no longer a valid target

For example, if your white list specifies that only Microsoft Windows hosts are allowed on your network, and the system detects that the host is now running Mac OS X, the system generates a white list event. In addition, the host attribute associated with the white list changes its value from **Compliant** to **Non-Compliant** for that host.

For the host in this example to come back into compliance, one of the following must occur:

- you edit the white list so that the Mac OS X operating system is allowed
- you manually change the operating system definition of the host to Microsoft Windows
- the system detects that the operating system has changed back to Microsoft Windows

In any case, the host attribute associated with the white list changes its value from **Non-Compliant** to **Compliant** for that host.

As another example, if your compliance white list disallows the use of FTP, and you then delete FTP from the application protocols network map or from an event view, hosts running FTP become compliant. However, if the system detects the application protocol again, the system generates a white list event and the hosts become non-compliant.

Note that if the system generates an event that contains insufficient information for the white list, the white list does not trigger. For example, consider a scenario where your white list specifies that you allow only TCP FTP traffic on port 21. Then, the system detects that port 21, using the TCP protocol, has become active on one of the white list targets, but the system is unable to determine whether the traffic is FTP. In this scenario, the white list does not trigger until either the system identifies the traffic as something other than FTP traffic or you use the host input feature to designate the traffic as non-FTP traffic.

IMPORTANT! During the initial evaluation of a white list, the system does **not** generate white list events for non-compliant hosts. If you want to generate white list events for all non-compliant targets, you must purge the Defense Center database. This causes the hosts on your network and their associated clients, application protocols, web applications, and protocols to be rediscovered, which may trigger white list events. For more information, see [Purging Discovery Data from the Database](#) on page 2236.

Finally, you can configure the system to trigger responses automatically when it detects a white list violation. Responses include remediations (such as running an Nmap scan), alerts (email, SNMP, and syslog alerts), or combination of alerts and remediations. For more information, see [Adding Responses to Rules and White Lists](#) on page 1517.

Creating Compliance White Lists

LICENSE: FireSIGHT

When you create a white list, you can survey either your entire network or a specific network segment. Surveying the network populates the white list with one host profile for each operating system that the system has detected on the network segment. By default, these host profiles allow all of the clients, application protocols, web applications, and protocols that the system has detected on the applicable operating systems.

Then, you must specify the targets of the white list. You can configure a white list to evaluate all the hosts on your monitored network, or you can restrict the white list to evaluate only certain network segments or even individual hosts. You can further restrict the white list so that it evaluates only hosts that have a certain host attribute or that belong to a certain VLAN. If you surveyed your network, by default the network segment that you surveyed represents the white list targets. You can edit or delete the surveyed network, or you can add new targets.

Next, create host profiles that represent compliant hosts. Host profiles in a white list specify which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the target hosts. You can configure the global host profile, edit the host profiles created by any network survey you performed, as well as add new host profiles, and add and edit shared host profiles.

Finally, save the white list and add it to an active correlation policy. The system begins evaluating the target hosts for compliance, generating white list events when a host violates the white list, and triggering any responses you have configured to white list violations. For a more detailed introduction to compliance white lists, see [Understanding Compliance White Lists](#) on page 1532.

TIP! You can also create a white list from a table view of hosts. For more information, see [Creating a Compliance White List Based on Selected Hosts](#) on page 1408.

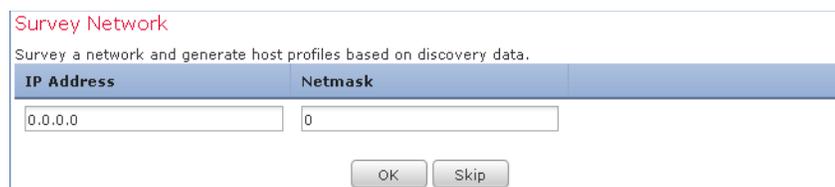
To create a compliance white list:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.
The White List page appears.

2. Click **New White List**.

The Survey Network page appears.



3. Optionally, survey your network:

- To survey your network, see [Surveying Your Network](#) on page 1543.
- To create a white list without surveying your network, click **Skip** and continue with the next step.

The Create White List page appears.



4. In the **Name** field, type a name for the new white list.

5. In the **Description** field, type a short description of the white list.

6. To allow jailbroken mobile devices on your network, enable **Allow Jailbroken Mobile Devices**. To cause all jailbroken devices evaluated by the white list to generate a white list violation, disable the option.

7. Specify the targets for the white list. You can edit or delete the targets created by a network survey as well as add new targets. Optionally, further restrict targets based on host attributes or VLAN ID. For more information, see [Configuring Compliance White List Targets](#) on page 1545.

8. Create host profiles that represent compliant hosts. You can configure the global host profile, edit the host profiles created by a network survey, as well as add new host profiles and add and edit shared host profiles. For more information, see [Configuring Compliance White List Host Profiles](#) on page 1548.

9. Click **Save White List** to save your white list.

The white list is saved. You can now add it to an active correlation policy to begin evaluating the target hosts for compliance, generating white list events when a host violated the white list, and, optionally triggering responses to white list violations. For more information, see [Creating Correlation Policies](#) on page 1513.

Surveying Your Network

LICENSE: FireSIGHT

When you begin creating a compliance white list, you can survey either your entire network or a specific network segment.

Surveying your network gathers data from the database about the application protocols, clients, web applications, and protocols running on the different detected operating systems. Then, the system creates one host profile within the white list for each detected operating system. By default, these host profiles allow all of the detected clients, application protocols, web applications, and protocols that the system has detected on each applicable operating systems.

This creates a baseline white list so that you do not have to manually create and configure multiple host profiles. After you survey your network, you can then edit or delete the host profiles that the survey created to suit your needs; you can also add any other host profiles you might need.

Note that you can survey your network at any time during the white list creation process. This can add additional allowed clients, application protocols, web applications, and protocols to the host profiles that already exist, and can create additional host profiles if the survey detects hosts running operating systems that were not detected during the initial survey. If you resurvey your network within a white list that is used within an active correlation policy, and the survey changes either your targets or host profiles, the target hosts are re-evaluated when you save the white list. Although this re-evaluation may bring some hosts into compliance, it does not generate any white list events.

To begin creating a compliance white list by surveying your network:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.

The White List page appears.

2. Click **New White List**.

The Survey Network page appears.

IP Address	Netmask
0.0.0.0	0

OK Skip

3. Do you want to survey your network?

- If **yes**, continue with the next step.
- If **no**, click **Skip**.

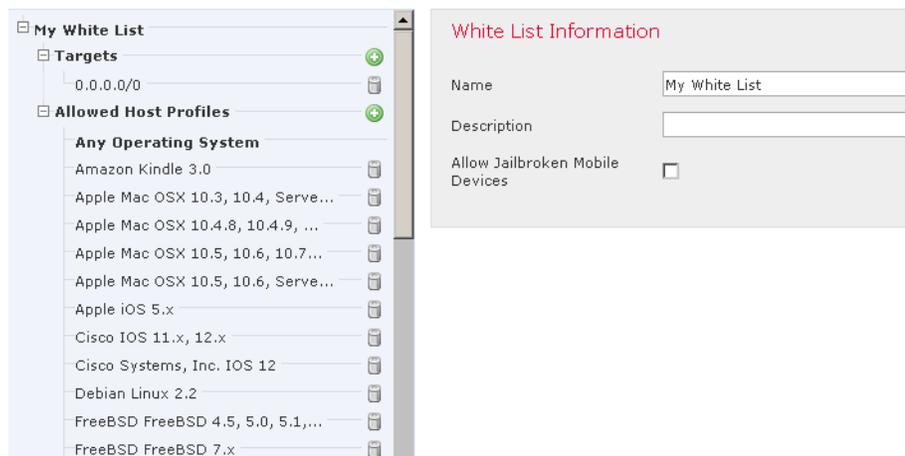
The Create White List page appears and displays a blank white list. Continue with the procedure in the next section, [Providing Basic White List Information](#).

4. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in special notation such as CIDR) that represent the hosts you want to survey.

Make sure to specify a network that you configured the system to monitor in the network discovery policy. For information on using IP address notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

TIP! To survey the entire monitored network, use the default values of 0.0.0.0/0 and ::/0.

5. Click **OK**.
The Create White List page appears.



The white list is pre-populated; its targets are the hosts in the network you surveyed and its allowed host profiles are those of the targets.

6. To survey additional networks, click **Survey Network** and repeat steps 4 and 5 for each additional network you want to survey.

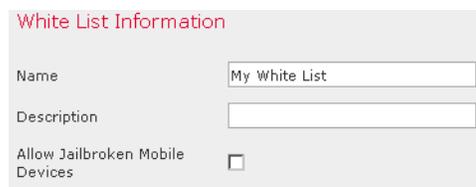
Surveying an additional network can add additional allowed clients, application protocols, web applications, and protocols to the host profiles that already exist, and can create additional host profiles if the survey detects hosts running operating systems that were not detected during the initial survey. Surveying an additional network also adds a target to the white list that represents the hosts in the network segment that you surveyed. You can then edit or delete this target.

7. Continue with the next section, [Providing Basic White List Information](#).

Providing Basic White List Information

LICENSE: FireSIGHT

You must give each white list a name, and, optionally, a short description. In addition, you can choose whether jailbroken mobile devices should cause a white list violation.



White List Information

Name

Description

Allow Jailbroken Mobile Devices

To provide basic white list information:

ACCESS: Admin

1. In the **Name** field, type a name for the new white list.
2. In the **Description** field, type a short description of the white list.
3. To allow jailbroken mobile devices on your network, enable **Allow Jailbroken Mobile Devices**. To cause all jailbroken devices evaluated by the white list to generate a white list violation, disable the option.
4. Continue with the next section, [Configuring Compliance White List Targets](#).

Configuring Compliance White List Targets

LICENSE: FireSIGHT

When you create a compliance white list, you must specify the portions of your network it applies to. You can use a white list to evaluate all the hosts on your monitored network, or you can restrict the white list to evaluate only certain network segments or even individual hosts. You can further restrict the white list so that it evaluates only hosts that have a certain host attribute or that belong to a certain VLAN. A host that is eligible to be evaluated by a white list is called a *target*. For a more detailed introduction to white list targets, see [Understanding White List Targets](#) on page 1533.

When you are finished creating compliance white list targets, continue with [Configuring Compliance White List Host Profiles](#) on page 1548.

IMPORTANT! If you change or delete a host attribute from a host and that modification means that the host is no longer a valid target, the host is no longer evaluated by the white list and is considered neither compliant nor non-compliant.

For information on how to modify and delete targets, see:

- [Modifying Existing Targets](#) on page 1548
- [Deleting Existing Targets](#) on page 1548

When you create a target for a compliance white list, you specify the criteria a host must meet to be evaluated against the white list. A valid target:

- must be in one of the IP address ranges you specify. You can also exclude ranges of IP addresses.
- must have at least one of the host attributes you specify.
- must belong to one of the VLANs you specify.

Note that if you add a target to a white list that is used by an active correlation policy, after you save the white list, the new target hosts are evaluated for compliance. However, this evaluation does not generate white list events.

To create a compliance white list target:

ACCESS: Admin

1. On the Create White List Page, next to **Targets**, click the add icon (+).
The settings for the new target appear.

Target: New Target

Name

Targeted Networks +

IP Address	Netmask	Exclude
There are no network restrictions		

Targeted Host Attributes +

Attribute	Value
There are no Host Attribute restrictions	

Targeted VLANs +

VLAN ID
No VLAN restrictions

TIP! You can also create a new target by surveying a network segment. On the Create White List page, click **Survey Network**, then follow steps 4 and 5 in [Surveying Your Network](#) on page 1543. The new target is created and is named according to the IP addresses you specified. Click the target you just created and continue with the rest of this procedure to rename the target, add or exclude additional networks, and add host attribute or VLAN restrictions.

2. In the **Name** field, type a name for the new target.
3. Target a specific set of IP addresses by clicking **Add** next to **Targeted Networks**.

Targeted Networks +

IP Address	Netmask	Exclude
<input type="text"/>	<input type="text"/>	<input type="checkbox"/>

4. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in special notation, such as CIDR) that represent the hosts you want to target or exclude from targeting.

You should make sure that you specify a network that you configured the system to monitor in your network discovery policy. For information on using IP address notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

TIP! To target the entire monitored network, use `0.0.0.0/0` and `::/0`.

5. If you want to exclude the network from monitoring, select **Exclude**.
6. To add additional networks, repeat steps 4 and 5.
7. Target hosts that have a specific host attribute by clicking **Add** next to **Targeted Host Attributes**.

Attribute	Value
Location	New York

8. From the **Attribute** and **Value** drop-down lists, specify the host attribute.
9. To add additional host attributes, repeat steps 7 and 8.

A host must have at least one of the host attributes you specify to be evaluated against the white list.

10. Target hosts that belong to a specific VLAN by clicking **Add** next to **Targeted VLANs**.

VLAN ID
<input type="text"/>

11. In the **VLAN ID** field, specify the VLAN IDs of the hosts you want to evaluate against the white list. This can be any integer between 0 and 4095 for 802.1q VLANs.
12. To add additional VLAN IDs, repeat steps 10 and 11.

The host must be a member of one of the VLANs you specify to be evaluated against the white list.

TIP! To remove a network, host attribute restriction, or VLAN restriction, click the delete icon () next to the element you want to delete.

Modifying Existing Targets

LICENSE: FireSIGHT

After you modify a target, you must save the white list for your changes to take effect. Note that if you modify a target in a white list that is used by an active correlation policy, after you save the white list, any new target hosts are evaluated for compliance. However, this evaluation does not generate white list events. In addition, the system changes the white list host attribute of previously valid targets to **Not Evaluated**.

To modify an existing target:

ACCESS: Admin

1. On the Create White List page, under **Targets**, click the target you want to modify.

The settings for the target appear.

2. Make changes as needed.

You can rename the target, add or exclude additional networks, and add host attribute or VLAN restrictions. For more information, see [Configuring Compliance White List Targets](#) on page 1545.

Deleting Existing Targets

LICENSE: FireSIGHT

After you delete a target, you must save the white list for your changes to take effect. Note that if you delete a target from a white list that is used by an active correlation policy, the system changes the white list host attribute of previously valid targets to **Not Evaluated**.

To delete a white list target:

ACCESS: Admin

1. Next to the target you want to delete, click the delete icon ().
2. When prompted, confirm that you want to delete the target.

The target is deleted.

Configuring Compliance White List Host Profiles

LICENSE: FireSIGHT

Host profiles in a compliance white list specify which operating systems, clients, application protocols, web applications, and protocols are allowed to run on the

target hosts. There are three kinds of host profiles you can configure in a white list:

- global host profiles, which specify the application protocols, clients, web applications, and protocols that are allowed to run on target hosts, regardless of the host's operating system
- host profiles for specific operating systems, which specify not only which operating systems are allowed to run on your network, but also the application protocols, clients, web applications, and protocols that are allowed to run on those operating systems
- shared host profiles, which function exactly like the host profiles for specific operating systems, except they are not tied to a single white list; you can use them across multiple white lists

For a more detailed introduction to compliance white list host profiles, see [Understanding White List Host Profiles](#) on page 1534.

When you are finished creating compliance white list host profiles, you can add the white list to an active correlation policy to begin evaluating the target hosts for compliance, generating white list events when a host violated the white list, and optionally, triggering responses based on white list violations.

For information on how to create, modify, and delete compliance white list host profiles, see:

- [Configuring the Global Host Profile](#) on page 1549
- [Creating Host Profiles for Specific Operating Systems](#) on page 1550
- [Adding a Shared Host Profile to a Compliance White List](#) on page 1558
- [Modifying Existing Host Profiles](#) on page 1559
- [Deleting Existing Host Profiles](#) on page 1562

Configuring the Global Host Profile

LICENSE: FireSIGHT

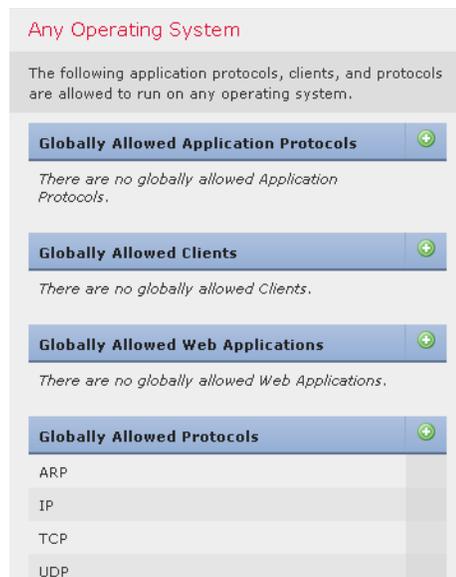
Every white list contains a global host profile, which specifies the application protocols, clients, web applications, and protocols that are allowed to run on target hosts, regardless of the host's operating system. For a more detailed introduction to the global host profile, see [Understanding the Global Host Profile](#) on page 1535.

To configure the global host profile:

ACCESS: Admin

1. On the Create White List page, under **Allowed Host Profiles**, click **Any Operating System**.

The settings for the global host profile appear.



2. To specify the application protocols you want to allow, follow the directions in [Adding an Application Protocol to a Host Profile](#) on page 1552.
3. To specify the clients you want to allow, follow the directions in [Adding a Client to a Host Profile](#) on page 1554.
4. To specify the web applications you want to allow, follow the directions in [Adding a Web Application to a Host Profile](#) on page 1556.
5. To specify the protocols you want to allow, follow the directions in [Adding a Protocol to a Host Profile](#) on page 1557.

Note that ARP, IP, TCP, and UDP are always allowed.

Creating Host Profiles for Specific Operating Systems

LICENSE: FireSIGHT

Host profiles for specific operating systems indicate not only which operating systems are allowed to run on your network, but also the application protocols, clients, web applications, and protocols that are allowed to run on those operating systems. For a more detailed introduction, see [Understanding Host Profiles for Specific Operating Systems](#) on page 1535.

To create a new compliance white list host profile for a specific operating system:

ACCESS: Admin

1. Next to **Allowed Host Profiles**, click the add icon (+).

The settings for the new host profile appear.

Host Profile: New Host Profile

Name: New Host Profile

OS Vendor: 12Planet

OS Name: Any

Version: Any

Allowed Application Protocols (+)

Allow all Application Protocols

Allowed Clients (+)

Allow all Clients

Allowed Web Applications (+)

Allow all Web Applications

Allowed Protocols (+)

ARP

IP

TCP

UDP

2. In the **Name** field, type a descriptive name for the host profile.
3. From the **OS Vendor**, **OS Name**, and **Version** drop-down lists, pick the operating system and version for which you want to create a host profile.
4. Specify the application protocols you want to allow. You have three options:
 - To allow all application protocols, leave the **Allow all Application Protocols** check box selected.
 - To allow no application protocols, clear the **Allow all Application Protocols** check box.
 - To allow specific application protocols, follow the directions in [Adding an Application Protocol to a Host Profile](#) on page 1552.
5. Specify the clients you want to allow. You have three options:
 - To allow all clients, leave the **Allow all Clients** check box selected.
 - To allow no clients, clear the **Allow all Clients** check box.
 - To allow specific clients, follow the directions in [Adding a Client to a Host Profile](#) on page 1554.

6. Specify the web applications you want to allow. You have three options:
 - To allow all web applications, leave the **Allow all Web Applications** check box selected.
 - To allow no web applications, clear the **Allow all Web Applications** check box.
 - To allow specific web applications, follow the directions in [Adding a Web Application to a Host Profile](#) on page 1556.

7. Specify the protocols you want to allow.

To add a protocol, next to **Allowed Protocols**, follow the directions in [Adding a Protocol to a Host Profile](#) on page 1557. Note that ARP, IP, TCP, and UDP are always allowed.

Adding an Application Protocol to a Host Profile

LICENSE: FireSIGHT

You can configure a compliance white list, using either a shared host profile or a host profile that belongs to a single white list, to allow certain application protocols to run on specific operating systems. You can also configure a white list to allow certain application protocols to run on any valid target; these are called globally allowed application protocols.

For any allowed application protocol, you can either specify the type of application protocol that you want to allow — FTP and SSH are examples of application protocol types — or you can allow a custom application protocol by specifying an application protocol type of **any**. You must also specify the protocol the allowed application protocol uses (TCP or UDP). You can allow the application protocol on any port, or restrict it to a port that you specify.

Optionally, you can require that the application protocol server have a specific vendor or version. For example, you could allow SSH to run on Linux hosts on port 22. You could also restrict the particular vendor and version to OpenSSH 4.2.

To add an application protocol to a compliance white list host profile:

ACCESS: Admin

1. While you are creating or modifying a white list host profile, click the add icon (+) next to **Allowed Application Protocols** (or next to **Globally Allowed Application Protocols** if you are modifying the Any Operating System host profile).

A pop-up window appears. The application protocols listed are:

- application protocols that you created within the white list
- application protocols that existed in the network map when you surveyed your networks as described in [Surveying Your Network](#) on page 1543
- application protocols that are used by other host profiles in the white list, which may include built-in application protocols created by the VRT for use in the default white list



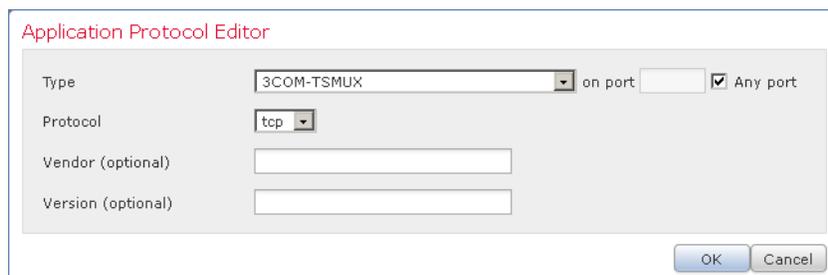
2. You have two options:

- To add an application protocol already in the list, select it and click **OK**. Use Ctrl or Shift while clicking to select multiple application protocols. You can also click and drag to select multiple adjacent application protocols.

The application protocol is added. Note that if you added a built-in application protocol, its name appears in italics. You can skip the rest of the procedure, or optionally, to change any of the application protocol's values (such as the port or protocol), click the application protocol you just added to display the application protocol editor.

- To add a new application protocol, select **<New Application Protocol>** and click **OK**.

The application protocol editor appears.



3. From the **Type** drop-down list, select the application protocol type. For custom application protocols, select **any**.
4. Specify the application protocol port. You have two options:
 - To allow the application protocol to run on any port, check the **Any port** check box.
 - To allow the application protocol to run only on a specific port, type the port number in the **port** field.
5. From the **Protocol** drop-down list, select the protocol: **TCP** or **UDP**.
6. Optionally, in the **Vendor** and **Version** fields, specify a vendor and version for the application protocol.

If you do not specify a vendor or version, the white list allows all vendors and versions as long as the type and protocol match. Note that if you restrict the vendor and version, you must make sure to specify them exactly as they would appear in an event view or in the application protocols network map.
7. Click **OK**.

The application protocol is added. Note that you must save the white list for your changes to take effect.

If you added an application protocol to a white list that is used by an active correlation policy, after you save the white list, the target hosts are re-evaluated. Although this re-evaluation may bring some hosts into compliance, it does not generate any white list events.

Adding a Client to a Host Profile

LICENSE: FireSIGHT

You can configure a compliance white list, using either a shared host profile or a host profile that belongs to a single white list, to allow certain client applications to run on specific operating systems. You can also configure a white list to allow certain clients to run on any valid target; these are called globally allowed clients.

Optionally, you can require that the client be a specific version. For example, you could allow only Microsoft Internet Explorer 8.0 to run on Microsoft Windows hosts.

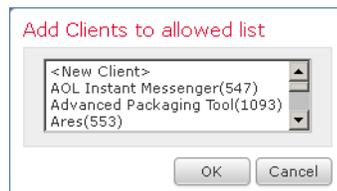
To add a client to a compliance white list host profile:

ACCESS: Admin

1. While you are creating or modifying a white list host profile, click the add icon (+) next to **Allowed Clients** (or next to **Globally Allowed Clients** if you are modifying the Any Operating System host profile).

A pop-up window appears. The clients listed are:

- clients that you created within the white list
- clients that were running on hosts in the network map when you surveyed your networks as described in [Surveying Your Network](#) on page 1543
- clients that are used by other host profiles in the white list, which may include built-in clients created by the VRT for use in the default white list



2. You have two options:
 - To add a client already in the list, select it and click **OK**. Use Ctrl or Shift while clicking to select multiple clients. You can also click and drag to select multiple adjacent clients.

The client is added. Note that if you added a built-in client, its name appears in italics. You can skip the rest of the procedure, or optionally, to change any of the client's values (such as its version), click the client you just added to display the client editor.

- To add a new client, select **<New Client>** and click **OK**.

The client editor appears.



3. From the **Client** drop-down list, select the client.
4. Optionally, in the **Version** field, specify a version for the client.

If you do not specify a version, the white list allows all versions as long as the name matches. Note that if you restrict the version, you must specify it exactly as it would appear in a table view of clients.

5. Click **OK**.

The client is added. Note that you must save the white list for your changes to take effect.

If you added a client to a white list that is used by an active correlation policy, after you save the white list, the target hosts are re-evaluated. Although this re-evaluation may bring some hosts into compliance, it does not generate any white list events.

Adding a Web Application to a Host Profile

LICENSE: FireSIGHT

You can configure a compliance white list, using either a shared host profile or a host profile that belongs to a single white list, to allow certain web applications to run on specific operating systems. You can also configure a white list to allow certain web applications to run on any valid target; these are called globally allowed web applications.

To add a web application to a compliance white list host profile:

ACCESS: Admin

1. While you are creating or modifying a white list host profile, click the add icon (+) next to **Allowed Web Applications** (or next to **Globally Allowed Web Applications** if you are modifying the Any Operating System host profile).

A pop-up window appears, listing all web applications detected by the system.

Add Web Applications to allowed list



2. Select a web application and click **OK**. Use Ctrl or Shift while clicking to select multiple web applications. You can also click and drag to select multiple adjacent web applications.

The web application is added. Note that you must save the white list for your changes to take effect.

If you added a web application to a white list that is used by an active correlation policy, after you save the white list, the target hosts are re-evaluated. Although this re-evaluation may bring some hosts into compliance, it does not generate any white list events.

Adding a Protocol to a Host Profile

LICENSE: FireSIGHT

You can configure a compliance white list, using either a shared host profile or a host profile that belongs to a single white list, to allow certain protocols to run on specific operating systems. You can also configure a white list to allow certain protocols to run on any valid target; these are called globally allowed protocols. Note that ARP, IP, TCP, and UDP are always allowed to run on any host; you cannot disallow them.

For any allowed protocol, you must specify its type (Network or Transport) and number.

To add a protocol to a compliance white list host profile:

ACCESS: Admin

1. While you are creating or modifying a white list host profile, click the add icon (+) next to **Allowed Protocols** (or next to **Globally Allowed Protocols** if you are modifying the Any Operating System host profile).

A pop-up window appears. The protocols listed are:

- protocols that you created within the white list
- protocols that were running on hosts in the network map when you surveyed your networks as described in [Surveying Your Network](#) on page 1543
- protocols that are used by other host profiles in the white list, which may include built-in protocols created by the VRT for use in the default white list



2. You have two options:

- To add a protocol already in the list, select it and click **OK**. Use Ctrl or Shift while clicking to select multiple protocols. You can also click and drag to select multiple adjacent protocols.

The protocol is added. Note that if you added a built-in protocol, its name appears in italics. You can skip the rest of the procedure, or optionally, to change any of the protocol's values (such as the type or number) click the protocol you just added to display the protocol editor.

- To add a new protocol, select **<New Protocol>** and click **OK**.

The protocol editor appears.



3. From the **Type** drop-down list, select the protocol type: **Network** or **Transport**.

4. Specify the protocol. You have two options:

- Select a protocol from the drop-down list.
- Select **Other (manual entry)** to specify a protocol that is not in the list. For network protocols, type the appropriate number as listed in <http://www.iana.org/assignments/ethernet-numbers/>. For transport protocols, type the appropriate number as listed in <http://www.iana.org/assignments/protocol-numbers/>.

5. Click **OK**.

The protocol is added. Note that you must save the white list for your changes to take effect.

If you added a protocol to a white list that is used by an active correlation policy, after you save the white list, the target hosts are re-evaluated.

Although this re-evaluation may bring some hosts into compliance, it does not generate any white list events.

Adding a Shared Host Profile to a Compliance White List

LICENSE: FireSIGHT

Shared host profiles are also tied to specific operating systems, but you can use them across white lists. That is, if you create multiple white lists but want to use the same host profile to evaluate hosts running a particular operating system across the white lists, use a shared host profile.

You can add any of the built-in shared host profiles to your compliance white lists, or you can add shared host profiles that you created. For more information, see [Understanding Shared Host Profiles](#) on page 1537 and [Creating Shared Host Profiles](#) on page 1565.

To add a shared host profile to a compliance white list:

ACCESS: Admin

1. On the Create White List page, click **Add Shared Host Profile**.

The Add Shared Host Profile page appears.



2. From the **Name** drop-down list, select the shared host profile you want to add to your white list, and click **OK**.

The shared host profile is added to your white list and the Create White List page appears again. The shared host profile's name appears in italics under Allowed Host Profiles.

TIP! You can edit a shared host profile from within a white list that uses it by clicking on the profile name under Allowed Host Profiles. For more information, see [Modifying Existing Host Profiles](#) on page 1559.

Modifying Existing Host Profiles

LICENSE: FireSIGHT

After you modify a host profile within a compliance white list, you must save the white list for your changes to take effect.

If a host profile you modify belongs to a white list used in an active correlation policy, modifying the profile may bring hosts into or out of compliance but does **not** generate white list events. Further, modifying a shared host profile affects every white list that uses it. This may bring hosts into or out of compliance not only in the white list you are working with, but in other white lists as well.

TIP! As with other shared host profiles, you can edit the built-in host profiles used by the default white list. You can also reset them to their factory defaults. For more information, see [Resetting Built-In Host Profiles to Their Factory Defaults](#) on page 1571.

To modify an existing host profile:

ACCESS: Admin

1. On the Create White List page, click the name of the host profile you want to modify.

The settings for the host profile appear. Note that if you are editing a shared host profile, an **Edit** link appears next to the name of the host profile. If you are editing a built-in host profile, the built-in host profile icon (🔒) also appears.

Host Profile: Debian Linux 2.2

Name	<input type="text" value="Debian Linux 2.2"/>
OS Vendor	<input type="text" value="Debian"/>
OS Name	<input type="text" value="Linux"/>
Version	<input type="text" value="2.2"/>

Allowed Application Protocols	
Telnet/7008 TCP	

Allowed Clients	
No Clients are allowed	
<input type="checkbox"/> Allow all Clients	

Allowed Web Applications	
No Web Applications are allowed	
<input type="checkbox"/> Allow all Web Applications	

Allowed Protocols	
ARP	
IP	
TCP	
UDP	

2. You have two options:
 - If you are modifying a shared host profile, click **Edit**.
A pop-up window appears. Make changes as needed as described in the table below. Click **Save All Profiles** to save the profile, then click **Done** to close the pop-up window.
For more information editing shared host profiles, see [Modifying a Shared Host Profile](#) on page 1567.
 - If you are modifying either the white list's global host profile or a host profile for a specific operating system, perform one of the actions described in the procedures below.

To rename the host profile:

ACCESS: Admin

- ▶ Type a new name in the **Name** field.

To change the operating system for the host profile:

ACCESS: Admin

- ▶ Select the new operating system and version from the **OS Vendor**, **OS Name**, and **Version** drop-down lists.

If you change these values, you may also want to rename the host profile. Note that a white list's global host profile has no operating system associated with it, so you cannot change it.

To add an application protocol:

ACCESS: Admin

- ▶ Follow the directions in [Adding an Application Protocol to a Host Profile](#) on page 1552.

To add a client:

ACCESS: Admin

- ▶ Follow the directions in [Adding a Client to a Host Profile](#) on page 1554.

To add a web application:

ACCESS: Admin

- ▶ Follow the directions in [Adding a Web Application to a Host Profile](#) on page 1556.

To add a protocol:

ACCESS: Admin

- ▶ Follow the directions in [Adding a Protocol to a Host Profile](#) on page 1557.

To allow all application protocols:

ACCESS: Admin

- ▶ Under **Allowed Application Protocols**, select the **Allow all Application Protocols** check box.

Note that the check box does not appear until you delete any application protocols you have previously allowed.

To allow all clients:

ACCESS: Admin

- ▶ Under **Allowed Clients**, select the **Allow all Clients** check box.

Note that the check box does not appear until you delete any clients you have previously allowed.

To allow all web applications:

ACCESS: Admin

- ▶ Under **Allowed Web Applications**, select the **Allow all Web Applications** check box.

Note that the check box does not appear until you delete any web applications you have previously allowed.

To modify an application protocol, client, web application, or protocol:

ACCESS: Admin

- ▶ Click the element you want to modify.

For more information on the properties you can change, see:

- [Adding an Application Protocol to a Host Profile](#) on page 1552
- [Adding a Client to a Host Profile](#) on page 1554
- [Adding a Protocol to a Host Profile](#) on page 1557

IMPORTANT! The changes you make to an application protocol, client, web application, or protocol are reflected in every host profile that uses that element.

To delete an application protocol, client, web application, or protocol:

ACCESS: Admin

- ▶ Next to the element you want to delete, click the delete icon ().

To survey your network:

ACCESS: Admin

- ▶ Click **Survey Network**. Surveying your network can add additional allowed clients, application protocols, and protocols to the host profiles that already exist, and can create additional host profiles if the survey detects hosts running operating systems that were not detected during any initial survey. For more information, see [Surveying Your Network](#) on page 1543.

Deleting Existing Host Profiles

LICENSE: FireSIGHT

After you delete a host profile from a compliance white list, you must save the white list for your changes to take effect. Note that deleting a shared host profile removes it from the white list, but does not delete the profile or remove it from any other white lists that use it. You cannot delete a white list's global host profile.

If the host profile you delete belongs to one or more white lists used in an active correlation policy, deleting the profile may force hosts out of compliance, but does **not** generate white list events.

To delete a compliance white list host profile:

ACCESS: Admin

1. On the Create White List page, next to the host profile you want to delete, click the delete icon ().
2. When prompted, confirm that you want to delete the host profile.
The host profile is deleted.

Managing Compliance White Lists

LICENSE: FireSIGHT

Use the White List page to manage compliance white lists. You can create, modify, and delete white lists, including the default white list. You can also edit any shared host profiles you have created, as well as the built-in shared host profiles, and add new shared host profiles.



Name	
Default White List Recommended White List	 
Microsoft Windows Only allows only Microsoft Windows operating systems	 
My White List	 

For more information, see:

- [Creating Compliance White Lists](#) on page 1541
- [Modifying a Compliance White List](#) on page 1563
- [Deleting a Compliance White List](#) on page 1564
- [Working with Shared Host Profiles](#) on page 1564

Modifying a Compliance White List

LICENSE: FireSIGHT

When you modify a compliance white list that is included in an active correlation policy, the system re-evaluates the target hosts. Note that the system does **not** generate white list events — and therefore does not trigger any responses you associated with the white list — during this re-evaluation, even if the white list is included in an active correlation policy and a previously compliant host becomes non-compliant with the updated white list.

To modify an existing compliance white list:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.
The White List page appears.

2. Next to the white list you want to modify, click the edit icon (✎).
The Create White List page appears.
3. Make modifications as necessary and click **Save White List**.
The white list is updated.

Deleting a Compliance White List

LICENSE: FireSIGHT

You cannot delete a compliance white list that you are using in one or more correlation policies; you must first delete the white list from all policies where it is used. For information on deleting a white list from a policy, see [Editing a Correlation Policy](#) on page 1520.

Deleting a white list also removes the host attribute associated with the white list from all hosts on your network.

To delete an existing compliance white list:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.
The White List page appears.
2. Next to the white list you want to delete, click the delete icon (🗑️).
The white list is deleted.

Working with Shared Host Profiles

LICENSE: FireSIGHT

Shared host profiles specify which operating systems, clients, application protocols, web applications, and protocols are allowed to run on target hosts across multiple white lists. That is, if you create multiple white lists but want to use the same host profile to evaluate hosts running a particular operating system across the white lists, use a shared host profile. Note that the default white list uses a special category of shared host profiles, called *built-in host profiles*.

For a more detailed introduction to shared host profiles, see [Understanding Shared Host Profiles](#) on page 1537.

You can create, modify, and delete shared host profiles. In addition, if you modify or delete any of the built-in shared host profiles, or modify or delete any of the built-in application protocols, protocols, or clients, you can reset them to their factory defaults. For more information, see:

- [Creating Shared Host Profiles](#) on page 1565
- [Modifying a Shared Host Profile](#) on page 1567

- [Deleting a Shared Host Profile](#) on page 1571
- [Resetting Built-In Host Profiles to Their Factory Defaults](#) on page 1571

After you create a shared host profile, you can add it to multiple white lists. For more information, see [Adding a Shared Host Profile to a Compliance White List](#) on page 1558.

Creating Shared Host Profiles

LICENSE: FireSIGHT

Create a shared host profile if you want to use the same host profile to evaluate hosts running a particular operating system across multiple white lists.

TIP! You can also create a shared host profile for your compliance white lists using the host profile of a specific host. For more information, see [Creating a White List Host Profile from a Host Profile](#) on page 1361.

To create a shared host profile:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.

The White List page appears.

2. Click **Edit Shared Profiles**.

The Edit Shared Profiles page appears.



3. Optionally, survey your network.

Surveying your network creates several baseline shared white lists based on the data the system has accumulated about your network. This saves you from manually creating and configuring multiple shared host profiles. You have two options:

- To survey your network, click **Survey Network**. For more information, see [Surveying Your Network](#) on page 1543.

The system creates one or more baseline shared host profiles. You can edit or delete these shared host profiles as described in [Modifying a Shared Host Profile](#) on page 1567 and [Deleting a Shared Host Profile](#) on page 1571. To add any other shared host profiles you might need, continue with the next step.

- To skip surveying your network, continue with the next step.

- Next to **Shared Host Profiles**, click the add icon (+).
- The settings for the new shared host profile appear.

Host Profile: New Host Profile

Name: New Host Profile

OS Vendor: 12Planet

OS Name: Any

Version: Any

Allowed Application Protocols (+)

No Application Protocols are allowed

Allow all Application Protocols

Allowed Clients (+)

No Clients are allowed

Allow all Clients

Allowed Web Applications (+)

No Web Applications are allowed

Allow all Web Applications

Allowed Protocols (+)

ARP

IP

TCP

UDP

- In the **Name** field, type a descriptive name for the shared host profile.
- From the **OS Vendor**, **OS Name**, and **Version** drop-down lists, pick the operating system and version for which you want to create a shared host profile.
- Specify the application protocols you want to allow. You have three options:
 - To allow all application protocols, select the **Allow all Application Protocols** check box.
 - To allow no application protocols, leave the **Allow all Application Protocols** check box cleared.
 - To allow specific application protocols, next to **Allowed Application Protocols**, follow the directions in [Adding an Application Protocol to a Host Profile](#) on page 1552.
- Specify the clients you want to allow. You have three options:
 - To allow all clients, select the **Allow all Clients** check box.
 - To allow no clients, leave the **Allow all Clients** check box cleared.
 - To allow specific clients, follow the directions in [Adding a Client to a Host Profile](#) on page 1554.

9. Specify the web applications you want to allow. You have three options:
 - To allow all web applications, select the **Allow all Web Applications** check box.
 - To allow no web applications, leave the **Allow all Web Applications** check box cleared.
 - To allow specific web applications, follow the directions in [Adding a Web Application to a Host Profile](#) on page 1556.
10. Specify the protocols you want to allow.

To add a protocol, next to **Allowed Protocols**, follow the directions in [Adding a Protocol to a Host Profile](#) on page 1557. Note that ARP, IP, TCP, and UDP are always allowed.
11. Click **Save all Profiles** to save your changes.

The shared host profile is created. You can now add the shared host profile to any compliance white list.

Modifying a Shared Host Profile

LICENSE: FireSIGHT

Modifying a shared host profile changes the profile for all the white lists it belongs to. For the white lists that use the shared host profile and are also used in an active correlation policy, modifying a shared host profile may bring hosts into or out of compliance, but does **not** generate white list events.

To modify a shared host profile:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.

The White List page appears.

2. Click **Edit Shared Profiles**.

The Edit Shared Profiles page appears.



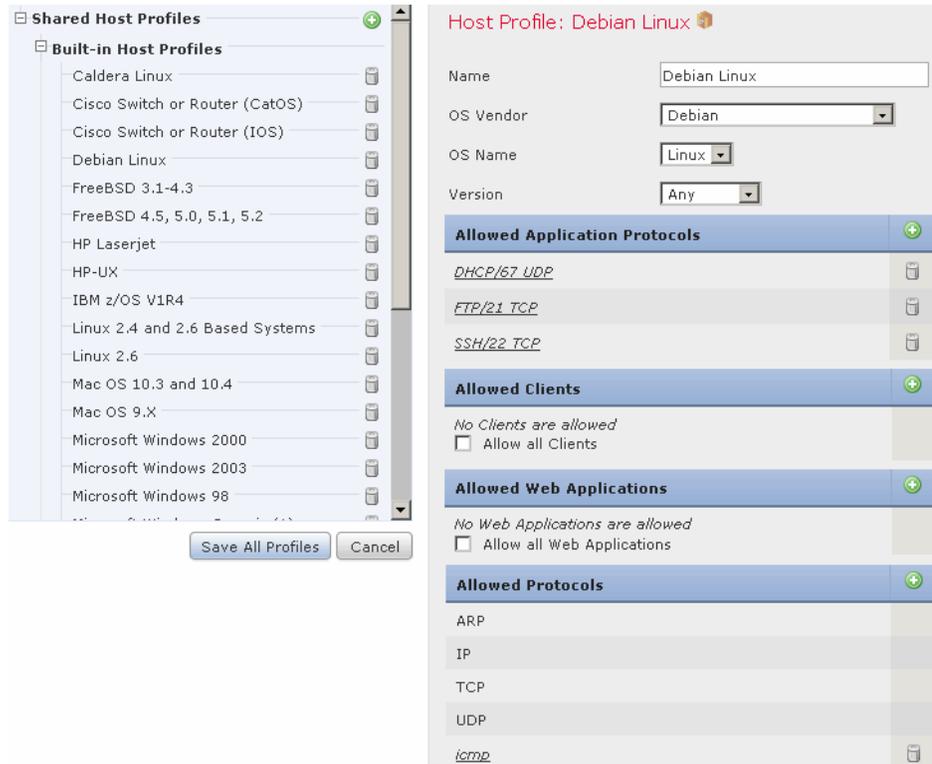
3. Do you want to edit one of the built-in shared host profiles?

- If yes, expand **Built-in Host Profiles** to display those host profiles.



- If no, continue with the next step.

- Click the name of the shared host profile you want to modify.
The host profile appears.



- Perform any of the actions described in the table below.

To...	You CAN...
rename the host profile	type a new name in the Name field.
change the operating system	select the new operating system and version from the OS Vendor , OS Name , and Version drop-down lists. If you change these values, you may also want to rename the host profile.
add an application protocol	follow the directions in Adding an Application Protocol to a Host Profile on page 1552.
add a client	follow the directions in Adding a Client to a Host Profile on page 1554.
add a web application	follow the directions in Adding a Web Application to a Host Profile on page 1556.

To...	You CAN...
add a protocol	follow the directions in Adding a Protocol to a Host Profile on page 1557.
allow all application protocols	under Allowed Application Protocols , select the Allow all Application Protocols check box. Note that the check box does not appear until you delete any application protocols you have previously allowed.
allow all clients	under Allowed Clients , select the Allow all Clients check box. Note that the check box does not appear until you delete any clients you have previously allowed.
allow all web applications	under Allowed Web Applications , select the Allow all Web Applications check box. Note that the check box does not appear until you delete any clients you have previously allowed.
modify an application protocol, client, web application, or protocol	click the element you want to modify. For more information on the properties you can change, see: <ul style="list-style-type: none">• Adding an Application Protocol to a Host Profile on page 1552• Adding a Client to a Host Profile on page 1554• Adding a Web Application to a Host Profile on page 1556• Adding a Protocol to a Host Profile on page 1557 <p>IMPORTANT! The changes you make to an application protocol, client, or protocol are reflected in every host profile that uses that element.</p>
delete an application protocol, client, web application, or protocol	next to the element you want to delete, click the delete icon ().
survey your network	click Survey Network . Surveying your network can add additional allowed clients, application protocols, web applications, and protocols to the host profiles that already exist, and can create additional host profiles if the survey detects hosts running operating systems that were not detected during any initial survey. For more information, see Surveying Your Network on page 1543.

6. Click **Save all Profiles** to save your changes.
The shared host profile is saved.

Deleting a Shared Host Profile

LICENSE: FireSIGHT

If the shared host profile you delete belongs to one or more white lists used in an active correlation policy, deleting the profile may force hosts out of compliance, but does **not** generate white list events.

TIP! If you delete a built-in shared host profile that is used by the default white list, you can restore it by resetting the built-in profiles to their factory defaults. For more information, see [Resetting Built-In Host Profiles to Their Factory Defaults](#) on page 1571.

To delete a shared host profile:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.
The White List page appears.
2. Click **Edit Shared Profiles**.
The Edit Shared Profiles page appears.
3. Do you want to delete one of the built-in shared host profiles?
 - If yes, expand **Built-in Host Profiles** to display those host profiles.
 - If no, continue with the next step.
4. Next to the shared host profile you want to delete, click the delete icon ().
Confirm that you want to delete the shared host profile.
5. Click **Save all Profiles** to save your changes.
The shared host profile is deleted and removed from all compliance white lists that use it.

Resetting Built-In Host Profiles to Their Factory Defaults

LICENSE: FireSIGHT

The default white list uses a special category of shared host profiles, called *built-in host profiles*. Built-in host profiles use built-in application protocols, protocols, and clients. You can use any these elements as-is in both the default white list and in any custom white list that you create, or you can modify them to suit your needs. For more information, see [Understanding Shared Host Profiles](#).

If you make changes to the built-in profiles, application protocols, protocols, web applications, or clients that you need to undo, you can reset to factory defaults. When you reset to factory defaults, the following things occur:

- **All** built-in host profiles, application protocols, protocols, and clients that you modified are reset to their factory defaults.
- **All** built-in host profiles, application protocols, protocols, and clients that you deleted are restored.
- **All** white lists (including the default white list) that are used by active correlation policies and that used any of the reset built-in host profiles, application protocols, protocols, or clients are re-evaluated. Although this re-evaluation may change the compliance of some hosts into compliance, it does not generate any white list events.

To reset built-in host profiles, application protocols, protocols, and clients:

ACCESS: Admin

1. Select **Policies > Correlation**, then click **White List**.

The White List page appears.

2. Click **Edit Shared Profiles**.

The Edit Shared Profiles page appears.

3. Click **Built-in Host Profiles**.

The Built-in Host Profiles page appears.



4. Click **Reset to Factory Defaults**.

5. Confirm that you want to reset to factory defaults by clicking **OK**.

All built-in host profiles, application protocols, protocols, and clients are reset to factory defaults. Any white list that is used by an active correlation policy and that used any of the reset built-in host profiles, application protocols, protocols, or clients is re-evaluated.

Working with White List Events

LICENSE: FireSIGHT

When the system generates a discovery event that indicates that a host is out of compliance with a white list that is included in an activated correlation policy, a white list event is generated. White list events are a special kind of correlation

event, and are logged to the correlation event database. You can search, view, and delete white list events.

TIP! For information on configuring the number of events saved in the database, see [Configuring Database Event Limits](#) on page 1978. Note that white list events are stored in the correlation event database.

For more information, see the following sections:

- [Viewing White List Events](#) on page 1573
- [Understanding the White List Events Table](#) on page 1575
- [Searching for Compliance White List Events](#) on page 1576

Viewing White List Events

LICENSE: FireSIGHT

You can use the Defense Center to view a table of compliance white list events. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access white list events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of white list events. You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The following table describes some of the specific actions you can perform on an white list events workflow page.

Compliance White List Event Actions

To...	YOU CAN...
view the host profile for a host	click the host profile icon () that appears next to the IP address.
view user profile information	click the user icon () that appears next to the user identity. For more information, see Understanding User Details and Host History on page 1449.
sort and constrain events on the current workflow page	find more information in Sorting Drill-Down Workflow Pages on page 1835.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1836.

Compliance White List Event Actions (Continued)

To...	YOU CAN...
navigate between pages in the current workflow, keeping the current constraints	click the appropriate page link at the top left of the workflow page. For more information, see Using Workflow Pages on page 1812.
learn more about the columns that appear	find more information in Understanding the White List Events Table on page 1575.
modify the time and date range for displayed events	find more information in see Setting Event Time Constraints on page 1820.
drill down to the next page in the workflow, constraining on a specific value	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • on a drill-down page that you created in a custom workflow, click a value within a row. Note that clicking a value within a row in a table view constrains the table view and does not drill down to the next page. • To drill down to the next workflow page constraining on some users, select the check boxes next to the users you want to view on the next workflow page, then click View. • To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include “Table View” in the page name. For more information, see Constraining Events on page 1830.</p>
delete white list events from the system	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To delete some events, select the check boxes next to the events you want to delete, then click Delete. • To delete all events in the current constrained view, click Delete All, then confirm you want to delete all the events.
navigate to other event views to view associated events	find more information in Navigating Between Workflows on page 1836.

To view compliance white list events:

ACCESS: Admin/Any Security Analyst/Discovery Admin

- ▶ Select **Analysis > Correlation > White List Events**.

The first page of the default white list events workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

Understanding the White List Events Table

LICENSE: FireSIGHT

You can use the correlation policy feature to build *correlation policies* that let the system respond in real time to threats on your network. Correlation policies describe the type of activity that constitutes a policy violation, which include compliance white list violations. For more information on correlation policies, see [Configuring Correlation Policies and Rules](#) on page 1459.

When a compliance white list is violated, the system generates a white list event. The fields in the white list events table are described in the following table.

Compliance White List Event Fields

FIELD	DESCRIPTION
Time	The date and time that the white list event was generated.
IP Address	The IP address of the non-compliant host.
User	The identity of any known user logged in to the non-compliant host.
Port	The port, if any, associated with the event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of white list violations, this field is blank.
Description	A description of how the white list was violated. For example: Client "AOL Instant Messenger" is not allowed. Violations that involve an application protocol indicate the application protocol name and version, as well as the port and protocol (TCP or UDP) it is using. If you restrict prohibitions to a particular operating system, the description includes the operating system name. For example: Server "ssh / 22 TCP (OpenSSH 3.6.1p2)" is not allowed on Operating System "Linux Linux 2.4 or 2.6".
Policy	The name of the correlation policy that was violated, that is, the correlation policy that includes the white list.
White List	The name of the white list.
Priority	The priority specified by the policy or white list that triggered the policy violation. For information on setting correlation rule and policy priorities, see Providing Basic Policy Information on page 1515 and Setting Rule and White List Priorities on page 1516.

Compliance White List Event Fields (Continued)

FIELD	DESCRIPTION
Host Criticality	The user-assigned host criticality of the host that is out of compliance with the white list: None , Low , Medium , or High . For more information on host criticality, see Working with the Predefined Host Attributes on page 1369.
Device	The name of the managed device that detected the white list violation.
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for Compliance White List Events

LICENSE: FireSIGHT

You can search for specific compliance white list events. You may want to create searches customized for your network environment, then save them to re-use later. The following table describes the search criteria you can use.

Compliance White List Event Search Criteria

FIELD	SEARCH CRITERIA RULES
Policy	Enter the name of a correlation policy to return all events caused by violations of white lists included in that policy.
White List	Enter the name of a white list to return all events caused by violations of that white list.
Description	Enter the white list event description.
Priority	Specify the priority of the white list event, which is determined either by the priority of the white list in a correlation policy or by the priority of the correlation policy itself. Note that the white list priority overrides the priority of its policy. Enter none for no priority. For information on setting correlation rule and policy priorities, see Providing Basic Policy Information on page 1515 and Setting Rule and White List Priorities on page 1516.
IP Address	Specify an IP address of a host that has become non-compliant with a white list.

Compliance White List Event Search Criteria (Continued)

FIELD	SEARCH CRITERIA RULES
User	Specify the identity of the user logged in to a host that has become non-compliant with a white list.
Port	Specify the port, if any, associated with the discovery event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol).
Host Criticality	Specify the host criticality of the source host involved in the white list event: None , Low , Medium , or High . For more information on host criticality, see Working with the Predefined Host Attributes on page 1369.
Device	Type the name of the device or device group that detected the white list violation.

To search for compliance white list events:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **White List Events**.
The page reloads with the appropriate constraints.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.

If you do not enter a name, one is created automatically when you save the search.

4. Enter your search criteria in the appropriate fields, as described in the [Compliance White List Event Search Criteria table](#) on page 1576, and keeping in mind the following additional points:
 - All fields accept negation (!).
 - All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
 - Many fields accept one or more asterisks (*) as wild cards.
 - Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
 - Click the add object icon (🟢+) that appears next to a search field to use an object as a search criterion.

For more information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.

Your search results appear in the default white list events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click (**switch workflow**) by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

Working with White List Violations

LICENSE: FireSIGHT

The system keeps track of the ways in which hosts on your network violate the compliance white lists in active correlation policies. You can search and view these records.

For more information, see the following sections:

- [Viewing White List Violations](#) on page 1579
- [Understanding the White List Violations Table](#) on page 1581
- [Searching for White List Violations](#) on page 1582

Viewing White List Violations

LICENSE: FireSIGHT

You can use the Defense Center to view a table of white list violations. Then, you can manipulate the event view depending on the information you are looking for. The page you see when you access white list violations differs depending on the workflow you use. There are two predefined workflows:

- The Host Violation Count workflow provides a series of pages that list all the hosts that violate at least one white list. The first page sorts the hosts based on the number of violations per host, with the hosts with the greatest number of violations at the top of the list. If a host violates more than one white list, there is a separate row for each violated white list. The workflow also contains a table view of white list violations that lists all violations with the most recently detected violation at the top of the list. Each row in the table contains a single detected violation.
- The White List Violations workflow includes a table view of white list violations that lists all violations with the most recently detected violation at the top of the list. Each row in the table contains a single detected violation.

Both predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs. For more information, see [Creating Custom Workflows](#) on page 1841.

The following table describes some of the specific actions you can perform on a white list violations workflow page.

Compliance White List Violations Actions

To...	YOU CAN...
view the host profile for a host	click the host profile icon () that appears next to the IP address.
sort and constrain events on the current workflow page	find more information in Sorting Drill-Down Workflow Pages on page 1835.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1836.
navigate between pages in the current workflow, keeping the current constraints	click the appropriate page link at the top left of the workflow page. For more information, see Using Workflow Pages on page 1812.
learn more about the columns that appear	find more information in Understanding the White List Violations Table on page 1581.
drill down to the next page in the workflow	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To drill down to the next workflow page constraining on a specific value, click a value within a row. Note that this only works on drill-down pages. Clicking a value within a row in a table view constrains the table view and does not drill down to the next page. • To drill down to the next workflow page constraining on some events, select the check boxes next to the events you want to view on the next workflow page, then click View. • To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include “Table View” in the page name.</p> <p>For more information, see Constraining Events on page 1830.</p>
navigate to other event views to view associated events	find more information in Navigating Between Workflows on page 1836.

To view compliance white list violations:

ACCESS: Admin/Any Security Analyst/Discovery Admin

- ▶ Select **Analysis > Correlation > White List Violations**.

The first page of the default white list violations workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

Understanding the White List Violations Table

LICENSE: FireSIGHT

You can use the correlation policy feature to build *correlation policies* that let the system respond in real time to threats on your network. Correlation policies describe the type of activity that constitutes a policy violation, which include compliance white list violations. For more information on correlation policies, see [Configuring Correlation Policies and Rules](#) on page 1459.

When a compliance white list is violated, the system records the violation. Note that you can not set event time constraints in the table view because the table view displays only the current host violations on your network. The fields in the white list violations table are described in the [Compliance White List Violation Fields](#) table.

Compliance White List Violation Fields

FIELD	DESCRIPTION
Time	The date and time that the white list violation was detected.
IP Address	The relevant IP address of the non-compliant host.
Type	The type of white list violation, that is, whether the violation occurred as a result of a non-compliant: <ul style="list-style-type: none">• operating system (os)• application protocol (server)• client (client)• protocol (protocol)• web application (web)
Information	Any available vendor, product, or version information associated with the white list violation. For example, if you have a white list that allows only Microsoft Windows hosts, the Information field describes the operating systems of the hosts that are not running Microsoft Windows. For protocols that violate a white list, the Information field also indicates whether the violation is due to a network or transport protocol.
Port	The port, if any, associated with the event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of white list violations, this field is blank.
Protocol	The protocol, if any, associated with the event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol). For other types of white list violations, this field is blank.

Compliance White List Violation Fields (Continued)

FIELD	DESCRIPTION
White List	The name of the white list that was violated.
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Searching for White List Violations

LICENSE: FireSIGHT

You can search for specific compliance white list violations. You may want to create searches customized for your network environment, then save them to re-use later. The following table describes the search criteria you can use.

Compliance White List Violations Search Criteria

FIELD	SEARCH CRITERIA RULES
Time	Specify the date and time that the white list was violated.
IP Address	Specify an IP address of a host that has become non-compliant with a white list.
White List	Enter the name of a white list to return all violations from that white list.
Type	Enter the type of white list violation: <ul style="list-style-type: none"> • enter os (or operating system) to search for violations based on operating systems • enter server to search for violations based on application protocols • enter client to search for violations based on clients • enter protocol to search for violations based on protocols • enter web application to search for violations based on web applications
Information	Enter white list violation information.

Compliance White List Violations Search Criteria (Continued)

FIELD	SEARCH CRITERIA RULES
Port	Specify the port, if any, associated with the discovery event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol).
Protocol	Specify the protocol, if any, associated with the discovery event that triggered an application protocol white list violation (a violation that occurred as a result of a non-compliant application protocol).

To search for compliance white list violations:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **White List Violations**.
The page reloads with the appropriate constraints.

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: White List Violations

Name: Search 1, My Search

Save As Private:

Constraint

Time: > 2009-07-16 13:00:31, < today at 4:30pm

IP Address: 192.168.1.0/24, 192.168.1.3, 2001:0db8:85a3::1370

White List: My White List

Type: server, client, operating system, protocol

Information: http *, Microsoft Windows XP

Port: 1-1024, 6000-6011, 180

Protocol: tcp, udp

Buttons: Search, Save As New Search

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.

4. Enter your search criteria in the appropriate fields, as described in the [Compliance White List Event Search Criteria table](#) on page 1576, and keeping in mind the following additional points:
 - All fields accept negation (!).
 - All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
 - Many fields accept one or more asterisks (*) as wild cards.
 - Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
 - Click the add object icon (🟢+) that appears next to a search field to use an object as a search criterion.

For more information on search syntax, including using objects in searches, see [Searching for Events](#) on page 1770.

5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default white list violations workflow. To use a different workflow, including a custom workflow, click (**switch workflow**) by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

CHAPTER 38

CREATING TRAFFIC PROFILES

A traffic profile is just that—a profile of the traffic on your network, based on connection data collected over a time span that you specify. You can use connection data collected by your devices, the connection data exported by any or all of your NetFlow-enabled devices, or both.

After you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

Keep in mind that the Sourcefire 3D System uses connection data to create traffic profiles and trigger correlation rules based on traffic profile changes. You cannot include connections that you do not log to the Defense Center database in traffic profiles. The system uses only end-of-connection data to populate connection summaries (see [Understanding Connection Summaries](#) on page 542), which the system then uses to create connection graphs and traffic profiles. Therefore, if you want to create and use traffic profiles, make sure you log connection events at the end of connections.

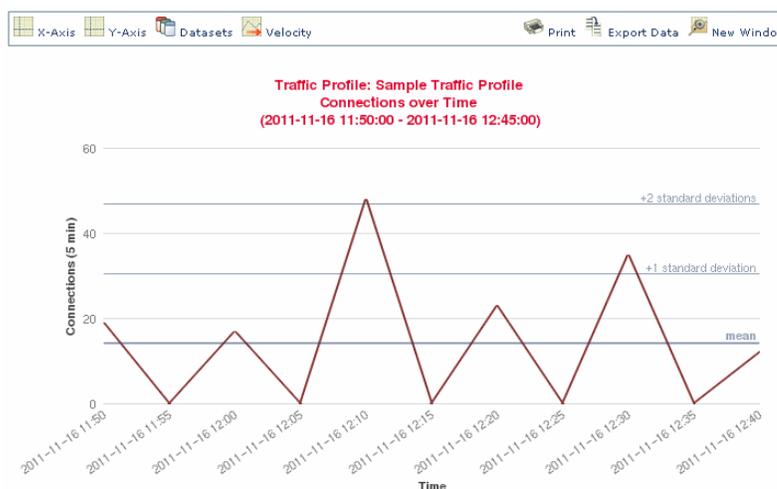
The time span used to collect data to build your traffic profile is called the profiling time window (PTW). The PTW is a sliding window; that is, if your PTW is one week (the default), your traffic profile includes connection data collected over the last week. You can change the PTW to be as short as an hour or as long as several weeks.

When you first activate a traffic profile, it collects and evaluates connection data according to the criteria you have set, for a learning period equal in time to the PTW. The Defense Center does not evaluate rules you have written against the traffic profile until the learning period is complete.

You can create profiles using all the traffic on a monitored network segment, or you can create more targeted profiles using criteria based on the data in the connection events. For example, you could set the profile conditions so that the traffic profile only collects data where the detected session uses a specific port, protocol, or application. Or, you could add a host profile qualification to the traffic profile to collect data only for hosts that exhibit a host criticality of **high**.

Finally, when you create a traffic profile, you can specify inactive periods—periods in which connection data do not affect profile statistics and rules written against the profile do not trigger. You can also change how often the traffic profile aggregates and calculates statistics on collected connection data.

The following graphic shows a traffic profile with a PTW of one day and a sampling rate of five minutes.



After you create and activate a traffic profile and its learning period is complete, you can create correlation rules that trigger when you detect anomalous traffic. For example, you could write a rule that triggers if the amount of data traversing your network (measured in packets, KBytes, or number of connections) suddenly spikes to three standard deviations above the mean amount of traffic, which could indicate an attack or other security policy violation. Then, you could include that rule in a correlation policy to alert you of the traffic spike or to perform a remediation in response. For information on using traffic profiles to detect abnormal network traffic, see [Creating Rules for Correlation Policies](#) on page 1461.

You create traffic profiles on the Traffic Profiles page.

Name	Progress	PTW	Sort by
Sample Traffic Profile one day - IPv4 private networks	<div style="width: 100%;"></div> 100 %	(1 day)	State
HTTP traffic profile	<div style="width: 0%;"></div> 0 %	(1 hour)	

The slider icon next to each profile indicates whether the profile is active. If you want to base a correlation rule on a traffic profile change, you must activate the profile. If the slider icon is blue with a check mark, the profile is active. If it is gray with an x, the profile is inactive. For more information, see [Activating and Deactivating Traffic Profiles](#) on page 1595.

The progress bar shows the status of the traffic profile's learning period. When the progress bar reaches 100%, correlation rules written against the profile will trigger.

TIP! You can sort traffic profiles by state (active versus inactive) or alphabetically by name using the **Sort by** drop-down list.

For more information, see:

- [Providing Basic Profile Information](#) on page 1587
- [Specifying Traffic Profile Conditions](#) on page 1588
- [Adding a Host Profile Qualification](#) on page 1590
- [Setting Profile Options](#) on page 1593
- [Saving a Traffic Profile](#) on page 1595
- [Activating and Deactivating Traffic Profiles](#) on page 1595
- [Editing a Traffic Profile](#) on page 1596
- [Understanding Condition-Building Mechanics](#) on page 1596

Providing Basic Profile Information

LICENSE: FireSIGHT

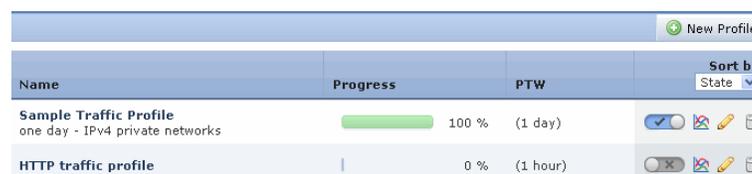
When you create a traffic profile, you must give it a name and, optionally, a short description.

To begin creating a traffic profile:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then click **Traffic Profiles**.

The Traffic Profiles page appears.



Name	Progress	PTW	Sort by
Sample Traffic Profile one day - IPv4 private networks	<div style="width: 100%;"></div> 100 %	(1 day)	State
HTTP traffic profile	<div style="width: 0%;"></div> 0 %	(1 hour)	

2. Click **New Profile**.

The Create Profile page appears.

The screenshot shows the 'Create Profile' page with three main sections: 'Profile Information', 'Profile Conditions', and 'Profile Options'.
- **Profile Information:** Includes 'Profile Name' and 'Profile Description' text boxes. A '+ Add Host Profile Qualification' button is in the top right.
- **Profile Conditions:** Features a blue header 'Collect connection information for all traffic that matches the following conditions:'. Below it are '+ Add condition' and '+ Add complex condition' buttons. A dropdown menu with a red 'X' icon is visible.
- **Profile Options:** Includes 'Profiling Time Window' (Maintain data for this profile for the last 1 week(s)), 'Sampling Rate' (Sample data every 05 minutes), and 'Inactive Periods' (There are no defined inactive periods. To add an inactive period, click "Add Inactive Period"). A '+ Add Inactive Period' button is in the top right.
At the bottom right are 'Save', 'Save & Activate', and 'Cancel' buttons.

3. In the **Profile Name** field, type a name of up to 255 characters for the new traffic profile.
4. In the **Profile Description** field, type a short description of up to 255 characters of the new traffic profile.
5. Continue with [Specifying Traffic Profile Conditions](#).

Specifying Traffic Profile Conditions

LICENSE: FireSIGHT

Profile conditions constrain the kinds of connection data you want the traffic profile to track. A simple traffic profile that profiles all the traffic on a monitored network segment has no conditions. In contrast, traffic profiles can be complex, with multiple nested conditions.

For example, the traffic profile conditions in the following graphic collects HTTP connections on the 10.4.x.x subnet.

The screenshot shows the 'Profile Conditions' section with a blue header 'Collect connection information for all traffic that matches the following conditions:'. Below it are '+ Add condition' and '+ Add complex condition' buttons. A dropdown menu with 'AND' is visible. Two conditions are listed:
- Application Protocol is HTTP
- Initiator/Responder IP is in 10.4.0.0/16
A 'Copy Settings' button is in the top right.

You build traffic profile conditions in the **Profile Conditions** section of the Create Profile page. See [Understanding Condition-Building Mechanics](#) on page 1596 for information on building conditions. Also, the syntax you can use to build

conditions is fully described in [Syntax for Traffic Profile Conditions](#) on page 1589.

TIP! If you want to use the settings from an existing traffic profile, click **Copy Settings** and, in the pop-up window, select the traffic profile you want to use and click **Load**.

Syntax for Traffic Profile Conditions

LICENSE: FireSIGHT

The [Syntax for Profile Conditions](#) table describes how to build a traffic profile condition.

Keep in mind that NetFlow records do not contain information about which host in the connection is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

Syntax for Profile Conditions

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Initiator IP, Responder IP, or Initiator/Responder IP	Use a specific IP address or CIDR notation to specify a range of IP addresses. See Specifying IP Addresses in Searches on page 1776 for a description of the syntax allowed for IP addresses. Note, however, that you cannot use the Local or remote keywords to specify IP addresses that are or are not in the networks you are monitoring.
Responder Port/ ICMP Code	Type the port number or ICMP code.
Transport Protocol	Type TCP or UDP as the transport protocol.
Connection Type	Specify in the traffic profile whether you want to use connection data collected by your Sourcefire devices or by NetFlow-enabled devices. If you do not specify a connection type, the traffic profile includes both.
NetFlow Device	Select the NetFlow-enabled device whose data you want to use to create the traffic profile. If you did not add any NetFlow-enabled devices to your deployment (using the local configuration), the NetFlow Device drop-down list is blank.

Syntax for Profile Conditions (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Application Protocol	Select an application protocol name from the drop-down list of available protocols.
Client	Select a client name from the drop-down list of available clients.
Web Application	Select a web application name from the drop-down list of available web applications.

Adding a Host Profile Qualification

LICENSE: FireSIGHT

You can constrain any traffic profile with information from the host profile of the tracked hosts. This constraint is called a *host profile qualification*. For example, as shown in the following graphic, you could collect connection data only for hosts that are assigned a host criticality of **high**.



To use a host profile qualification, the host must exist in the database and the host profile property you want to use as a qualification must already be included in the host profile. For example, if you configure a correlation policy rule to trigger when an intrusion event is generated on a host running Windows, the rule only triggers if the host is already identified as Windows when the intrusion event is generated.

To add a host profile qualification:

ACCESS: Admin/Discovery Admin

1. On the Create Profile page, click **Add Host Profile Qualification**.

The Host Profile Qualification section appears.



2. Build the host profile qualification's conditions.

You can create a single, simple condition, or you can create more elaborate constructs by combining and nesting conditions. See [Understanding Condition-Building Mechanics](#) on page 1596 for information building conditions.

The syntax you can use to build conditions is described in [Syntax for Host Profile Qualifications](#) on page 1591.

TIP! To remove a host profile qualification, click **Remove Host Profile Qualification**.

Syntax for Host Profile Qualifications

LICENSE: FireSIGHT

When you build a host profile qualification condition, you must first select the host you want to use to constrain your traffic profile. You can select either **Responder Host** or **Initiator Host**. After you select the host role, continue building your host profile qualification condition, as described in the [Syntax for Host Profile Qualifications](#) table.

Although you can configure the network discovery policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. In addition, if your traffic profile uses connection data exported by NetFlow-enabled devices, keep in mind that NetFlow records do not contain information about which host in the connection is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known. For more information, see [Differences Between NetFlow and FireSIGHT Data](#) on page 1272.

To match against *implied* or generic clients, create a host profile qualification based on the application protocol used by the server responding to the client. When the client list on a host that acts as the initiator or source of a connection includes an application protocol name followed by **client**, that client may actually be an implied client. In other words, the system reports that client based on server response traffic that uses the application protocol for that client, not on detected client traffic.

For example, if the system reports **HTTPS client** as a client on a host, create a host profile qualification for **Responder Host** where **Application Protocol** is set to **HTTPS**,

because HTTPS client is reported as a generic client based on the HTTPS server response traffic sent by the responder or destination host.

Syntax for Host Profile Qualifications

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
Host Type	Select one or more host types from the drop-down list. You can choose between a normal host or one of several types of network device.
NETBIOS Name	Type the NetBIOS name of the host.
Operating System > OS Vendor	Select one or more operating system vendor names from the drop-down list.
Operating System > OS Name	Select one or more operating system names from the drop-down list.
Operating System > OS Version	Select one or more operating system versions from the drop-down list.
Network Protocol	Type the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers .
Transport Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers .
Host Criticality	Select the host criticality from the list that appears. You can select None , Low , Medium , or High . For more information on host criticality, see Working with the Predefined Host Attributes on page 1369.
VLAN ID	Type the VLAN ID number of the host.
Application Protocol > Application Protocol	Select an application protocol from the drop-down list.
Application Protocol > Application Port	Type the application protocol port number.
Application Protocol > Protocol	Select the protocol from the drop-down list.
Client > Client	Select a client from the drop-down list.
Client > Client Version	Type the client version.
Web Application	Select a client from the drop-down list.

Syntax for Host Profile Qualifications (Continued)

IF YOU SPECIFY...	SELECT AN OPERATOR, THEN...
MAC Address > MAC Address	Type all or part of the MAC address of the host.
MAC Address > MAC Type	Select whether the MAC type is ARP/DHCP Detected . That is, select whether the system positively identified the MAC address as belonging to the host (is ARP/DHCP Detected), whether the system is seeing many hosts with that MAC address because, for example, there is a router between the device and the host (is not ARP/DHCP Detected), or whether the MAC type is irrelevant (is any).
MAC Vendor	Type all or part of the MAC vendor of hardware used by the host.
any available host attribute, including the default compliance white list host attribute	Specify the appropriate value, which depends on the type of host attribute you select: <ul style="list-style-type: none"> • If the host attribute type is Integer, enter an integer value in the range defined for the attribute. • If the host attribute type is Text, and enter a text value. • If the host attribute type is List, select a valid list string from the drop-down list. • If the host attribute type is URL, enter a URL value. <p>For more information on host attributes, see Working with User-Defined Host Attributes on page 1370.</p>

Setting Profile Options

LICENSE: FireSIGHT

The profiling time window (PTW) is the sliding time window, equal in length to the learning period, that the Sourcefire 3D System uses to calculate statistics for the traffic profile. The default PTW is one week, but you can change it to be as short as an hour or as long as several weeks.

Also, traffic profiles are based on aggregated connection data. By default, traffic profiles generate statistics on connection events generated by the system over five-minute intervals. However, you can set this sampling rate anywhere between the default five minutes and one hour.

Keep in mind that you should set your PTW and sampling rate so that your traffic profiles contain enough data to be statistically meaningful. For example, a PTW of one day with a sampling rate of one hour would only contain 24 data points, which may not be enough for accurate analysis of network traffic patterns.

TIP! Your PTW should include at least 100 data points.

You can also set up inactive periods in traffic profile. For example, consider a network infrastructure where all the workstations are backed up at midnight every night. The backup takes about 30 minutes and spikes the network traffic. In that case, you might want to set up a recurring inactive period for your traffic profile to coincide with the scheduled backups. During inactive periods, the traffic profile collects data (so you can see the traffic on the traffic profile graphs), but that data is not used when calculating profile statistics. You can set up inactive periods to recur daily, weekly, or monthly. Inactive periods can be as short as five minutes or as long as one hour. Traffic profile graphs plotted over time show inactive periods as a shaded region.



ACCESS: Admin/Discovery Admin

Follow the directions in the [Profile Options](#) table to set profile options.

Profile Options

To...	YOU CAN...
change the profiling time window	in the Profiling Time Window field, type the number of hours, days, or weeks. Then choose hour(s) , day(s) , or week(s) from the drop-down list.
change the sampling rate	select the rate from the Sampling Rate drop-down list.
add an inactive period	click Add Inactive Period . Then, using the drop-down lists, specify when and how often you want the traffic profile to refrain from collecting data.
delete an inactive period	click Delete next to the inactive period you want to delete.

Saving a Traffic Profile

LICENSE: FireSIGHT

Use the following procedure to save a traffic profile.

To save a traffic profile:

ACCESS: Admin/Discovery Admin

- ▶ You have two options:
 - To save the profile without activating it, click **Save**.
 - To save the profile and start collecting data immediately, click **Save & Activate**.

Activating and Deactivating Traffic Profiles

LICENSE: FireSIGHT

When you want it to begin profiling the traffic on a monitored network segment, you must activate a traffic profile.

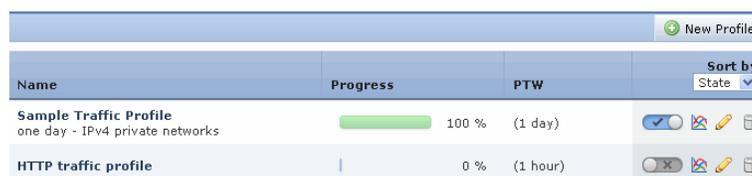
Deactivate a profile when you want it to stop collecting and evaluating connection data. Rules written against deactivated traffic profiles do not trigger. In addition, deactivating a traffic profile deletes all data collected and aggregated by the profile. If you later reactivate a deactivated traffic profile, you must wait the length of its PTW before rules written against it will trigger.

To activate or deactivate a traffic profile:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then click **Traffic Profiles**.

The Traffic Profiles page appears.



Name	Progress	PTW	Sort by
Sample Traffic Profile one day - IPv4 private networks	<div style="width: 100%;"></div> 100 %	(1 day)	State
HTTP traffic profile	<div style="width: 0%;"></div> 0 %	(1 hour)	

2. You have two options:
 - To activate an inactive traffic profile, click **Activate** next to the profile.
 - To deactivate an active traffic profile, click **Deactivate** next to the profile. Confirm that you want to deactivate the profile by clicking **OK**.

Editing a Traffic Profile

LICENSE: FireSIGHT

You cannot substantially edit an active traffic profile; if the traffic profile is active you can only change its name and description. To edit a traffic profile's conditions options, you must first deactivate it. Note that deactivating a traffic profile deletes all the data it has collected.

For more information on activating and deactivating traffic profiles, see [Activating and Deactivating Traffic Profiles](#) on page 1595.

To edit a traffic profile:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then click **Traffic Profiles**.
The Traffic Profiles page appears.
2. Next to the traffic profile you want to edit, click **Edit**.
The Create Profile page appears.
3. Make changes to the profile and click **Save**.
Your profile is updated.

Understanding Condition-Building Mechanics

LICENSE: FireSIGHT

You build traffic profiles by specifying the conditions they use to collect data. You can create either simple conditions or more elaborate constructs with nested conditions.

For example, if you want to create a traffic profile that collects data for your entire monitored network segment, you can create a very simple profile with no conditions, as shown in the following graphic.

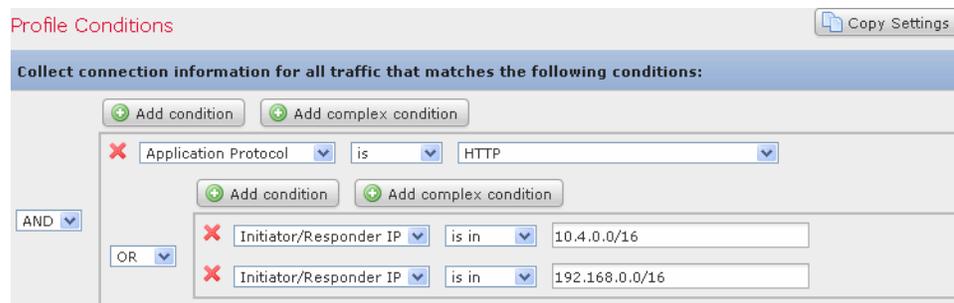
The screenshot displays the configuration interface for a traffic profile. It is divided into two main sections: "Profile Information" and "Profile Conditions".

- Profile Information:** This section contains two text input fields. The "Profile Name" field is filled with "Simple Traffic Profile". The "Profile Description" field is filled with "Collects all connection data on the". To the right of these fields is a button labeled "Add Host Profile Qualification".
- Profile Conditions:** This section features a "Copy Settings" button. Below it is a blue header bar with the text "Collect connection information for all traffic that matches the following conditions:". Underneath this bar are two buttons: "Add condition" and "Add complex condition". At the bottom of the section, there is a red 'X' icon followed by a dropdown menu.

If you wanted to constrain the profile and collect data only for the 10.4.x.x network, you can add a single condition, as shown in the following graphic.



But the following traffic profile, which collects HTTP activity on the 10.4.x.x network and the 192.168.x.x network, has three conditions, with the last constituting a complex condition.



The syntax you can use within conditions varies depending on the element you are creating, but the mechanics are the same. For more information, see:

- [Building a Single Condition](#) on page 1597
- [Adding and Linking Conditions](#) on page 1599
- [Using Multiple Values in a Condition](#) on page 1602

Building a Single Condition

LICENSE: FireSIGHT

Most conditions have three parts: a category, an operator, and a value. Some conditions are more complex and contain several categories, each of which may have their own operators and values.

For example, the following traffic profile collects information on the 10.4.x.x network. The category of the condition is **Initiator/Responder IP**, the operator is **is in**, and the value is **10.4.0.0/16**.



The following steps explain how to build this traffic profile condition.

To build a single condition:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then click **Traffic Profiles**.
The Traffic Profiles page appears.
2. Click **New Profile**.
The Create Profile page appears.
3. Under **Profile Conditions**, begin building the profile's single condition by selecting **Initiator/Responder IP** from the first (category) drop-down list.
4. Select **is in** from the second (operator) drop-down list.

TIP! When the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a range of IP addresses, as expressed in CIDR notation. For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

5. Type **10.4.0.0/16** in the text field.

In contrast, the following host profile qualification is more complex; it constrains a traffic profile such that it collects connection data only if the responding host in the detected connection is running a version of Microsoft Windows.

Responder Host	Operating System	has the following properties
<input type="text"/>	OS Vendor	is <input type="text" value="Microsoft"/>
<input type="text"/>	OS Name	is <input type="text" value="Windows"/>
<input type="text"/>	OS Version	is <input type="text" value="any"/>

The following steps explain how to build this host profile qualification.

To build this host profile qualification:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then click **Traffic Profiles**.
The Traffic Profiles page appears.
2. Click **New Profile**.
The Create Profile page appears.
3. Click **Add Host Profile Qualification**.

4. Under **Host Profile Qualification**, in the first condition, specify the host whose information you want to collect.
In this example, select **Responder Host** because we only want information on responding hosts in a connection.
5. Begin specifying the details of the operating system of the host by choosing the **Operating System** category.
Three subcategories appear: **OS Vendor**, **OS Name**, and **OS Version**.
6. To specify that the host can be running any version of Microsoft Windows, use the same operator for all three subcategories: **is**.
7. Finally, specify the values for the subcategories.
Select **Microsoft** as the value for **OS Vendor**, **Windows** as the value for **OS Name**, and leave **any** as the value for **OS Version**.

Note that the categories you can choose from depend on whether you are building traffic profile conditions or a host profile qualification. In addition, a condition's available operators depend on the category you choose. Finally, the syntax you can use to specify a condition's value depends on the category and operator. Sometimes you must type the value in a text field. Other times, you can pick a value from a drop-down list.

IMPORTANT! Where the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list. For more information, see [Using Multiple Values in a Condition](#) on page 1602.

For information on the syntax for building traffic profile conditions and host profile qualifications, see:

- [Syntax for Traffic Profile Conditions](#) on page 1589
- [Syntax for Host Profile Qualifications](#) on page 1591

Adding and Linking Conditions

LICENSE: FireSIGHT

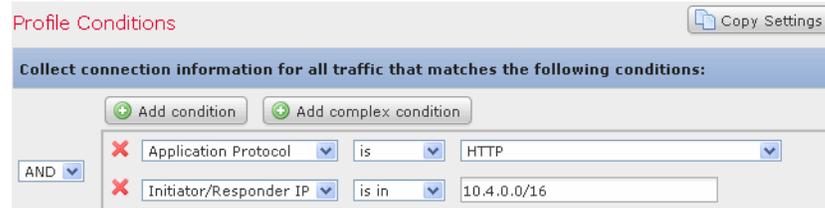
You can create simple traffic profile conditions and host profile qualifications, or you can create more elaborate constructs by combining and nesting conditions.

When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together:

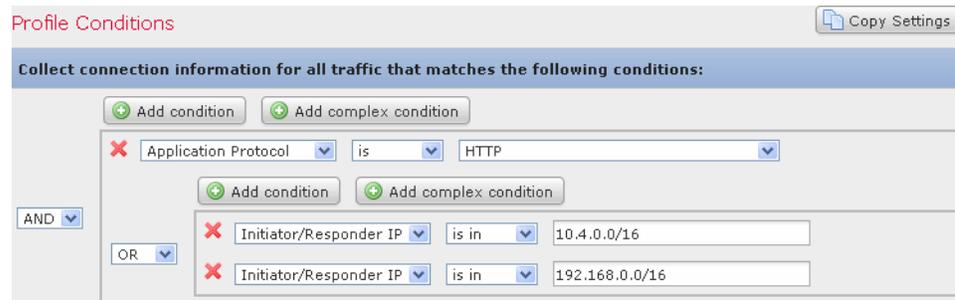
- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

For example, the following traffic profile contains two conditions linked by **AND**. This means that the traffic profile collects connection data only if both conditions

are true. In this example, it collects HTTP connections for all hosts with IP addresses in the 10.4.x.x subnet.



In contrast, the following traffic profile, which collects connection data for HTTP activity in either the 10.4.x.x network or the 192.168.x.x network, has three conditions, with the last constituting a complex condition.



Logically, the above traffic profile is evaluated as follows:
 (A and (B or C))

WHERE...	IS THE CONDITION THAT STATES...
A	Application Protocol Name is HTTP
B	IP Address is in 10.4.0.0/16
C	IP Address is in 192.168.0.0/16

To add a single condition:

ACCESS: Admin/Discovery Admin

- ▶ To add a single condition, click **Add condition** above the current condition. A new condition is added to the same logical level as the current set of conditions. By default, it is linked to the conditions on its level with the **OR** operator, though you can change the operator to **AND**.

For example, if you add a simple condition to the following host profile qualification:

Host Profile Qualification Remove Host Profile Qualification

Only collect connection information with the following properties:

Add condition Add complex condition

× Responder Host Host Criticality is High

The result is:

Host Profile Qualification Remove Host Profile Qualification

Only collect connection information with the following properties:

Add condition Add complex condition

OR × Responder Host Host Criticality is High

× [Empty Condition]

To add a complex condition:

ACCESS: Admin/Discovery Admin

- ▶ Click **Add complex condition** above the current condition.

A complex condition is added below the current set of conditions. The complex condition comprises two subconditions, which are linked to each other with the opposite operator from the one used to link the conditions on the level above it.

For example, if you add a complex condition to the following host profile qualification:

The screenshot shows a 'Host Profile Qualification' window with a 'Remove Host Profile Qualification' button. Below the title bar, it says 'Only collect connection information with the following properties:'. There are two buttons: 'Add condition' and 'Add complex condition'. Below these is a single condition: 'Responder Host' (with a red 'X' icon), 'Host Criticality', 'is', and 'High'.

The result is:

The screenshot shows the 'Host Profile Qualification' window after adding a complex condition. The main condition is 'Responder Host is High'. Below it, there is a complex condition structure. On the left, there is an 'OR' dropdown. Inside the complex condition, there is an 'AND' dropdown, followed by two empty condition fields, each with a red 'X' icon. There are also 'Add condition' and 'Add complex condition' buttons above the subconditions.

To link conditions:

ACCESS: Admin/Discovery Admin

- ▶ Use the drop-down list to the left of a set of conditions:
 - To require that all conditions on the level that the operator controls are met, select **AND**.
 - To require that only one of the conditions on the level that the operator controls is met, select **OR**.

Using Multiple Values in a Condition

LICENSE: FireSIGHT

When you are building a condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list. For example, if you want to add a host profile qualification to a traffic profile that requires that a host be running some flavor of UNIX, instead of constructing multiple conditions linked with the OR operator, use the following procedure.

To include multiple values in one condition:

ACCESS: Admin/Discovery Admin

1. Build a condition, choosing **is in** or **is not in** as the operator. The drop-down list changes to a text field.

2. Click anywhere in the text field or on the **Edit** link. A pop-up window appears.

3. Under **Available**, use Ctrl or Shift while clicking to select multiple values. You can also click and drag to select multiple adjacent values.
4. Click the right arrow (>) to move the selected entries to **Selected**.
5. Click **OK**.

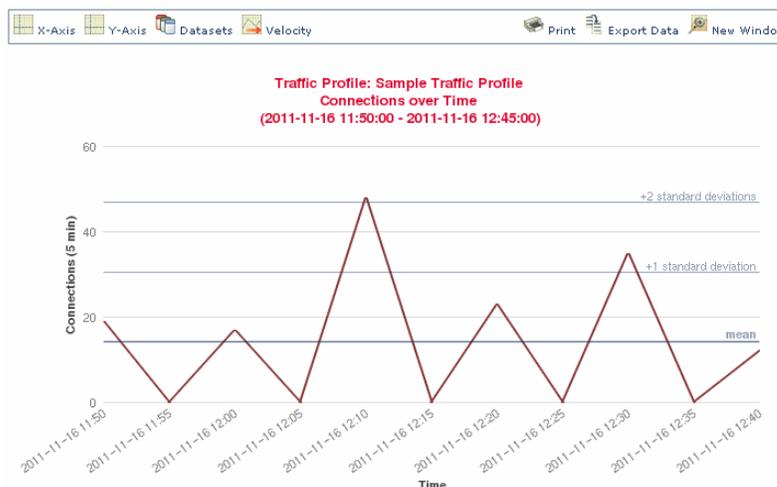
Your selections appear in the value field of your condition on the Create Profile page.

Viewing Traffic Profiles

LICENSE: FireSIGHT

Because traffic profiles are based on connection data, you can view graphs of traffic profiles. The following graphic shows a traffic profile with a PTW of one

week, a sampling rate of five minutes, and a daily half-hour inactive period from midnight to 12:30 AM.



You can perform almost all the same actions on traffic profile graphs that you can perform on connection data graphs. However, because traffic profiles are based on aggregated data (connection summaries), you cannot examine the individual connection events on which the graphs are based. In other words, you cannot drill down to a connection data table view from a traffic profile graph. See [Viewing Connection Data Graphs and Tables](#) on page 555 for more information. In addition, traffic profiles appear as detached graphs. For more information, see [Detaching Connection Graphs](#) on page 569.

In addition, traffic profile graphs plotted over time show the mean (average) y-axis value as a bold horizontal line. Graphs over time also plot the values of the first four standard deviations above and below the mean, assuming that network traffic is distributed normally. By default, these statistics are calculated over the PTW, but if you alter the time settings for a graph, the Defense Center recalculates the statistics. Rules written against traffic profile statistics, however, are always evaluated against the statistics for the PTW.

To view a traffic profile graph for a traffic profile:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Correlation**, then click **Traffic Profiles**.

The Traffic Profiles page appears.

				New Profile
Name	Progress	PTW	Sort by State	
Sample Traffic Profile one day - IPv4 private networks	<div style="width: 100%; height: 10px; background-color: green;"></div> 100 %	(1 day)	<input checked="" type="checkbox"/>	
HTTP traffic profile	<div style="width: 0%; height: 10px; background-color: gray;"></div> 0 %	(1 hour)	<input type="checkbox"/>	

2. Next to the traffic profile for which you want to view the graph, click the graph icon ().

The graph for the traffic profile appears in a separate browser window.

CHAPTER 39

CONFIGURING REMEDIATIONS

When a correlation policy violation occurs, you can configure the Sourcefire 3D System to initiate one or multiple responses, which include remediations (such as running an Nmap scan) and various types of alerts.

The most basic kind of response you can launch is an alert. Alerts notify you, via email, a SNMP trap server, or syslog, of a policy violation. For information on creating alerts, see [Configuring External Alerting](#) on page 524.

Another kind of response you can launch is a remediation. A remediation is a program that the Defense Center runs when your network traffic violates a correlation policy. The Sourcefire 3D System ships with predefined remediations, which perform actions such as blocking a host at the firewall or router when it violates a policy or scanning the host.

When the Defense Center launches a remediation, it generates a remediation status event. You can search, view, and delete remediation status events, as you would any other event.

The Sourcefire 3D System also provides a flexible API that allows you to create custom remediation modules to respond to correlation policy violations. For example, if you are running a Linux-based firewall, you could write and upload a remediation module that dynamically updates the `iptables` file on the Linux server so that traffic violating a correlation policy is blocked. For more information about writing your own remediation modules, refer to the *Sourcefire Remediation API Guide*.

For more information, see:

- [Creating Remediations](#) on page 1607
- [Working with Remediation Status Events](#) on page 1633

Creating Remediations

LICENSE: FireSIGHT

In addition to alerts, which are simple notifications of a correlation policy violation, you can also configure responses called *remediations*. Remediations are programs that the Defense Center runs when a correlation policy is violated. These programs use information provided in the event that triggered the violation to perform a specific action.

The Sourcefire 3D System ships with several predefined remediation modules:

- The Cisco IOS Null Route module, which, if you are running Cisco routers that use Cisco IOS® Version 12.0 or higher, allows you to dynamically block traffic sent to an IP address or network that violates a correlation policy.

See [Configuring Remediations for Cisco IOS Routers](#) on page 1608 for more information.

- The Cisco PIX Shun module, which, if you are running Cisco PIX® Firewall Version 6.0 or higher, allows you to dynamically block traffic sent from an IP address that violates a correlation policy.

See [Configuring Remediations for Cisco PIX Firewalls](#) on page 1617 for more information.

- The Nmap Scanning module, which allows you to actively scan specific targets to determine operating systems and servers running on those hosts.

See [Configuring Nmap Remediations](#) on page 1623 for more information.

- The Set Attribute Value module, which allows you to set a host attribute on a host where a correlation event occurs.

See [Configuring Set Attribute Remediations](#) on page 1629.

You can create multiple instances for each remediation module, where each instance represents a connection to a specific appliance. For example, if you have four Cisco IOS routers where you want to send remediations, you should configure four instances of the Cisco IOS remediation module.

When you create an instance, you specify the configuration information necessary for the Defense Center to establish a connection with the appliance. Then, for each configured instance, you add remediations that describe the actions you want the appliance to perform when a policy is violated.

After they are configured, you can add remediations to what are called response groups, or you can assign the remediations specifically to rules within correlation policies. When the system executes these remediations, it generates a remediation status event, which includes details such as the remediation name, the policy and rule that triggered it, and the exit status message. For more information on these events, see [Working with Remediation Status Events](#) on page 1633.

In addition to the default modules that Sourcefire provides, you can write custom remediation modules that perform other specific tasks when policy violations

trigger. Refer to the *Sourcefire Remediation API Guide* for more information about writing your own remediation modules and installing them on the Defense Center. If you are installing a custom module, you can use the Modules page to install, view, and delete new modules.

To install a new module on the Defense Center:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Modules**.

The Modules page appears.

Installed Remediation Modules

Module Name	Version	Description	
Cisco IOS Null Route	1.0	Block an IP address in a Cisco IOS router	 
Cisco PIX Shun	1.0	Shun an IP address in the PIX firewall	 
Nmap Remediation	2.0	Perform an Nmap Scan	 
Set Attribute Value	1.0	Set an Attribute Value	 



2. Click **Browse** to navigate to the location where you saved the file that contains the custom remediation module (refer to the *Sourcefire Remediation API Guide* for more information).

3. Click **Install**.

The custom remediation module installs.

To view or delete a module from the Defense Center:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Modules**.

The Modules page appears.

2. Perform one of the following actions:

- Click **View** to view the module.
The Module Detail page appears.
- Click **Delete** next to the module you want to delete. You **cannot** delete default modules provided by Sourcefire.

The remediation module is deleted.

Configuring Remediations for Cisco IOS Routers

LICENSE: FireSIGHT

Sourcefire provides a Cisco IOS Null Route remediation module that allows you to block a single IP address or an entire block of addresses using Cisco's "null route" command when a correlation policy is violated. This forwards all traffic sent to the

host or network listed as the source or destination host in the event that violated the correlation policy to the router's NULL interface, causing it to be dropped (note that this will not block traffic sent **from** the violating host or network).

The Cisco IOS Null Route remediation module supports Cisco routers running Cisco IOS 12.0 and higher. You must have level 15 administrative access to the router to execute Cisco IOS remediations.

IMPORTANT! A destination-based remediation only works if you configure it to launch when a correlation rule that is based on a connection event or intrusion event triggers. Discovery events only transmit source hosts.

WARNING! When a Cisco IOS remediation is activated, there is no timeout period. To remove the blocked IP address or network from the router, you must manually clear the routing change from the router itself.

To create remediations for routers running Cisco IOS:

ACCESS: Admin/Discovery Admin

1. Enable Telnet on the Cisco router.
Refer to the documentation provided with your Cisco router or IOS software for more information about enabling Telnet.
2. On the Defense Center, add a Cisco IOS Null Route instance for each Cisco IOS router you plan to use with the Defense Center.
See [Adding a Cisco IOS Instance](#) on page 1609 for the procedures.
3. Create specific remediations for each instance, based on the type of response you want to elicit on the router when correlation policies are violated.

Each available remediation type is described in the following sections:

- [Cisco IOS Block Destination Remediations](#) on page 1612
- [Cisco IOS Block Destination Network Remediations](#) on page 1614
- [Cisco IOS Block Source Remediations](#) on page 1615
- [Cisco IOS Block Source Network Remediations](#) on page 1616

4. Begin assigning Cisco IOS remediations to specific correlation policy rules.

Adding a Cisco IOS Instance

LICENSE: FireSIGHT

After you configure Telnet access on the Cisco IOS router (refer to the documentation provided with your Cisco router or IOS software for more information about enabling Telnet access), you can add an instance to the

Defense Center. If you have multiple routers where you want to send remediations, you must create a separate instance for each router.

To add a Cisco IOS instance:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.

The Instances page appears.

Configured Instances

Instance Name	Module Name	Version	
IOS_01 Cisco IOS Null Route	Cisco IOS Null Route	1.0	 
NmapTest	Nmap Remediation	2.0	 
PIX_01 Cisco PIX Shun	Cisco PIX Shun	1.0	 
Value_01 Set Attribute Value	Set Attribute Value	1.0	 

Add a New Instance

Select a module type:

2. From the **Add a New Instance** list, select **Cisco IOS Null Route (v1.0)** and click **Add**.

The Edit Instance page appears.

Edit Instance

Instance Name

Module

Description

Router IP

Username *(optional)*

Connection Password
Retype to confirm

Enable Password
Retype to confirm

White List
(an optional list of networks)

3. In the **Instance Name** field, enter a name for the instance.
The name you choose should contain no spaces or special characters and should be descriptive. For example, if you intend to connect more than one Cisco IOS router, you will have multiple instances, so you may want to choose a name such as `IOS_01` and `IOS_02`.
4. In the **Router IP** field, enter the IP address of the Cisco IOS router you want to use for the remediation.
5. In the **Username** field, enter the Telnet user name for the router. This user must have level 15 administrative access on the router.
6. In the **Connection Password** fields, enter the Telnet user's user password. The password entered in both fields must match.
7. In the **Enable Password** fields, enter the Telnet user's enable password. This is the password used to enter privileged mode on the router. The password entered in both fields must match.
8. In the **White List** field, enter IP addresses that you want to exempt from the remediation, one per line. You can also use CIDR notation or a specific IP address. For example, the following white list would be accepted by the system:

```
10.1.1.152
172.16.1.0/24
```

Note that this white list is not associated with any compliance white lists you have created. For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

9. Click **Create**.
The instance is created and remediations appear in the Configured Remediations section of the page. You must add specific remediations for them to be used by correlation policies. See the following sections for more information:
 - [Cisco IOS Block Destination Remediations](#) on page 1612
 - [Cisco IOS Block Destination Network Remediations](#) on page 1614
 - [Cisco IOS Block Source Remediations](#) on page 1615
 - [Cisco IOS Block Source Network Remediations](#) on page 1616

Cisco IOS Block Destination Remediations

LICENSE: FireSIGHT

The Cisco IOS Block Destination remediation allows you to block traffic sent from the router to the destination host in a correlation event.

IMPORTANT! Do **not** use this remediation as a response to a correlation rule that is based on a discovery event; discovery events only transmit a source host and not a destination host. You can use this remediation in response to correlation rules that are based on connection events or intrusion events.

To add the remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.

The Instances page appears.

Configured Instances

Instance Name	Module Name	Version	
IOS_01 Cisco IOS Null Route	Cisco IOS Null Route	1.0	 
NmapTest	Nmap Remediation	2.0	 
PIX_01 Cisco PIX Shun	Cisco PIX Shun	1.0	 
Value_01 Set Attribute Value	Set Attribute Value	1.0	 

Add a New Instance

Select a module type

2. Next to the instance where you want to add the remediation, click the view icon (🔍).

If you have not yet added an instance, see [Adding a Cisco IOS Instance](#) on page 1609.

The Edit Instance page appears.

Edit Instance

Instance Name

Module Cisco IOS Null Route (v1.0)

Description

Router IP

Username (optional)

Connection Password
Retype to confirm

Enable Password
Retype to confirm

White List
(an optional list of networks)

3. In the **Configured Remediations** section, select **Block Destination** and click **Add**. The Edit Remediation page appears.

Edit Remediation

Remediation Name

Remediation Type Block Destination

Description

4. In the **Remediation Name** field, enter a name for the remediation.
The name you choose cannot contain spaces or special characters and should be descriptive. For example, if you have multiple Cisco IOS router instances and multiple remediations for each instance, you may want to specify a name such as `IOS_01_BlockDest`.
5. Optionally, in the **Description** field, enter a description of the remediation.

6. Click **Create**, then click **Done**.
The remediation is added.

Cisco IOS Block Destination Network Remediations

LICENSE: FireSIGHT

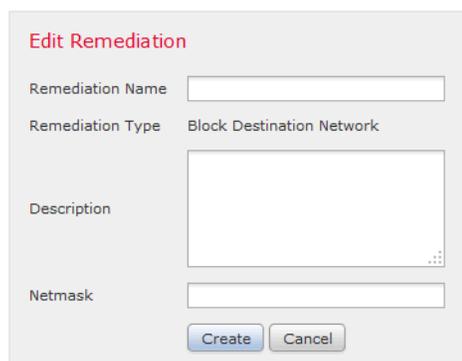
The Cisco IOS Block Destination Network remediation allows you to block any traffic sent from the router to the network of the destination host in a correlation event.

IMPORTANT! Do **not** use this remediation as a response to a correlation rule that is based on a discovery event; discovery events only transmit a source host and not a destination host. You can use this remediation in response to correlation rules that are based on connection events or intrusion events.

To add the remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.
The Instances page appears.
2. Next to the instance where you want to add the remediation, click **View**.
If you have not yet added an instance, see [Adding a Cisco IOS Instance](#) on page 1609.
The Edit Instance page appears.
3. In the **Configured Remediations** section, select **Block Destination Network** and click **Add**.
The Edit Remediation page appears.



The screenshot shows the 'Edit Remediation' form. It has a title 'Edit Remediation' in red. Below the title are four fields: 'Remediation Name' (a text input field), 'Remediation Type' (a dropdown menu with 'Block Destination Network' selected), 'Description' (a large text area with a small icon in the bottom right corner), and 'Netmask' (a text input field). At the bottom of the form are two buttons: 'Create' and 'Cancel'.

4. In the **Remediation Name** field, enter a name for the remediation.
The name you choose cannot contain spaces or special characters and should be descriptive. For example, if you have multiple Cisco IOS router instances and multiple remediations for each instance, you may want to specify a name such as `IOS_01_BlockDestNet`.

5. Optionally, in the **Description** field, enter a description of the remediation.
6. In the **Netmask** field, enter the subnet mask or use CIDR notation to describe the network that you want to block traffic to.
For example, to block traffic to an entire Class C network when a single host triggered a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.
As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.
7. Click **Create**, then click **Done**.
The remediation is added.

Cisco IOS Block Source Remediations

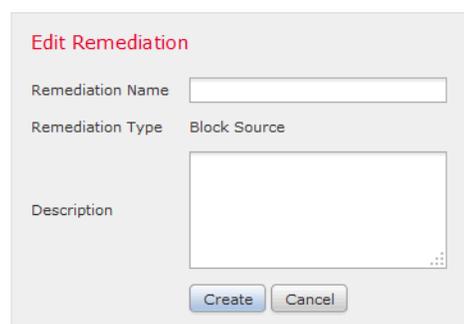
LICENSE: FireSIGHT

The Cisco IOS Block Source remediation allows you to block any traffic sent from the router to the source host included in a correlation event that violates a correlation policy. The source host is the source IP address in the connection event or intrusion event upon which the correlation rule is based, or the host IP address in a discovery event.

To add the remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.
The Instances page appears.
2. Next to the instance where you want to add the remediation, click **View**.
If you have not yet added an instance, see [Adding a Cisco IOS Instance](#) on page 1609.
The Edit Instance page appears.
3. In the **Configured Remediations** section, select **Block Source** and click **Add**.
The Edit Remediation page appears.



The screenshot shows a web form titled "Edit Remediation". It includes a text input field for "Remediation Name", a dropdown menu for "Remediation Type" currently set to "Block Source", and a large text area for "Description". At the bottom of the form are two buttons: "Create" and "Cancel".

4. In the **Remediation Name** field, enter a name for the remediation.
The name you choose cannot contain spaces or special characters and should be descriptive. For example, if you have multiple Cisco IOS router instances and multiple remediations for each instance, you may want to specify a name such as `IOS_01_BlockSrc`.
5. Optionally, in the **Description** field, enter a description of the remediation.
6. Click **Create**, then click **Done**.
The remediation is added.

Cisco IOS Block Source Network Remediations

LICENSE: FireSIGHT

The Cisco IOS Block Source Network remediation allows you to block any traffic sent from the router to the network of the source host in a correlation event. The source host is the source IP address in the connection event or intrusion event upon which the correlation rule is based, or the host IP address in a discovery event.

To add the remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.
The Instances page appears.
2. Next to the instance where you want to add the remediation, click **View**.
If you have not yet added an instance, see [Adding a Cisco IOS Instance](#) on page 1609.
The Edit Instance page appears.
3. In the **Configured Remediations** section, select **Block Source Network** and click **Add**.
The Edit Remediation page appears.

The screenshot shows a web form titled "Edit Remediation". The form is light gray with a white background for the input fields. At the top left, the title "Edit Remediation" is in red. Below the title, there are four main sections: "Remediation Name" with a text input field, "Remediation Type" with a dropdown menu currently showing "Block Source Network", "Description" with a large text area, and "Netmask" with a text input field. At the bottom of the form, there are two buttons: "Create" (highlighted in blue) and "Cancel".

4. In the **Remediation Name** field, enter a name for the remediation.
The name you choose should contain no spaces or special characters and should be descriptive. For example, if you have multiple Cisco IOS router instances and multiple remediations for each instance, you may want to specify a name such as `IOS_01_BlockSourceNet`.
5. Optionally, in the **Description** field, enter a description of the remediation.
6. In the **Netmask** field, enter the subnet mask or CIDR notation that describes the network that you want to block traffic to.
For example, to block traffic to an entire Class C network when a single host triggered a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.
7. Click **Create**, then click **Done**.
The remediation is added.

Configuring Remediations for Cisco PIX Firewalls

LICENSE: FireSIGHT

Sourcefire provides a Cisco PIX Shun remediation module that allows you to block an IP address or network using Cisco's "shun" command. This blocks all traffic sent from either the source or destination host that violated the correlation policy and closes all current connections (note that this will not block traffic sent through the firewall *to* the host).

The Cisco PIX Shun remediation module supports Cisco PIX Firewall 6.0 and higher. You must have level 15 administrative access or higher to launch Cisco PIX remediations.

IMPORTANT! A destination-based remediation only works if you configure it to launch when a correlation rule that is based on a connection event or intrusion event triggers. Discovery events only transmit source hosts.

WARNING! When a Cisco PIX remediation is activated, no timeout period is used. To unblock the IP address or network, you must manually remove the rule from the firewall.

To create remediations for Cisco PIX firewalls:

ACCESS: Admin/Discovery Admin

1. Enable Telnet or SSH (Sourcefire recommends SSH) on the firewall.
Refer to the documentation provided with your Cisco PIX firewall for more information about enabling SSH or Telnet.
2. On the Defense Center, add a Cisco PIX Shun instance for each Cisco PIX firewall you plan to use with the Defense Center.
See [Adding a Cisco PIX Instance](#) on page 1618 for the procedures.
3. Create specific remediations for each instance, based on the type of response you want to elicit on the firewall when correlation policies are violated.

The available remediation types are described in the following sections:

- [Cisco PIX Block Destination Remediations](#) on page 1620
- [Cisco PIX Block Source Remediations](#) on page 1622

4. Begin assigning Cisco PIX remediations to specific correlation policy rules.

Adding a Cisco PIX Instance

LICENSE: FireSIGHT

After you configure SSH or Telnet on the Cisco PIX firewall, you can add an instance to the Defense Center. If you have multiple firewalls you want to send remediations to, you must create a separate instance for each firewall.

IMPORTANT! Sourcefire recommends that you use an SSH connection instead of a Telnet connection. Data transmitted using SSH is encrypted, making it much more secure than Telnet.

To add a Cisco PIX instance:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.

The Instances page appears.

Configured Instances

Instance Name	Module Name	Version	
IOS_01 Cisco IOS Null Route	Cisco IOS Null Route	1.0	 
NmapTest	Nmap Remediation	2.0	 
PIX_01 Cisco PIX Shun	Cisco PIX Shun	1.0	 
Value_01 Set Attribute Value	Set Attribute Value	1.0	 

Add a New Instance

Select a module type

2. From the **Add a New Instance** list, select **Cisco PIX Shun** and click **Add**.

The Edit Instance page appears.

Edit Instance

Instance Name

Module

White List
(an optional list of networks)

Protocol

3. In the **Instance Name** field, type a name for the instance.
The name you choose cannot contain spaces or special characters and should be descriptive. For example, if you intend to connect more than one Cisco firewall, you will have multiple instances, so you may want to choose a name such as **PIX_01**, **PIX_02**, and so on.
4. Optionally, type a description for the instance in the **Description** field.
5. In the **PIX IP** field, enter the IP address of the Cisco PIX firewall you want to use for the remediation.
6. If you require a specific username other than the default (**pix**), type it in the **Username** field.
7. In the **Connection Password** fields, enter the password required to connect to the firewall using SSH or Telnet. The password entered in both fields must match.

8. In the **Enable Password** fields, enter the SSH or Telnet enable password. This is the password used to enter privileged mode on the firewall. The password entered in both fields must match.
9. In the **White List** field, enter IP addresses that you want to exempt from the remediation, one on each line. You can also use CIDR notation or a specific IP address. For example, the following white list is accepted by the system:

```
10.1.1.152
172.16.1.0/24
```

Note that this white list is not associated with any compliance white lists you have created. For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

10. From the **Protocol** list, select the method you want to use to connect to the firewall.
11. Click **Create**.

The instance is created and remediations appear in the Configured Remediations section of the page. You must add specific remediations for them to be used in correlation policies. See the following sections for more information:

- [Cisco PIX Block Destination Remediations](#) on page 1620
- [Cisco PIX Block Source Remediations](#) on page 1622

Cisco PIX Block Destination Remediations

LICENSE: FireSIGHT

The Cisco PIX Block Destination remediation allows you to block traffic sent from the destination host in a correlation event.

IMPORTANT! Do **not** use this remediation as a response to a correlation rule that is based on a discovery event; discovery events only transmit a source host and not a destination host. You can use this remediation in response to correlation rules that are based on connection events or intrusion events.

To add the remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.

The Instances page appears.

Configured Instances

Instance Name	Module Name	Version	
IOS_01 Cisco IOS Null Route	Cisco IOS Null Route	1.0	 
NmapTest	Nmap Remediation	2.0	 
PIX_01 Cisco PIX Shun	Cisco PIX Shun	1.0	 
Value_01 Set Attribute Value	Set Attribute Value	1.0	 

Add a New Instance

Select a module type

2. Next to the instance where you want to add the remediation, click **View**.
If you have not yet added an instance, see [Adding a Cisco PIX Instance](#) on page 1618.

The Edit Instance page appears.

Edit Instance

Instance Name

Module

White List
(an optional list of networks)

Protocol

3. In the **Configured Remediations** section, select **Block Destination** and click **Add**.

The Edit Remediation page appears.

Edit Remediation

Remediation Name

Remediation Type

Description

4. In the **Remediation Name** field, enter a name for the remediation.
The name you choose cannot contain spaces or special characters and should be descriptive. For example, if you have multiple Cisco PIX firewall instances and multiple remediations for each instance, you may want to specify a name such as `PIX_01_BlockDest`.
5. Optionally, in the **Description** field, enter a description of the remediation.
6. Click **Create**, then click **Done**.
The remediation is added.

Cisco PIX Block Source Remediations

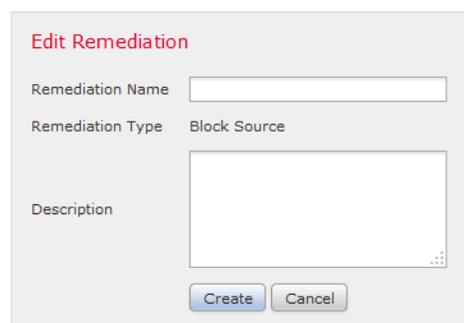
LICENSE: FireSIGHT

The Cisco PIX Block Source remediation allows you to block any traffic sent from the source host included in the event that violates a correlation policy. The source host is the source IP address in the connection event or intrusion event upon which the correlation rule is based, or the host IP address in a discovery event.

To add the remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.
The Instances page appears.
2. Next to the instance where you want to add the remediation, click **View**.
If you have not yet added an instance, see [Adding a Cisco PIX Instance](#) on page 1618.
The Edit Instance page appears.
3. In the **Configured Remediations** section, select **Block Source** and click **Add**.
The Edit Remediation page appears.



4. In the **Remediation Name** field, enter a name for the remediation.
The name you choose cannot contain spaces or special characters and should be descriptive. For example, if you have multiple Cisco PIX firewall instances and multiple remediations for each instance, you may want to specify a name such as `PIX_01_BlockSrc`.

5. Optionally, in the **Description** field, enter a description of the remediation.
The remediation is added.

Configuring Nmap Remediations

LICENSE: FireSIGHT

You can respond to a correlation event by scanning the host where the triggering event occurred. You can choose to scan only the port from the event that triggered the correlation event.

To set up Nmap scanning in response to a correlation event, you must first create an Nmap scan instance, then add an Nmap scan remediation. You can then configure Nmap scanning as responses to violations of rules within the policy.

See the following sections:

- [Adding an Nmap Scan Instance](#) on page 1623
- [Nmap Scan Remediations](#) on page 1625

Adding an Nmap Scan Instance

LICENSE: FireSIGHT

You can set up a separate scan instance for each Nmap module that you want to use to scan hosts on your network for operating system and server information. You can set up scan instances for the local Nmap module on your Defense Center and for any managed devices you want to use to run scans remotely. The results of each scan are always stored on the Defense Center where you configure the scan, even if you run the scan from a remote managed device. To prevent accidental or malicious scanning of mission-critical hosts, you can create a blacklist for the instance to indicate the hosts that should never be scanned with the instance.

Note that you cannot add a scan instance with the same name as any existing scan instance.

To create a scan instance:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.
The Instances page appears.

2. Select **Nmap Remediation (v1.0)** from the **Add a module type** drop-down list and click **Add**.

The Edit Instance page appears.

The screenshot shows the 'Edit Instance' form with the following fields and values:

- Instance Name:** (empty text input)
- Module:** Nmap Remediation (v2.0)
- Description:** (empty text area)
- Black Listed Scan hosts (an optional list of networks):** (empty text area)
- Remote Device Name (optional):** (empty text input)

Buttons: Create, Cancel

3. In the **Instance Name** field, enter a name that includes 1 to 63 alphanumeric characters, with no spaces and no special characters other than underscore (`_`) and dash (`-`).
4. In the **Description** field, specify a description that includes 0 to 255 alphanumeric characters, including spaces and special characters.
5. Optionally, in the **Black Listed Scan hosts** field, specify any hosts or networks that should *never* be scanned with this scan instance, using the following syntax:
 - For IPv6 hosts, an exact IP address (for example, `2001:DB8::fedd:eeff`)
 - For IPv4 hosts, an exact IP address (for example, `192.168.1.101`) or an IP address range using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between `192.168.1.1` and `192.168.1.254`, inclusive)

If you specifically target a scan to a host that is in a blacklisted network, that scan will not run. For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.

6. Optionally, to run the scan from a remote managed device instead of the Defense Center, specify the name or IP address of the managed device in the **Remote Device Name** field.
7. Click **Create**.
The scan instance is created.

Nmap Scan Remediations

LICENSE: FireSIGHT

You can define the settings for an Nmap scan by creating an Nmap remediation. An Nmap remediation can be used as a response in a correlation policy, run on demand, or scheduled to run at a specific time. In order for the results of an Nmap scan to appear in the network map, the scanned host must already exist in the network map. Note that NetFlow, the host input feature, and the system itself can add hosts to the network map.

For more information on the specific settings in an Nmap remediation, see [Understanding Nmap Remediations](#) on page 1693.

Note that Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host for operating system and server data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up-to-date. For more information, see [Automating Nmap Scans](#) on page 1936. Also note that if the host is deleted from the network map, any Nmap scan results for that host are discarded.

For general information about Nmap functionality, refer to the Nmap documentation at <http://insecure.org>.

To create a Nmap remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.

The Scanners page appears.

Nmap Scan Instances

 Nmap_Sample_Instance	  
Test_Remediation Test Remediation	  
 NmapTest	  

2. Click **Add Remediation** next to the scan instance where you want to add a remediation.

The Edit Remediation page appears.

The screenshot shows the 'Edit Remediation' form with the following fields and options:

- Remediation Name: [Text Input]
- Remediation Type: Nmap Scan
- Description: [Text Area]
- Default NSE scripts: On Off
- Timing Template (Higher Is Faster): 3
- Buttons: Create, Cancel

3. In the **Remediation Name** field, type a name for the remediation that includes 1 to 63 alphanumeric characters, with no spaces and no special characters other than underscore (_) and dash (-).
4. In the **Description** field, type a description for the remediation that includes 0 to 255 alphanumeric characters, including spaces and special characters.
5. If you plan to use this remediation in response to a correlation rule that triggers on an intrusion event, a connection event, or a user event, configure the **Scan Which Address(es) From Event?** option.
 - Select **Scan Source and Destination Addresses** to scan the hosts represented by the source IP address and the destination IP address in the event.
 - Select **Scan Source Address Only** to scan the host represented by the event's source IP address.
 - Select **Scan Destination Address Only** to scan the host represented by the event's destination IP address.

If you plan to use this remediation in response to a correlation rule that triggers on a discovery event or a host input event, by default the remediation scans the IP address of the host involved in the event; you do not need to configure this option.

IMPORTANT! Do **not** assign a Nmap remediation as a response to a correlation rule that triggers on a traffic profile change.

6. Configure the **Scan Type** option:
 - To scan quickly in stealth mode on hosts where the `admin` account has raw packet access or where IPv6 is not running, by initiating TCP connections but not completing them, select **TCP Syn Scan**.
 - To scan by using a system `connect()` call, which can be used on hosts where the `admin` account on your Defense Center does not have raw packet access or where IPv6 is running, select **TCP Connect Scan**.
 - To send an ACK packet to check whether ports are filtered or unfiltered, select **TCP ACK Scan**.
 - To send an ACK packet to check whether ports are filtered or unfiltered but also determine whether a port is open or closed, select **TCP Window Scan**.
 - To identify BSD-derived systems using a FIN/ACK probe, select **TCP Maimon Scan**.
7. Optionally, to scan UDP ports in addition to TCP ports, select **On** for the **Scan for UDP ports** option.

TIP! A UDP portscan takes more time than a TCP portscan. To speed up your scans, leave this option disabled.

8. If you plan to use this remediation in response to correlation policy violations, configure the **Use Port From Event** option:
 - Select **On** to scan the port in the correlation event, rather than the ports you specify in step 12.
If you scan the port in the correlation event, note that the remediation scans the port on the IP addresses that you specified in step 8. These ports are also added to the remediation's dynamic scan target.
 - Select **Off** to scan only the ports you will specify in step 12.
9. If you plan to use this remediation in response to correlation policy violations and want to run the scan using the appliance running the detection engine that detected the event, configure the **Scan from reporting detection engine** option:
 - To scan from the appliance running the reporting detection engine, select **On**.
 - To scan from the appliance configured in the remediation, select **Off**.
10. Configure the **Fast Port Scan** option:
 - To only scan ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the managed device that does the scanning, ignoring other port settings, select **On**.
 - To scan all TCP ports, select **Off**.

11. In the **Port Ranges and Scan Order** field, type the ports you want to scan by default, using Nmap syntax, in the order you want to scan those ports.
Specify values from 1 to 65535. Separate ports using commas or spaces. You can also use a hyphen to indicate a port range. When scanning for both TCP and UDP ports, preface the list of TCP ports you want to scan with a T and the list of UDP ports with a U. For example, to scan ports 53 and 111 for UDP traffic, then scan ports 21-25 for TCP traffic, enter **U:53,111,T:21-25**.
Note that the **Use Port From Event** option overrides this setting when the remediation is launched in response to a correlation policy violation, as described in step 8.
12. To probe open ports for server vendor and version information, configure **Probe open ports for vendor and version information**:
 - Select **On** to scan open ports on the host for server information to identify server vendors and versions.
 - Select **Off** to continue using server information for the host.
13. If you choose to probe open ports, set the number of probes used by selecting a number from the **Service Version Intensity** drop-down list:
 - To use more probes for higher accuracy with a longer scan, select a higher number.
 - To use fewer probes for less accuracy with a faster scan, select a lower number.
14. To scan for operating system information, configure **Detect Operating System** settings:
 - Select **On** to scan the host for information to identify the operating system.
 - Select **Off** to continue using operating system information for the host.
15. To determine whether or not host discovery occurs and whether port scans are only run against available hosts, configure **Treat All Hosts As Online**:
 - To skip the host discovery process and run a port scan on every host in the target range, select **On**.
 - To perform host discovery using the settings for **Host Discovery Method** and **Host Discovery Port List** and skip the port scan on any host that is not available, select **Off**.
16. Select the method to be used when Nmap tests to see if a host is present and available:
 - To send an empty TCP packet with the SYN flag set and elicit an RST response on a closed port or a SYN/ACK response on an open port on available hosts, select **TCP SYN**.Note that this option scans port 80 by default and that TCP SYN scans are less likely to be blocked by a firewall with stateful firewall rules.

- To send an empty TCP packet with the ACK flag set and elicit an RST response on available hosts, select **TCP ACK**.
Note that this option scans port 80 by default and that TCP ACK scans are less likely to be blocked by a firewall with stateless firewall rules.
 - To send a UDP packet to elicit port unreachable responses from closed ports on available hosts, select **UDP**. This option scans port 40125 by default.
17. If you want to scan a custom list of ports during host discovery, type a list of ports appropriate for the host discovery method you selected, separated by commas, in **Host Discovery Port List**.
 18. Configure the **Default NSE Scripts** option to control whether to use the default set of Nmap scripts for host discovery and server, operating system, and vulnerability discovery:
 - To run the default set of Nmap scripts, select **On**.
 - To skip the default set of Nmap scripts, select **Off**.See <http://nmap.org/nsedoc/categories/default.html> for the list of default scripts.
 19. To set the timing of the scan process, select a timing template number; select a higher number for a faster, less comprehensive scan and a lower number for a slower, more comprehensive scan.
 20. Click **Save**, then click **Done**.
The remediation is created.

Configuring Set Attribute Remediations

LICENSE: FireSIGHT

You can respond to a correlation event by setting a host attribute value on the host where the triggering event occurred. For text host attributes, you can choose to use the description from the event as the attribute value. For more information on host attributes, see [Working with the Predefined Host Attributes](#) on page 1369 and [Working with User-Defined Host Attributes](#) on page 1370.

To configure setting an attribute value in response to a correlation event, you must first create a set attribute instance, then add a set attribute remediation. You can then configure attribute value updates as responses to violations of rules within the policy.

For more information, see the following sections:

- [Adding a Set Attribute Value Instance](#) on page 1630
- [Set Attribute Value Remediations](#) on page 1631

Adding a Set Attribute Value Instance

LICENSE: FireSIGHT

You can set up an instance to set attribute values in response to correlation rule violations.

To create a set attribute instance:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.

The Instances page appears.

Configured Instances

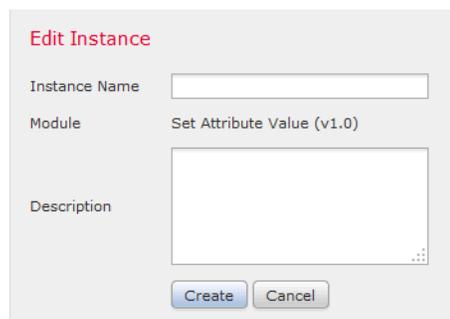
Instance Name	Module Name	Version	
IOS_01 Cisco IOS Null Route	Cisco IOS Null Route	1.0	 
NmapTest	Nmap Remediation	2.0	 
PIX_01 Cisco PIX Shun	Cisco PIX Shun	1.0	 
Value_01 Set Attribute Value	Set Attribute Value	1.0	 

Add a New Instance



2. Select **Set Attribute Value (v1.0)** from the **Add a module type** drop-down list and click **Add**.

The Edit Instance page appears.



3. In the **Instance Name** field, enter a name that includes 1 to 63 alphanumeric characters, with no spaces and no special characters other than underscore (_) and dash (-).
4. In the **Description** field, specify a description that includes 0 to 255 alphanumeric characters, including spaces and special characters.
5. Click **Create**.

The instance is created.

Set Attribute Value Remediations

LICENSE: FireSIGHT

You can create a set attribute value remediation for each attribute value you want to be able to set in response to a correlation rule violation. If the attribute you want to set is a text attribute, you can set the remediation to use the description from the event as the attribute value.

To create a set attribute value remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Instances**.

The Instances page appears.

Configured Instances

Instance Name	Module Name	Version	
IOS_01 Cisco IOS Null Route	Cisco IOS Null Route	1.0	 
NmapTest	Nmap Remediation	2.0	 
PIX_01 Cisco PIX Shun	Cisco PIX Shun	1.0	 
Value_01 Set Attribute Value	Set Attribute Value	1.0	 

Add a New Instance

Select a module type

2. Click **View** next to the scan instance where you want to add a remediation.

The Edit Instance page appears.

Edit Instance

Instance Name Value_01

Module Set Attribute Value (v1.0)

Description

3. Select **Set Attribute Value** from the **Add a new remediation of type** drop-down list. The Edit Remediation page appears.

The screenshot shows the 'Edit Remediation' form with the following fields and values:

- Remediation Name:** [Empty text box]
- Remediation Type:** Set Attribute Value
- Description:** [Empty text area]
- Update Which Host(s) From Event?:** Update Source and Destination Hosts
- Attribute Name:** [Empty text box]
- Use Description From Event For Attribute Value (text attributes only):** On Off
- Attribute Value:** [Empty text box]

Buttons: Create, Cancel

4. In the **Remediation Name** field, type a name for the remediation that includes 1 to 63 alphanumeric characters, with no spaces and no special characters other than underscore (_) and dash (-).
5. In the **Description** field, type a description for the remediation that includes 0 to 255 alphanumeric characters, including spaces and special characters.
6. If you plan to use this remediation in response to a correlation rule that triggers on an intrusion event, user event, or a connection event, configure the **Update Which Host(s) From Event** option.
 - Select **Update Source and Destination Hosts** to update the attribute value on the hosts represented by the source IP address and the destination IP address in the event.
 - Select **Update Source Host Only** to update the attribute value on the host represented by the event's source IP address.
 - Select **Update Destination Host Only** to update the attribute value on the host represented by the event's destination IP address.

If you plan to use this remediation in response to a correlation rule that triggers on a discovery event or host input event, by default the remediation scans the IP address of the host involved in the event; you do not need to configure this option.

7. Configure the **Use Description From Event For Attribute Value (text attributes only)** option:
 - To use the description from the event as the attribute value, select **On**.
 - To use the Attribute Value setting for the remediation as the attribute value, select **Off**.
8. If you are not planning to use the event description, type the attribute value you want to set in the **Attribute Value** field.
9. Click **Save**, then click **Done**.
The remediation is created.

Working with Remediation Status Events

LICENSE: FireSIGHT

When a remediation triggers, a remediation status event is generated. These events are logged to the database and can be viewed on the Remediation Status page. You can search, view, and delete remediation status events.

For more information, see:

- [Setting Event Time Constraints](#) on page 1820
- [Searching for Remediation Status Events](#) on page 1638

Viewing Remediation Status Events

LICENSE: FireSIGHT

The page you see when you access remediation status events differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of remediations. The table view contains a row for each remediation status event. You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The [Remediation Status Actions](#) table below describes some of the specific actions you can perform on a remediation status events workflow page.

Remediation Status Actions

To...	You CAN...
learn more about the columns that appear	find more information in Understanding the Remediation Status Table on page 1636.
modify the time and date range for displayed events	see Setting Event Time Constraints on page 1820. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This can occur even if you configured a sliding time window for the appliance.
sort and constrain the events	see Constraining Events on page 1830 and Sorting Drill-Down Workflow Pages on page 1835.
temporarily use a different workflow	click (switch workflow) by the workflow title. For more information, see Selecting Workflows on page 1809.
navigate to the correlation events view to see associated events	click Correlation Events . For more information, see Navigating Between Workflows on page 1836.
bookmark the current page so that you can quickly return to it	click Bookmark This Page . For more information, see Using Bookmarks on page 1838.
navigate to the bookmark management page	click View Bookmarks . For more information, see Using Bookmarks on page 1838.
generate a report based on the data in the table view	click Report Designer . For more information, see Creating a Report Template from an Event View on page 1725.

Remediation Status Actions (Continued)

To...	You CAN...
drill down to the next page in the workflow, constraining on a specific value	<p>use one of the following methods:</p> <ul style="list-style-type: none">• on a drill-down page that you created in a custom workflow, click a value within a row. Note that clicking a value within a row in a table view constrains the table view and does not drill down to the next page.• To drill down to the next workflow page constraining on some users, select the check boxes next to the users you want to view on the next workflow page, then click View.• To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include “Table View” in the page name.</p> <p>For more information, see Constraining Events on page 1830.</p>
delete remediation status events from the system	<p>use one of the following methods:</p> <ul style="list-style-type: none">• To delete some events, select the check boxes next to events you want to delete, then click Delete.• To delete all events in the current constrained view, click Delete All, then confirm you want to delete all the events.
search for remediation status events	<p>click Search. For more information, see Searching for Remediation Status Events on page 1638.</p>

To view remediation status events:

ACCESS: Admin

- ▶ Select **Analysis > Correlation > Status**.

The first page of the default remediations workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

TIP! If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu by the workflow title, then select **Remediation Status**.

Working with Remediation Status Events

LICENSE: FireSIGHT

You can change the layout of the event view or constrain the events in the view by a field value.

When you disable a column, it is disabled for the duration of your session (unless you add it back later). Note that when you disable the first column, the Count column is added.

Clicking a value within a row in a table view constrains the table view and does not drill down to the next page.

TIP! Table views always include “Table View” in the page name.

For more information, see the following topics:

- [Constraining Events](#) on page 1830.
- [Using Compound Constraints](#) on page 1832.
- [Sorting Drill-Down Workflow Pages](#) on page 1835.
- [Understanding the Remediation Status Table](#) on page 1636

Understanding the Remediation Status Table

LICENSE: FireSIGHT

You can configure the Defense Center to launch a variety of responses to policy violations and to discovery events. These responses include remediations, such as blocking a host at the firewall or router when it violates a policy. When a remediation triggers, a remediation status event is generated and logged to the database. For more information on remediations, see [Configuring Remediations](#) on page 1606.

The fields in the remediation status table are described in the [Remediation Status Fields](#) table.

Remediation Status Fields

FIELD	DESCRIPTION
Policy	The name of the correlation policy that was violated and triggered the remediation.
Remediation Name	The name of the remediation that was launched.

Remediation Status Fields (Continued)

FIELD	DESCRIPTION
Result Message	<p>A message that describes what happened when the remediation was launched. Status messages include:</p> <ul style="list-style-type: none"> • Successful completion of remediation • Error in the input provided to the remediation module • Error in the remediation module configuration • Error logging into the remote device or server • Unable to gain required privileges on remote device or server • Timeout logging into remote device or server • Timeout executing remote commands or servers • The remote device or server was unreachable • The remediation was attempted but failed • Failed to execute remediation program • Unknown/unexpected error <p>IMPORTANT! If custom remediation modules are installed, you may see additional status messages that are implemented by the custom module.</p>
Rule	The name of the rule that triggered the remediation.
Time	The date and time that the Defense Center launched the remediation
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

To display the table view of remediation status events:

ACCESS: Admin

- ▶ Select **Analysis > Correlation > Status**.

The table view appears. For information on working with remediation status events, see [Working with Remediation Status Events](#) on page 1633.

TIP! If you are using a custom workflow that does not include the table view of remediation status events, click **(switch workflow)** by the workflow title, then click **Remediation Status**.

Searching for Remediation Status Events

LICENSE: FireSIGHT

You can search for remediation status events to determine when and if a particular remediation was launched. You may want to create searches customized for your network environment, then save them to reuse later. The search criteria you can use are described in the [Remediation Status Search Criteria](#) table.

Remediation Status Search Criteria

SEARCH FIELD	DESCRIPTION
Result Message	<p>Enter the exact name of the result message (a message that describes what happened when the remediation was launched) you want to match. Valid status messages are:</p> <ul style="list-style-type: none">• Successful completion of remediation• Error in the input provided to the remediation module• Error in the remediation module configuration• Error logging into the remote device or server• Unable to gain required privileges on remote device or server• Timeout logging into remote device or server• Timeout executing remote commands or servers• The remote device or server was unreachable• The remediation was attempted but failed• Failed to execute remediation program• Unknown/unexpected error <p>IMPORTANT! If you installed custom remediation modules, you may be able to enter additional status messages implemented by the custom module.</p>
Time	<p>Specify the date and time the Defense Center launched the remediation. See Specifying Time Constraints in Searches on page 1775 for the syntax for entering time.</p>
Remediation Name	<p>Enter the exact name of the remediation that was launched. This is the name you specified when you created the remediation.</p>
Policy	<p>Enter the name of the correlation policy that triggered the remediation.</p>
Rule	<p>Enter the name of the correlation rule that triggered the remediation.</p>

For more information on searching, including how to load and delete saved searches, see [Searching for Events](#) on page 1770.

To search for remediation status events:

ACCESS: Admin

1. Select **Analysis & Reporting > Searches > Remediation Status**.

The Remediation Status search page appears.

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: Remediation Status

Name: Search 1, My Search

Save As Private:

Constraint

Result Message: string

Time: > 2009-07-16 13:00:31, < today at 4:30pm

Remediation Name: name

Policy: Policy 1, My Policy

Rule: Rule 1, My Rule

Buttons: Search, Save As New Search

TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

2. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is automatically created when you save the search.
3. Enter your search criteria in the appropriate fields, as described in the [Remediation Status Search Criteria](#) table. If you enter multiple criteria, the search returns only the records that match all the criteria.
4. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

5. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default remediation status workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

- Click **Save** if you are modifying an existing search and want to save your changes.
- Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

CHAPTER 40

ENHANCING NETWORK DISCOVERY

The information about your network traffic collected by the Sourcefire 3D System is most valuable to you when the system can correlate this information to identify the hosts on your network that are most vulnerable and most important.

As an example, if you have several devices on your network running a customized version of SuSE Linux, the system cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for SuSE Linux, you may want to create a custom fingerprint for one of the hosts that can then be used to identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for SuSE Linux in the fingerprint to associate that list with each host that matches the fingerprint.

The system also allows you to input host data from third-party systems directly into the network map, using the host input feature. However, third-party operating system or application data does not automatically map to vulnerability information. If you want to see vulnerabilities and perform impact correlation for hosts using third-party operating system, server, and application protocol data, you must map the vendor and version information from the third-party system to the vendor and version listed in the vulnerability database (VDB). You also may want to maintain the host input data on an ongoing basis. Note that even if you map application data to Sourcefire 3D System vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

If the system cannot identify application protocols running on hosts on your network, you can create user-defined application protocol detectors that allow the system to identify the applications based on a port or a pattern. You can also

import, activate, and deactivate certain application detectors to further customize the application detection capability of the Sourcefire 3D System.

You can also replace detection of operating system and application data using scan results from the Nmap active scanner or augment the vulnerability lists with third-party vulnerabilities. The system may reconcile data from multiple sources to determine the identity for an application. For more information on how the system does this, see [Understanding Current Identities](#) on page 1646. For more information on active scanning, see [Configuring Active Scanning](#) on page 1692.

For more information, see the following sections:

- [Assessing Your Detection Strategy](#) on page 1642
- [Enhancing Your Network Map](#) on page 1644
- [Using Custom Fingerprinting](#) on page 1649
- [Working with Application Detectors](#) on page 1663
- [Importing Host Input Data](#) on page 1680

Assessing Your Detection Strategy

LICENSE: FireSIGHT

Before you make any changes to the system's default detection capabilities, you should analyze what hosts are not being identified correctly and why, so you can decide what solution to implement. Use the following as a guide for your decision:

- [Are Your Managed Devices Correctly Placed?](#) on page 1642
- [Do Unidentified Operating Systems Have a Unique TCP Stack?](#) on page 1643
- [Can the Sourcefire 3D System Identify All Applications?](#) on page 1644
- [Have You Applied Patches that Fix Vulnerabilities?](#) on page 1644
- [Do You Want to Track Third-Party Vulnerabilities?](#) on page 1644

Are Your Managed Devices Correctly Placed?

LICENSE: FireSIGHT

If network devices such as load balancers, proxy servers, or NAT devices reside between the managed device and the unidentified or misidentified host, place a managed device closer to the misidentified host rather than using custom fingerprinting. Sourcefire does not recommend using custom fingerprinting in this scenario.

Do Unidentified Operating Systems Have a Unique TCP Stack?

LICENSE: FireSIGHT

If the system misidentifies a host, you should investigate why the host is misidentified to help you decide between creating and activating a custom fingerprint or substituting Nmap or host input data for discovery data.

WARNING! If you encounter misidentified hosts, contact your support representative before creating custom fingerprints.

If a host is running an operating system that is not detected by the system by default and does not share identifying TCP stack characteristics with existing detected operating systems, you should create a custom fingerprint.

For example, if you have a customized version of Linux with a unique TCP stack that the system cannot identify, you would benefit from creating a custom fingerprint, which allows the system to identify the host and continuing monitoring it, rather than using scan results or third-party data, which require you to actively update the data yourself on an ongoing basis.

Note that many open source Linux distributions use the same kernel, and as such, the system identifies them using the Linux kernel name. If you create a custom fingerprint for a Red Hat Linux system, you may see other operating systems (such as Debian Linux, Mandrake Linux, Knoppix, and so on) identified as Red Hat Linux, because the same fingerprint matches multiple Linux distributions.

You should not use a fingerprint in every situation. For example, a modification may have been made to a host's TCP stack so that it resembles or is identical to another operating system. For example, an Apple Mac OS X host is altered, making its fingerprint identical to a Linux 2.4 host, causing the system to identify it as Linux 2.4 instead of Mac OS X. If you create a custom fingerprint for the Mac OS X host, it may cause all legitimate Linux 2.4 hosts to be erroneously identified as Mac OS X hosts. In this case, if Nmap correctly identifies the host, you could schedule regular Nmap scans for that host.

If you import data from a third-party system using host input, you must map the vendor, product, and version strings that the third party uses to describe servers and application protocols to the Sourcefire definitions for those products. For more information, see [Managing Third-Party Product Mappings](#) on page 1682. Note that even if you map application data to Sourcefire 3D System vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

The system may reconcile data from multiple sources to determine the current identity for an operating system or application. For more information on how the system does this, see [Understanding Current Identities](#) on page 1646.

For Nmap data, you can schedule regular Nmap scans. For host input data, you can regularly run the Perl script for the import or the command line utility.

However, note that active scan data and host input data may not be updated with the frequency of discovery data.

Can the Sourcefire 3D System Identify All Applications?

LICENSE: FireSIGHT

If a host is correctly identified by the system but has unidentified applications, you can create a user-defined detector to provide the system with port and pattern matching information to help identify the application. For more information, see [Creating a User-Defined Application Protocol Detector](#) on page 1666.

Have You Applied Patches that Fix Vulnerabilities?

LICENSE: FireSIGHT

If the system correctly identifies a host but does not reflect applied fixes, you can use the host input feature to import patch information. When you import patch information, you must map the fix name to a fix in the database. For more information, see [Mapping Third-Party Product Fixes](#) on page 1685.

Do You Want to Track Third-Party Vulnerabilities?

LICENSE: FireSIGHT

If you have vulnerability information from a third-party system that you want to use for impact correlation, you can map the third-party vulnerability identifiers for servers and application protocols to vulnerability identifiers in the Sourcefire database and then import the vulnerabilities using the host input feature. For more information on using the host input feature, see the *Sourcefire 3D System Host Input API Guide*. For more information on mapping third-party vulnerabilities, see [Mapping Third-Party Vulnerabilities](#) on page 1687. Note that even if you map application data to Sourcefire 3D System vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

Enhancing Your Network Map

LICENSE: FireSIGHT

The Sourcefire 3D System builds the network map using data it detects by passively analyzing traffic. It also uses data added through active sources such as the host input feature and the Nmap scanner. Understanding how the system decides which data to use for an application or operating system identity can help you decide how best to augment the system's passive detection capabilities with active input sources.

For more information, see the following topics:

- [Understanding Passive Detection](#) on page 1645
- [Understanding Active Detection](#) on page 1645
- [Understanding Current Identities](#) on page 1646
- [Understanding Identity Conflicts](#) on page 1648

Understanding Passive Detection

LICENSE: FireSIGHT

Passive detection is the detection of host operating system, client, and application information through analysis of traffic passively collected by the system. The system uses information in the VDB to help it identify your network assets.

If the system cannot identify an operating system on a host, you can manually determine it and then create a custom server or client fingerprint to help the system recognize that operating system on other hosts with similar operating system characteristics.

The system uses all collected passive fingerprints for a host operating system to create a *derived fingerprint*. The system creates derived fingerprints by applying a formula which calculates the most likely identity using the confidence value of each collected fingerprint and the amount of corroborating fingerprint data between identities. Common elements are identified between identities.

If you use user-defined application detectors on your network, you can augment the system's application detection capabilities by creating custom detectors that provide the system with the information it needs to identify those applications. NetFlow can also add passively detected application information to the network map.

Note that the system does not use application protocol and operating system data that it classified as *unknown* because it is unable to interpret the data. The managed device reports the identity to the Defense Center as **unknown** and the identity data is not used to derive fingerprints.

Understanding Active Detection

LICENSE: FireSIGHT

Active detection is addition, to the network map, of data collected by active sources, such as host operating system and application information. For example, you can use the Nmap scanner to actively scan the hosts that you target on your network. Nmap discovers operating systems and applications on hosts.

In addition, the host input feature allows you to actively add *host input data* to the network map. There are two different categories of host input data:

- You can modify a host's operating system or application identity through the Sourcefire 3D System user interface. Data added through the interface is *user input data*.
- You can also import data using a command line utility. Imported data is *host import input data*.

The system retains one identity for each active source. When you run an Nmap scan instance, for example, the results of the previous scan are replaced with the new scan results. However, if you run an Nmap scan and then replace those results with data from a client whose results are imported through the command line, the system retains both the identities from the Nmap results and the identities from the import client. Then the system uses the priorities set in the system policy to determine which active identity to use as the current identity.

Note that user input is considered one source, even if it comes from different users. As an example, if UserA sets the operating system through the host profile, and then UserB changes that definition through the host profile, the definition set by UserB is retained, and the definition set by UserA is discarded. In addition, note that user input overrides all other active sources and is used as the current identity if it exists.

Understanding Current Identities

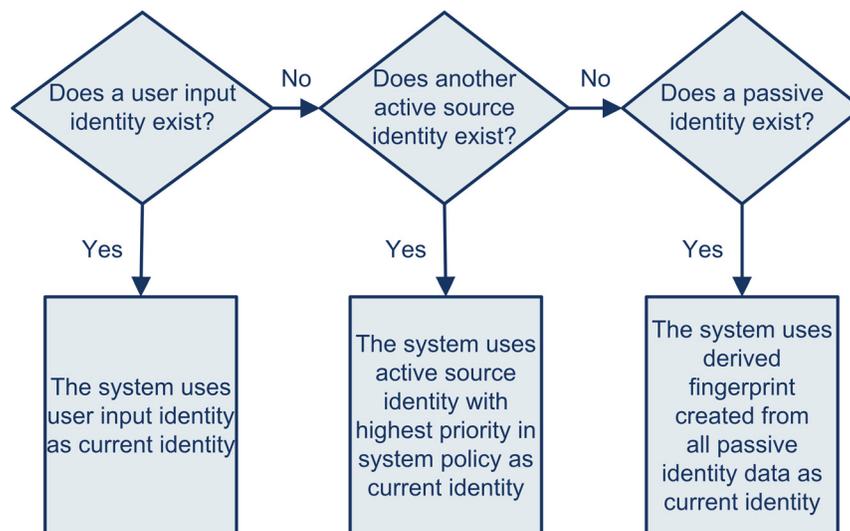
LICENSE: FireSIGHT

The *current identity* for an application or an operating system on a host is the identity that the system finds most likely to be correct.

The system uses the current identity for an operating system or application for the following purposes:

- to assign vulnerabilities to a host
- for impact assessment
- when evaluating correlation rules written against operating system identifications, host profile qualifications, and compliance white lists
- for display in the Hosts and Servers table views in workflows
- for display in the host profile
- to calculate the operating system and application statistics on the Discovery Statistics page

The system uses source priorities to determine which active identity should be used as the current identity for an application or operating system.



For example, if a user sets the operating system to Windows 2003 Server on a host, Windows 2003 Server is the current identity. Attacks which target Windows 2003 Server vulnerabilities on that host are given a higher impact, and the vulnerabilities listed for that host in the host profile include Windows 2003 Server vulnerabilities.

The database may retain information from several sources for the operating system or for a particular application on a host.

The system treats an operating system or application identity as the current identity when the source for the data has the highest source priority. Possible sources have the following priority order:

1. user
2. scanner and application (set in the network discovery policy)
3. managed devices
4. NetFlow

Note that a new higher priority application identity will not override a current application identity if it has less detail than the current identity.

In addition, note that when an identity conflict occurs, the resolution of the conflict depends on settings in the network discovery policy or on your manual resolution, as described in [Understanding Identity Conflicts](#) on page 1648.

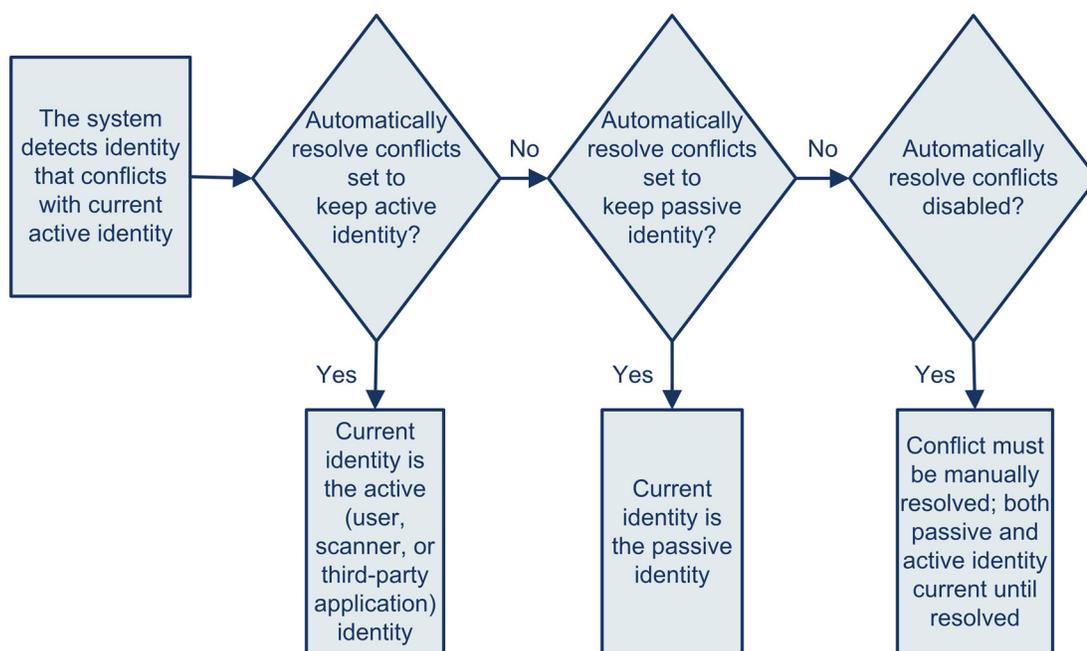
Understanding Identity Conflicts

LICENSE: FireSIGHT

An *identity conflict* occurs when the system reports a new passive identity that conflicts with the current active identity and previously reported passive identities. For example, the previous passive identity for an operating system is reported as Windows 2000, then an active identity of Windows XP becomes current. Next, the system detects a new passive identity of Ubuntu Linux 8.04.1. The Windows XP and the Ubuntu Linux identities are in conflict.

When an identity conflict exists for the identity of the host's operating system or one of the applications on the host, the system lists both conflicting identities as current and uses both for impact assessment until the conflict is resolved.

A user with Administrator privileges can resolve identity conflicts automatically by choosing to always use the passive identity or always use the active identity. Unless you disable automatic resolution of identity conflicts, identity conflicts are always automatically resolved.



A user with Administrator privileges can also configure the system to generate an event when an identity conflict occurs. That user can then set up a correlation policy with a correlation rule that uses an Nmap scan as a correlation response. When an event occurs, Nmap scans the host to obtain updated host operating system and application data.

Using Custom Fingerprinting

LICENSE: FireSIGHT

The Sourcefire 3D System includes operating system *fingerprints* that the system uses to identify the operating system on each host it detects. However, sometimes the system cannot identify a host operating system or misidentifies it because no fingerprints exist that match the operating system. To correct this problem, you can create a *custom fingerprint*, which provides a pattern of operating system characteristics unique to the unknown or misidentified operating system, to supply the name of the operating system for identification purposes.

If the system cannot match a host's operating system, it cannot identify the vulnerabilities for the host, because the system derives the list of vulnerabilities for each host from its operating system fingerprint. For example, if the system detects a host running Microsoft Windows, the system has a stored Microsoft Windows vulnerability list that it adds to the host profile for that host based on the detected Windows operating system.

As an example, if you have several devices on your network running a new beta version of Microsoft Windows, the system cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for Microsoft Windows, you may want to create a custom fingerprint for one of the hosts that can then be used to identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for Microsoft Windows in the fingerprint to associate that list with each host that matches the fingerprint.

When you create a custom fingerprint, you can add a customized display of operating system information, and you can select the operating system vendor, product name, and product version for the operating system which the system should use as a model for the vulnerability list for the fingerprint. The Defense Center lists the set of vulnerabilities associated with that fingerprint for any hosts running the same operating system. If the custom fingerprint you create does not have any vulnerabilities mappings in it, the system uses the fingerprint to assign the custom operating system information you provide in the fingerprint. When the system sees new traffic from a host that has already been detected and currently resides in the network map, the system updates the host with the new fingerprint information. the system also uses the new fingerprint to identify any new hosts with that operating system the first time they are detected.

Before attempting to fingerprint a host, you should determine why the host is not being identified correctly to decide whether custom fingerprinting is a viable solution. For more information, see [Assessing Your Detection Strategy](#) on page 1642.

You can create two types of fingerprints with the system:

- Client fingerprints, which identify operating systems based on the SYN packet that the host sends when it connects to a TCP application running on another host on the network.

See [Fingerprinting Clients](#) on page 1650 for information about how to obtain a client fingerprint for a host.

- Server fingerprints, which identify operating systems based on the SYN-ACK packet that the host uses to respond to an incoming connection to a running TCP application.

See [Fingerprinting Servers](#) on page 1655 for information about how to obtain a server fingerprint for a host.

After creating fingerprints, you must activate them before the system can associate them with hosts. See [Managing Fingerprints](#) on page 1660 for more information.

IMPORTANT! If both a client and server fingerprint match the same host, the client fingerprint is used.

Fingerprinting Clients

LICENSE: FireSIGHT

Client fingerprints identify operating systems based on the SYN packet a host sends when it connects to a TCP application running on another host on the network.

If the Defense Center does not have direct contact with monitored hosts, you can specify a device that is managed by the Defense Center and is closest to the host you intend to fingerprint when specifying client fingerprint properties.

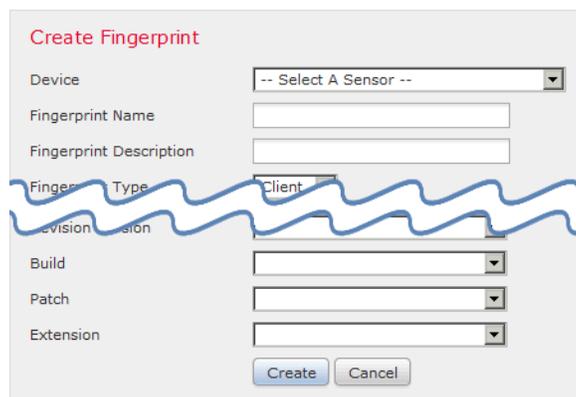
Before you begin the fingerprinting process, obtain the following information about the host you want to fingerprint:

- The number of network hops between the host and the Defense Center or the device you use to obtain the fingerprint. (Sourcefire strongly recommends that you directly connect the Defense Center or the device to the same subnet that the host is connected to.)
- The network interface (on the Defense Center or the device) that is connected to the network where the host resides.
- The actual operating system vendor, product, and version of the host.
- Access to the host in order to generate client traffic.

To obtain a client fingerprint for a host:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**.
The Custom Fingerprint page appears.
2. Click **Create Custom Fingerprint**.
The Create Custom Fingerprint page appears.



The screenshot shows the 'Create Fingerprint' form with the following fields and values:

- Device: -- Select A Sensor --
- Fingerprint Name: (empty)
- Fingerprint Description: (empty)
- Fingerprint Type: Client
- Revision: (empty)
- Build: (empty)
- Patch: (empty)
- Extension: (empty)

Buttons: Create, Cancel

3. From the **Device** drop-down list, select the Defense Center or the device that you want to use to collect the fingerprint.



Device: linden.example.com/10.10.10.10

4. In the **Fingerprint Name** field, type an identifying name for the fingerprint.



Fingerprint Name: (empty)

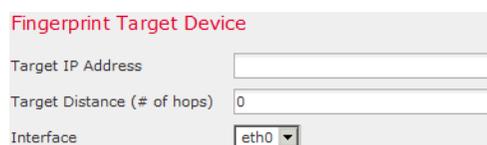
Fingerprint Description: (empty)

Fingerprint Type: Server

5. In the **Fingerprint Description** field, type a description for the fingerprint.
6. From the **Fingerprint Type** list, select **Client**.

7. In the **Target IP Address** field, type an IP address of the host you want to fingerprint. Note that the fingerprint will only be based on traffic to and from the host IP address you specify, not any of the host's other IP addresses (if it has any).

WARNING! You can capture IPv6 fingerprints only with appliances running Version 5.2 and later of the Sourcefire 3D System. These appliances must have IPv6 capability enabled. For information on enabling IPv6 on managed devices and Defense Centers, see [Configuring Network Settings](#) on page 2009.



Fingerprint Target Device

Target IP Address

Target Distance (# of hops)

Interface

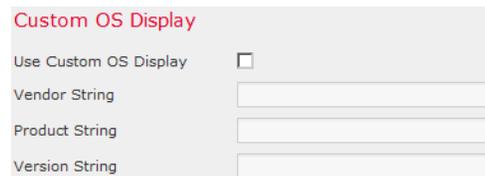
8. In the **Target Distance** field, enter the number of network hops between the host and the device that you selected in step 3 to collect the fingerprint.

WARNING! This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by the system.

9. From the **Interface** list, select the network interface that is connected to the network segment where the host resides.

WARNING! Sourcefire recommends that you do **not** use the sensing interface on a managed device for fingerprinting for several reasons. First, fingerprinting does not work if the sensing interface is on a span port. Also, if you use the sensing interface on a device, the device stops monitoring the network for the amount of time it takes to collect the fingerprint. You can, however, use the management interface or any other available network interfaces to perform fingerprint collection. If you do not know which interface is the sensing interface on your device, refer to the *Installation Guide* for the specific model you are using to fingerprint.

10. If you want to display custom information in the host profile for fingerprinted hosts (or if the host you want to fingerprint does not reside in the OS Vulnerability Mappings section), select **Use Custom OS Display** in the Custom OS Display section and provide the values you want to display in host profiles for the following:
- In the **Vendor String** field, type the operating system's vendor name. For example, the vendor for Microsoft Windows would be Microsoft.
 - In the **Product String** field, type the operating system's product name. For example, the product name for Microsoft Windows 2000 would be Windows.
 - In the **Version String** field, type the operating system's version number. For example, the version number for Microsoft Windows 2000 would be 2000.



The screenshot shows a configuration panel titled "Custom OS Display". It contains a checkbox labeled "Use Custom OS Display" which is currently unchecked. Below the checkbox are three text input fields: "Vendor String", "Product String", and "Version String".

11. In the OS Vulnerability Mappings section, select the operating system, product, and versions you want to use for vulnerability mapping.



The screenshot shows a section titled "OS Vulnerability Mappings". It contains two dropdown menus. The first is labeled "Vendor" and has a placeholder text "-- Select A Vendor --". The second is labeled "Extension" and is currently empty. A blue wavy line is drawn across the top of the form area.

For example, if you want your custom fingerprint to assign the list of vulnerabilities from Redhat Linux 9 to matching hosts, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the major version.

TIP! When creating a fingerprint, you assign a single vulnerability mapping for the fingerprint. After the fingerprint is created and activated, you can add additional vulnerability mappings for other versions of the operating system. See [Editing an Active Fingerprint](#) on page 1663 for more information.

You must specify a Vendor and Product name in this section if you want to use the fingerprint to identify vulnerabilities for matching hosts or if you do not assign custom operating system display information. To map vulnerabilities for all versions of an operating system, specify only the vendor and product name. For example, to add all versions of the Palm OS, you would select **PalmSource, Inc.** from the **Vendor** list, **Palm OS** from the **Product** list, and leave all other lists at their default settings.

IMPORTANT! Not all options in the **Major Version, Minor Version, Revision Version, Build, Patch, and Extension** drop-down lists may apply to the operating system you choose. In addition, if no definition appears in a list that matches the operating system you want to fingerprint, you can leave these values empty. Be aware that if you do not create any OS vulnerability mappings in a fingerprint, the system cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

12. Click **Create**.

The Custom Fingerprint status page reappears. The status page refreshes every ten seconds until it receives data from the host in question.

TIP! When you click **Create**, the status briefly shows **New**, then switches to **Pending**, where it remains until traffic is seen for the fingerprint, then the status switches to **Ready**.

13. Using the IP address you specified as the target IP address, access the host you are trying to fingerprint and initiate a TCP connection to the appliance.

For example, access the web interface of the Defense Center from the host you want to fingerprint or SSH into the Defense Center from the host. If you are using SSH, use the following command:

```
ssh -b localIPv6address DCmanagementIPv6address
```

where *localIPv6address* is the IPv6 address specified in step 7 that is currently assigned to the host and *DCmanagementIPv6address* is the management IPv6 address of the Defense Center.

The Custom Fingerprint page should then reload with a “Ready” status.

IMPORTANT! To create an accurate fingerprint, traffic **must** be seen by the appliance collecting the fingerprint. If you are connected through a switch, traffic to a system other than the appliance may not be seen by the system.

14. After the fingerprint is created, you must activate it before the Defense Center can use it to identify hosts. See [Managing Fingerprints](#) on page 1660 for more information.

Fingerprinting Servers

LICENSE: FireSIGHT

Server fingerprints identify operating systems based on the SYN-ACK packet that the host uses to respond to an incoming connection to a running TCP application. Before you begin, you should obtain the following information about the host you want to fingerprint:

- The number of network hops between the host and the appliance you use to obtain the fingerprint. Sourcefire strongly recommends that you directly connect an unused interface on the appliance to the same subnet that the host is connected to.
- The network interface (on the appliance) that is connected to the network where the host resides.
- The actual operating system vendor, product, and version of the host.
- An IP address that is not currently in use and is authorized on the network where the host is located.

TIP! If the Defense Center does not have direct contact with monitored hosts, you can specify a managed device that is closest to the host you intend to fingerprint when specifying server fingerprint properties.

To obtain a server fingerprint for a host:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**.
The Custom Fingerprint page appears.
2. Click **Create Custom Fingerprint**.
The Create Custom Fingerprint page appears.

Create Fingerprint

Device: -- Select A Sensor --

Fingerprint Name:

Fingerprint Description:

Fingerprint Type: Server

Build:

Patch:

Extension:

Buttons: Create, Cancel

3. From the **Device** list, select the Defense Center or the managed device that you want to use to collect the fingerprint.

Device: linden.example.com/10.10.10.10

4. In the **Fingerprint Name** field, type an identifying name for the fingerprint.

Fingerprint Name:

Fingerprint Description:

Fingerprint Type: Server

5. In the **Fingerprint Description** field, type a description for the fingerprint.
6. From the **Fingerprint Type** list, select **Server**.
Server fingerprinting options appear.

Fingerprint Target Device

Target IP Address:

Target Distance (# of hops): 0

Interface: eth0

Server Port:

7. In the **Target IP Address** field, type an IP address of the host you want to fingerprint. Note that the fingerprint will only be based on traffic to and from the host IP address you specify, not any of the host's other IP addresses (if it has any).

WARNING! You can capture IPv6 fingerprints only with appliances running Version 5.2 and later of the Sourcefire 3D System.

8. In the **Target Distance** field, enter the number of network hops between the host and the device that you selected in step 3 to collect the fingerprint.

WARNING! This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by the system.

9. From the **Interface** list, select the network interface that is connected to the network segment where the host resides.

WARNING! Sourcefire recommends that you do **not** use the sensing interface on a managed device for fingerprinting for several reasons. First, fingerprinting does not work if the sensing interface is on a span port. Also, if you use the sensing interface on a device, the device stops monitoring the network for the amount of time it takes to collect the fingerprint. You can, however, use the management interface or any other available network interfaces to perform fingerprint collection. If you do not know which interface is the sensing interface on your device, refer to the *Installation Guide* for the specific model you are using to fingerprint.

10. Click **Get Active Ports**.

If the system has detected any open ports on the host, they appear in the drop-down list.

11. In the **Server Port** field, type the port that you want the device selected to collect the fingerprint to initiate contact with, or select a port from the **Get Active Ports** drop-down list.

You can use any server port that you know is open on the host (for instance, 80 if the host is running a web server).

12. In the **Source IP Address** field, type an IP address that should be used to attempt to communicate with the host.



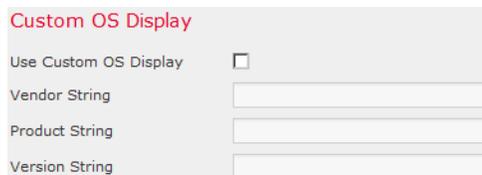
You should use a source IP address that is authorized for use on the network but is not currently being used, for example, a DHCP pool address that is currently not in use. This prevents you from temporarily knocking another host offline while you create the fingerprint.

In addition, you should exclude that IP address from monitoring in your network discovery policy while you create the fingerprint. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address. For more information, see [Understanding Discovery Data Collection](#) on page 1251.

13. In the **Source Subnet Mask** field, type the subnet mask for the IP address you are using.
14. If the **Source Gateway** field appears, enter the default gateway IP address that should be used to establish a route to the host.

The **Source Gateway** field appears if the target distance (number of hops) is 1 or higher and you are using an interface other than the management interface to connect to the network where the host resides.

15. If you want to display custom information in the host profile for fingerprinted hosts or if the fingerprint name you want to use does not exist in the OS Definition section, select **Use Custom OS Display** in the Custom OS Display section.



Provide the values you want to appear in host profiles for the following:

- In the **Vendor String** field, type the operating system's vendor name. For example, the vendor for Microsoft Windows would be Microsoft.
- In the **Product String** field, type the operating system's product name. For example, the product name for Microsoft Windows 2000 would be Windows.
- In the **Version String** field, type the operating system's version number. For example, the version number for Microsoft Windows 2000 would be 2000.

16. In the OS Vulnerability Mappings section, select the operating system, product, and versions you want to use for vulnerability mapping. For example, if you want your custom fingerprint to assign the list of vulnerabilities from Redhat Linux 9 to matching hosts, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.



The screenshot shows a section titled "OS Vulnerability Mappings". It contains two dropdown menus. The first is labeled "Vendor" and has a placeholder text "-- Select A Vendor --". The second is labeled "Extension" and is currently empty. The background of the form has a blue wavy pattern.

TIP! When creating a fingerprint, you assign a single vulnerability mapping for the fingerprint. After the fingerprint is created and activated, you can add additional vulnerability mappings for other versions of the operating system. See [Editing an Active Fingerprint](#) on page 1663 for more information.

You must specify a Vendor and Product name in this section if you want to use the fingerprint to identify vulnerabilities for matching hosts or if you do not assign custom operating system display information. To map vulnerabilities for all versions of an operating system, specify only the vendor and product name. For example, to add all versions of the Palm OS, you would select **PalmSource, Inc.** from the **Vendor** list, **Palm OS** from the **Product** list, and leave all other lists at their default settings.

IMPORTANT! Not all options in the **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** drop-down lists may apply to the operating system you choose. In addition, if no definition appears in a list that matches the operating system you want to fingerprint, you can leave these values empty. Be aware that if you do not create any OS vulnerability mappings in a fingerprint, the system cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

17. Click **Create**.
18. The Custom Fingerprint status page appears. It reloads every ten seconds and should reload with a "Ready" status.

IMPORTANT! If the target system stops responding during the fingerprinting process, the status shows an **ERROR: No Response** message. If you see this message, submit the fingerprint again. Wait three to five minutes (the time period may vary depending on the target system), click the edit icon (✎) to access the Custom Fingerprint page, and then click **Create**.

19. After the fingerprint is created, activate it and, optionally, add vulnerability mappings. See [Managing Fingerprints](#) on page 1660 for more information.

Managing Fingerprints

LICENSE: FireSIGHT

You can activate, deactivate, delete, view, and edit custom fingerprints. When creating a fingerprint, you assign a single vulnerability mapping for the fingerprint. For more information on creating a fingerprint, see [Fingerprinting Clients](#) on page 1650 and [Fingerprinting Servers](#) on page 1655. After the fingerprint is created and activated, you can edit the fingerprint to make changes or add vulnerability mappings.

To access the Custom Fingerprints page:

ACCESS: Admin/Discovery Admin

- ▶ Select **Policies > Network Discovery**, then click **Custom Operating Systems**. The Custom Fingerprint page appears.

Name	Status	Type	Vendor	Product	Version	
Sample Client Fingerprint Sample Client Fingerprint	New	Client	Red Hat, Inc.	Fedora Core	6	 

If the system is awaiting data to create a fingerprint, it automatically refreshes the page every 10 seconds until the fingerprint is created.

See the following sections for more information:

- [Activating Fingerprints](#) on page 1660
- [Deactivating Fingerprints](#) on page 1661
- [Deleting Fingerprints](#) on page 1661
- [Editing Fingerprints](#) on page 1662

Activating Fingerprints

LICENSE: FireSIGHT

After creating a custom fingerprint, you must activate it before the system can use it to identify hosts. After the new fingerprint is activated, the system uses it to re-identify previously discovered hosts and discover new hosts.

To activate a fingerprint:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**. The Custom Fingerprint page appears.

Name	Status	Type	Vendor	Product	Version	
Sample Client Fingerprint Sample Client Fingerprint	New	Client	Red Hat, Inc.	Fedora Core	6	 

2. Click the slider next to the fingerprint you want to activate.

IMPORTANT! The activate option is only available if the fingerprint you created is valid. If the slider is not available, try creating the fingerprint again.

The Defense Center activates the fingerprint and propagates it to all managed devices. The icon next to the fingerprint name changes to indicate that the fingerprint is active.

Deactivating Fingerprints

LICENSE: FireSIGHT

If you want to stop using a fingerprint, you can deactivate it. Deactivating a fingerprint causes a fingerprint to no longer be used, but allows it to remain on the system. When you deactivate a fingerprint, the operating system is marked as unknown for hosts that use the fingerprint. If the hosts are detected again and match a different active fingerprint, they are then identified by that active fingerprint.

Deleting a fingerprint removes it from the system completely. After deactivating a fingerprint, you can delete it.

To deactivate an active fingerprint:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**.
The Custom Fingerprint page appears.
2. Click the slider next to the active fingerprint you want to deactivate.
The Defense Center deactivates the fingerprint and propagates the deactivation to all managed devices.

Deleting Fingerprints

LICENSE: FireSIGHT

If you no longer have use for a fingerprint, you can delete it from the system. Note that you must deactivate fingerprints before you can delete them.

To delete a fingerprint:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**.
The Custom Fingerprint page appears.

Name	Status	Type	Vendor	Product	Version	
Sample Client Fingerprint Sample Client Fingerprint	New	Client	Red Hat, Inc.	Fedora Core	6	 

2. If the fingerprints you want to delete are active, click the slider icon next to each one to deactivate it.
3. Click the delete icon () next to the fingerprint you want to delete.
4. Click **OK** to confirm that you want to delete the fingerprint.
The fingerprint is deleted.

Editing Fingerprints

LICENSE: FireSIGHT

After you create a fingerprint, you can view or edit it. This allows you to make changes and resubmit the fingerprint or add additional vulnerability mappings to it. You can modify fingerprints whether they are active or inactive, but depending on a fingerprint's state, the things that can be modified differ.

If a fingerprint is *inactive*, you can modify all elements of the fingerprint and resubmit it to the Defense Center. This includes all properties you specified when creating the fingerprint, such as fingerprint type, target IP addresses and ports, vulnerability mappings, and so on. When you edit an inactive fingerprint and submit it, it is resubmitted to the system and, if it is a client fingerprint, you must resend traffic to the appliance before activating it. Note that you can select only a single vulnerability mapping for an inactive fingerprint. After you activate the fingerprint, you can map additional operating systems and versions to its vulnerabilities list.

If a fingerprint is *active*, you can modify the fingerprint name, description, custom operating system display, and map additional vulnerabilities to it.

For more information, see the following sections:

- [Editing an Inactive Fingerprint](#) on page 1662
- [Editing an Active Fingerprint](#) on page 1663

Editing an Inactive Fingerprint

LICENSE: FireSIGHT

If a fingerprint is inactive, you can modify its properties and resubmit it to the system. This includes making changes such as the type of fingerprint to use, the target system to fingerprint, and so on.

To edit inactive fingerprints:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**.
The Custom Fingerprint page appears.
2. Click the edit icon () next to the fingerprint you want to edit.
The Edit Custom Fingerprint page appears.

3. Make changes to the fingerprint as necessary:
 - If you are modifying a client fingerprint, see [Fingerprinting Clients](#) on page 1650 for more information about the options you can configure.
 - If you are modifying a server fingerprint, see [Fingerprinting Servers](#) on page 1655 for more information about the options you can configure.
4. Click **Save** to resubmit the fingerprint.

IMPORTANT! If you modified a client fingerprint, remember to send traffic from the host to the appliance gathering the fingerprint.

Editing an Active Fingerprint

LICENSE: FireSIGHT

When a fingerprint is active, you can change its name, description, and display label. In addition, you can manage vulnerability mappings, including adding and deleting vulnerability mappings.

To edit active fingerprints:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Network Discovery**, then click **Custom Operating Systems**.
The Custom Fingerprint page appears.
2. Click the edit icon () next to the fingerprint you want to edit.
The Edit Custom Fingerprint Product Mappings page appears.
3. Modify the fingerprint name, description, and custom OS display, if necessary.
4. If you want to delete a vulnerability mapping, click **Delete** next to the mapping in the Pre-Defined OS Product Maps section of the page.
5. If you want to add additional operating systems for vulnerability mapping, select the **Product** and, if applicable, the **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** and then click **Add OS Definition**.
The vulnerability mapping is added to the Pre-Defined OS Product Maps list.
6. Click **Save** to save your changes.

Working with Application Detectors

LICENSE: FireSIGHT

When the Sourcefire 3D System analyzes IP traffic, it uses detectors to identify the commonly used applications on your network. You use the Detectors page

(**Policies > Application Detectors**) to customize the detection capability of the Sourcefire 3D System.



The page provides information about each detector, including:

- the name of the detector
- the protocol (TCP, UDP, or both) of traffic that the detector inspects
- whether the type of the detector is application protocol, client, web application, or internal detector
- for port-based application detectors, the port used by the application traffic
- details regarding the detected application, including the name, description, risk, business relevance, tags, and categories associated with the application detected by the detector
- the state (active or inactive) of the detector

The system uses only active detectors to analyze application traffic.

You may notice that the listed detectors have different properties. For example, you can view the settings for some detectors but not others. Similarly, you can delete some detectors but not others. This is because there are several different types of Sourcefire-provided detectors, as described in the following sections.

Sourcefire-Provided Internal Detectors

Internal detectors are application detectors that are only delivered with updates to the Sourcefire 3D System. Internal detectors detect client, web application, or application protocol traffic, depending on the detector, but they are categorized as internal detectors rather than one of the other types because they are built-in detectors and cannot be deactivated.

Internal detectors are always on; you cannot deactivate, delete, or otherwise configure them. Examples of internal detectors are the Built-in Amazon detector and the Built-in AppleTalk detector.

Sourcefire-Provided Client Detectors

Sourcefire-provided *client detectors*, which detect client traffic, are delivered via VDB updates but may also be provided with updates to the Sourcefire 3D System. These detectors may also be provided by Sourcefire Professional Services as an importable detector.

You can activate and deactivate client detectors according to the needs of your organization. VDB updates may also activate or deactivate client detectors. You can export a client detector only if you import it.

The Google Earth and Immuneset detectors are examples of client detectors.

Sourcefire-Provided Web Application Detectors

Sourcefire-provided *web application detectors*, which detect web applications in payloads of HTTP traffic, are delivered via VDB updates but may also be provided with updates to the Sourcefire 3D System.

You can activate and deactivate web application detectors according to the needs of your organization. VDB updates may activate or deactivate web application detectors. Examples of web application detectors are the Blackboard and LiveJournal detectors.

Sourcefire-Provided Application Protocol (Port) Detectors

Port-based application protocol detectors, provided by Sourcefire, are based on detection of network traffic on well-known ports. These detectors are delivered via VDB updates but may also be provided with updates to the Sourcefire 3D System or provided by Sourcefire Professional Services as an importable detector.

You can activate and deactivate application protocol detectors according to the needs of your organization. You can also view a detector definition to use it as the basis for a custom detector. VDB updates may activate or deactivate application protocol detectors.

The chargen and finger detectors are examples of port detectors.

Sourcefire-Provided Application Protocol (FireSIGHT) Detectors

FireSIGHT-based application protocol detectors, provided by Sourcefire, are based on detection of network traffic using FireSIGHT application fingerprints. These detectors are delivered via VDB updates but may also be provided with updates to the Sourcefire 3D System.

You can activate and deactivate application protocol detectors according to the needs of your organization. VDB updates may activate or deactivate Sourcefire-provided application protocol detectors. Examples of FireSIGHT-based application protocol detectors are the Jabber and Steam detectors.

Application Protocol (Pattern) Detectors

Pattern-based application detectors are based on detection of patterns in packets from network traffic. These detectors can be provided by Sourcefire Professional Services as an importable detector or created by you. This allows you to enhance the system's detection capabilities with new pattern-based detectors without updating the Sourcefire 3D System as a whole.

You can activate and deactivate application protocol detectors according to the needs of your organization.

You have full control over imported and user-defined detectors; you can activate, deactivate, edit, import, export, and delete them. An example of a pattern-based detector is a user-defined detector using a pattern in the packet header to detect traffic for a custom application.

Keep in mind that the detector list may change depending on the version of the Sourcefire 3D System and the VDB you have installed, as well as on any individual detectors you may have imported or created. You should carefully read the release notes for each Sourcefire 3D System update as well as the advisories for each VDB update for information on updated detectors.

For more information, see:

- [Understanding Application Detection](#) on page 1263
- [Creating a User-Defined Application Protocol Detector](#) on page 1666
- [Managing Detectors](#) on page 1673

Creating a User-Defined Application Protocol Detector

LICENSE: FireSIGHT

If you use custom applications on your network, you can create user-defined application protocol detectors that provide the system with the information it needs to identify those applications. You can base application protocol detection on the port or ports used by application traffic, patterns within the traffic, or on both ports and patterns.

For example, if you expect traffic for a custom application protocol to use port 1180, you can create an application protocol detector that detects traffic on that port. As another example, if you know that the header for any packet containing application protocol traffic has a string of `ApplicationName` in it, you can create a detector that registers the ASCII string of `ApplicationName` as a pattern to match.

You can only create user-defined application detectors for application protocols, not for clients or for web applications.

WARNING! When you create and activate a new application detector, a short pause in traffic flow and processing may occur on your managed devices, which may also cause a few packets to pass uninspected.

User-defined application protocol detectors must use either a port or a pattern match; you cannot create a detector that uses neither, even if you base the detector on an existing detector. You can also create a detector that uses both

criteria; this increases the likelihood of correctly identifying traffic for that application protocol.

TIP! If you have already created a detector on another Defense Center, you can export it and then import it onto this Defense Center. You can then edit the imported detector to suit your needs. You can export and import user-defined detectors as well as detectors provided by Sourcefire Professional Services. However, you **cannot** export or import any other type of Sourcefire-provided detectors. For more information, see [Importing and Exporting Objects](#) on page 2225.

To create a user-defined application protocol detector:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Application Detectors**.

The Detectors page appears.

2. Click **Create Detector**.

The Create Detector page appears.

Please enter a name

Enter a description

Detector Information

Author: admin

Application Protocol: Select an application protocol [v] [Add]

State: Inactive

Type: Application Protocol: FireSIGHT

Detection Criteria

Protocol: TCP/UDP [v]

Port(s): Enter a port or comma-separated list of ports. Example: 23, 200

Detection Patterns [Add]

Pattern String Ty...	Pattern String	Offset
There are no patterns. Click "Add" to add a pattern.		

Packet Captures [Add]

Packet Capture Name
There are no packet captures. Click "Add" to add a packet capture.

3. Provide basic detector information, such as the detector name and description.

See [Providing Basic Application Protocol Detector Information](#) on page 1668.

4. Optionally, create a user-defined application for the detector.

See [Creating a User-Defined Application](#) on page 1669.

5. Provide detection criteria, including the protocol of traffic the detector should inspect and the port that the traffic uses.
See [Specifying Detection Criteria for Application Protocol Detectors](#) on page 1670.
6. Optionally, configure the detector to inspect traffic for matches to one or more patterns that occurs in traffic for that application protocol.
See [Adding Detection Patterns to an Application Protocol Detector](#) on page 1671.
7. Optionally, test the new detector against the contents of one or more PCAP files.
See [Testing an Application Protocol Detector Against Packet Captures](#) on page 1673.
8. Click **Save**.
The application protocol detector is saved.

IMPORTANT! You must activate the detector before the system can use it to analyze application protocol traffic. For more information, see [Activating and Deactivating Detectors](#) on page 1678. Note that if you include the application in an access control rule, the detector is automatically activated and cannot be deactivated while in use.

Providing Basic Application Protocol Detector Information

LICENSE: FireSIGHT

You must give each user-defined application protocol detector a name, as well as identify the application protocol you want to detect. Optionally, you can provide a brief description of the detector.

In addition to the information you provide, the Defense Center indicates whether the detector is active or inactive, and whether the detector is a port or pattern detector. If a detector identifies application protocol traffic by port and pattern, the Sourcefire 3D System considers it a pattern detector.



Please enter a name

Enter a description

Detector Information

Author: admin

Application Protocol:

State: Inactive

Type: Application Protocol: FireSIGHT

If you are editing an existing detector, the Defense Center also displays the detector's author. If you created a user-defined application protocol detector, you are the author. You are also the author for any detector that you import or that you edit and save.

To provide basic application protocol detector information:

ACCESS: Admin/Discovery Admin

1. On the Create Detector page, in the **Please enter a name** field, type a name for the detector.

Detector names must be unique within the protocol for the traffic you are inspecting. That is, you can create a TCP detector and a UDP detector with the same name, but you cannot create two TCP detectors with the same name.

2. Identify the application protocol you want to detect. You have the following options:
 - If you are creating a detector for an existing application protocol (for example, if you want to detect a particular application protocol on a non-standard port), select the application protocol from the **Application Protocol** drop-down list. Continue with the procedure in [Specifying Detection Criteria for Application Protocol Detectors](#) on page 1670.
 - If you are creating a detector for a custom application, continue with the procedure in the next section, [Creating a User-Defined Application](#).

Creating a User-Defined Application

LICENSE: FireSIGHT

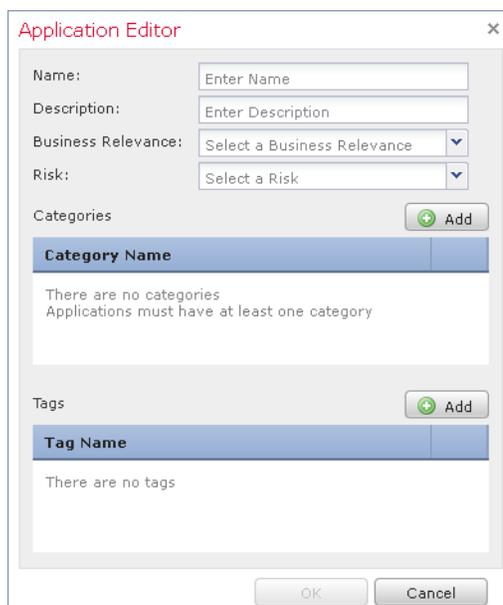
You can create a user-defined application to identify a custom application on your network. You can also create custom categories and custom tags to describe the application. Applications, categories, and tags created here are available in access control rules and in the application filter object manager as well.

For more information on application detection, including a discussion of application protocols and the categories, tags, risk levels, and business relevance used to describe them, see [Understanding Application Detection](#) on page 1263.

To create a user-defined application:

ACCESS: Admin/Discovery Admin

1. On the Create Detector page, click **Add**.
The Application Editor pop-up window appears.



2. Type a **Name** for the custom application.
3. Type a **Description** for the custom application.
4. Select a **Business Relevance**.
5. Select a **Risk**.
6. Click **Add** next to Categories to add a category and type a new category name or select an existing category from the **Categories** drop-down list.
7. Optionally, click **Add** next to Tags to add a tag and type a new tag name or select an existing tag from the **Tags** drop-down list.
Click **OK** to return to the Create Detector page.
8. Continue with the procedure in the next section, [Specifying Detection Criteria for Application Protocol Detectors](#).

Specifying Detection Criteria for Application Protocol Detectors

LICENSE: FireSIGHT

When creating a user-defined application protocol detector, you must specify the protocol of traffic (TCP, UDP, or both) the detector should inspect. Optionally, you can specify a port that the traffic uses.

Note that if you do not specify a port, you must configure the detector to inspect traffic for matches to one or more patterns, as described in [Adding Detection Patterns to an Application Protocol Detector](#) on page 1671.

To specify detection criteria for an application protocol detector:

ACCESS: Admin/Discovery Admin

1. On the Create Detector page, from the **Protocol** drop-down list, select the protocol for traffic the detector should inspect.
Detectors can inspect TCP, UDP, or TCP and UDP traffic.
2. Optionally, to identify application protocol traffic based on the port it uses, type a port from 1 to 65535 in the **Port(s)** field. To use multiple ports, separate them by commas.
3. You have the following options:
 - If you want to configure the application protocol detector to inspect traffic for matches to one or more patterns that occurs in traffic for that application protocol, continue with the procedure in the next section, [Adding Detection Patterns to an Application Protocol Detector](#).
 - If you want to test the new detector against the contents of one or more PCAP files, skip to [Testing an Application Protocol Detector Against Packet Captures](#) on page 1673.
 - If you are done creating the detector, click **Save**.

The application protocol detector is saved.

Note that you must activate the detector before the system can use it to analyze application protocol traffic. For more information, see [Activating and Deactivating Detectors](#) on page 1678.

Adding Detection Patterns to an Application Protocol Detector

LICENSE: FireSIGHT

If you know that the header for any packet containing application protocol traffic contains a particular pattern string, you can configure a user-defined application protocol detector to search for that pattern.

Application protocol detectors can search for ASCII or hexadecimal patterns, using any offset. You can also configure detectors to search for multiple patterns; in that case the application protocol traffic must match all of the patterns for the detector to positively identify the application protocol.

Note that if you do not specify a pattern, you must configure the detector to inspect traffic that uses one or more ports, as described in [Specifying Detection Criteria for Application Protocol Detectors](#) on page 1670.

To add a detection pattern to an application protocol detector:

ACCESS: Admin/Discovery Admin

1. On the Create Detector page, in the Detection Patterns section, click **Add**. The Add Pattern pop-up window appears.



2. Specify the pattern type you want to detect: **Ascii** or **Hex**.
3. Type a string of the type you specified in the **Pattern String** field.
4. Optionally, specify where in a packet the system should begin searching for the pattern; this is called the offset.

Type the offset (in bytes from the beginning of the packet payload) in the **Offset** field.

Because packet payloads start at byte 0, calculate the offset by subtracting 1 from the number of bytes you want to move forward from the beginning of the packet payload. For example, to look for the pattern in the fifth bit of the packet, type 4 in the **Offset** field.

5. Optionally, repeat steps 1 to 4 to add additional patterns.

TIP! To delete a pattern, click the delete icon () next to the pattern you want to delete.

6. You have the following options:
 - If you want to test the new detector against the contents of one or more PCAP files, continue with the procedure in the next section, [Testing an Application Protocol Detector Against Packet Captures](#).
 - If you are done creating the detector, click **Save**.

The application protocol detector is saved.

Note that you must activate the detector before the system can use it to analyze application protocol traffic. For more information, see [Activating and Deactivating Detectors](#) on page 1678.

Testing an Application Protocol Detector Against Packet Captures

LICENSE: FireSIGHT

If you have a packet capture (PCAP) file that contains packets with traffic from the application protocol you want to detect, you can test a user-defined application protocol detector against that PCAP file. Note that PCAP files must be 32KB or smaller; if you try to test your detector against a larger PCAP file, the Defense Center automatically truncates it.

To test an application protocol detector against a PCAP file:

ACCESS: Admin/Discovery Admin

1. On the Create Detector page, in the Packet Captures section, click **Add**.
A pop-up window appears.



2. Browse to the PCAP file and click **OK**.
The PCAP file appears in the Packet Captures file list.
3. To test your detector against the contents of the PCAP file, click the evaluate icon next to the PCAP file.
A message appears, indicating whether the test succeeded.
4. Optionally, repeat steps 1 to 3 to test the detector against additional PCAP files.

TIP! To delete a PCAP file, click the delete icon (🗑️) next to the file you want to delete.

5. To save the detector, click **Save**.

IMPORTANT! You must activate the detector before the system can use it to analyze application protocol traffic. For more information, see [Activating and Deactivating Detectors](#) on page 1678.

Managing Detectors

LICENSE: FireSIGHT

You view and manage detectors on the Detectors page.

From the Detectors page, you can:

- view details about the application the detector identifies
- sort, filter, and browse the list of detectors
- view a list of the Sourcefire-provided internal detectors
- view the properties of the Sourcefire-provided application protocol port detectors, and optionally save copies as new, user-defined detectors that you can modify
- create, modify, delete, and export user-defined application protocol detectors
- delete and export any application protocol detectors you individually imported
- activate and deactivate user-defined, imported, or Sourcefire-provided web application, client, and application protocol detectors

Note that you cannot modify or delete internal or Sourcefire-provided application protocol, client, or web application detectors and cannot deactivate internal detectors.

For more information, see:

- [Viewing Detector Details](#) on page 1674
- [Sorting the Detector List](#) on page 1675
- [Filtering the Detector List](#) on page 1675
- [Navigating to Other Detector Pages](#) on page 1678
- [Activating and Deactivating Detectors](#) on page 1678
- [Modifying Application Detectors](#) on page 1679
- [Deleting Detectors](#) on page 1680

Viewing Detector Details

LICENSE: FireSIGHT

You can view more detail on a detector from the application detectors list.

To view application detector details:

ACCESS: Admin/Discovery Admin

- ▶ Click the information icon () in the Details column.
The information pop-up window for the detector appears.



For more information on risk, business relevance, tags, and categories, see [Understanding Application Detection](#) on page 1263.

Sorting the Detector List

LICENSE: FireSIGHT

By default, the Detectors page lists detectors alphabetically by name. An up () or down arrow next to a column heading indicates that the page is sorted by that column in that direction.

To sort detectors:

ACCESS: Admin/Discovery Admin

- ▶ On the Detectors page, click the appropriate column heading.
The detectors are sorted in the direction indicated by the arrow that appears on the column heading. To sort in the opposite direction, click the heading again.

Filtering the Detector List

LICENSE: FireSIGHT

You can filter the detectors you display on the Detectors page by a single criterion, or a combination of multiple criteria. The filter you construct is shown at the top of the page. You can use several filter groups, separately or in combination, to filter the list of detectors.

Name

Finds detectors with names or descriptions containing the string you type. Strings can contain any alphanumeric or special character.

Custom Filter

Finds detectors matching a custom application filter created on the object management page. For more information, see [Working with Application Filters](#) on page 176.

Author

Finds detectors according to who created the detector. You can filter detectors by:

- any individual user who has created or imported a detector
- **Sourcefire**, which represents all Sourcefire-provided detectors *except* individually imported add-on detectors; you are the author for any detector that you import
- **Any User**, which represents all detectors not provided by Sourcefire

State

Finds detectors according to their state, that is, **Active** or **Inactive**. For more information, see [Activating and Deactivating Detectors](#) on page 1678.

Type

Finds detectors according to the detector type: **Application Protocol**, **Web Application**, **Client**, or **Internal Detector**.

Application protocol detectors have three subtypes you can use to further filter detectors:

- **Port** application protocol detectors include the Sourcefire-provided well-known port detectors, as well as any port-based user-defined application detectors.
- **Pattern** application protocol detectors include pattern-based or port-and-pattern-based user-defined application detectors.
- **FireSIGHT** application protocol detectors are application protocol fingerprint detectors provided by Sourcefire that can be activated and deactivated.

For more information on detector types, see [Working with Application Detectors](#) on page 1663.

Protocol

Finds detectors according to which traffic protocol the detector inspects. Detectors can inspect TCP, UDP, or TCP and UDP traffic.

Category

Finds detectors according to the categories assigned to the application they detect.

Tag

Finds detectors according to the tags assigned to the application they detect.

Risk

Finds detectors according to the risks assigned to the application they detect: **Very High, High, Medium, Low, and Very Low.**

Business Relevance

Finds detectors according to the business relevance assigned to the application they detect: **Very High, High, Medium, Low, and Very Low.**

To apply a filter:

ACCESS: Admin/Discovery Admin

1. On the Detectors page, expand the filter group you want to use to filter the detectors.
2. Type the name or select the specific filter you want to use. To select all filters in a group, right-click the group name and select **Check All**.
3. Optionally, if the filter you are using has subfilters, select the subfilter to further filter the detectors.

To remove a filter:

ACCESS: Admin/Discovery Admin

- ▶ Click the remove icon (✕) in the name of the filter in the **Filters** field or disable the filter in the filter list. To remove all filters in a group, right-click the group name and select **Uncheck All**.
The filter is removed and the results update.

To remove all filters:

ACCESS: Admin/Discovery Admin

- ▶ Click **Clear all** next to the list of filters applied to the detectors.

Navigating to Other Detector Pages

LICENSE: FireSIGHT

The Detectors page displays 25 detectors at a time. The following table explains how to view additional pages of detectors using the navigation links at the bottom of the page.

ACCESS: Admin/Discovery Admin

<< Page 1 of 30 >>

Navigating Detector Pages

To...	YOU CAN...
view the next page	click the right arrow icon (>).
view the previous page	click the left arrow icon (<).
view a different page	type the page number and press Enter.
jump to the last page	click the right end arrow icon (>).
jump to the first page	click the left end arrow icon (<).

Activating and Deactivating Detectors

LICENSE: FireSIGHT

You must activate a detector before you can use it to analyze network traffic. By default, all Sourcefire-provided detectors are activated.

You can activate multiple application detectors for each port to supplement the system's detection capability.

When you include an application in an access control rule in a policy and that policy is applied, if there is no active detector for that application, one or more detectors automatically activate. Similarly, while an application is in use in an applied policy, you cannot deactivate a detector if deactivating leaves no active detectors for that application.

WARNING! When you activate or deactivate an existing detector, a short pause in traffic flow and processing may occur on your managed devices, which may also cause a few packets to pass uninspected.

TIP! For improved performance, deactivate any application protocol, client, or web application detectors you are not interested in.

To activate or deactivate a detector:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Application Detectors.**

The Detectors page appears.

2. Locate the detector you want to activate or deactivate.

If the detector you want to activate or deactivate is not on the first page, you can find it by paging through the detector list or applying one or more filters. For more information, see [Managing Detectors](#) on page 1673.

3. You have the following options:

- To **activate** a detector, so that the system will use it when analyzing network traffic, click the deactivated slider () next to the detector.
- To **deactivate** a detector so that the system will not use it when analyzing network traffic, click the activated slider () next to the detector.

Note that some application detectors are required by other detectors. If you deactivate one of these detectors, a warning appears to indicate that the detectors that depend on it are also disabled.

Modifying Application Detectors

LICENSE: FireSIGHT

Use the following procedure to modify user-defined application detectors.

To modify an application detector:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Applications.**

The Detectors page appears.

2. Locate the detector you want to modify.

If the detector you want to modify is not on the first page, you can find it by paging through the detector list or applying one or more filters. For more information, see [Managing Detectors](#) on page 1673.

3. To modify a user-defined detector, click **Edit next to the detector you want to modify.**

The Edit Application Detector page appears.

4. Make changes to the detector.

See [Creating a User-Defined Application Protocol Detector](#) on page 1666 for information on the various configurations you can change.

5. You have the following options:
 - If you are modifying an inactive user-defined detector, either click **Save** to save your changes, or click **Save as New** to save the detector as a new, inactive user-defined detector.
 - If you are modifying an active user-defined detector, either click **Save and Reactivate** to save your changes and immediately start using the modified detector, or click **Save as New** to save the detector as a new, inactive user-defined detector.

IMPORTANT! The system only uses applications with active detectors to analyze application traffic. For more information, see [Activating and Deactivating Detectors](#) on page 1678.

Deleting Detectors

LICENSE: FireSIGHT

Use the following procedure to delete a detector. You can delete user-defined detectors as well as individually imported add-on detectors provided by Sourcefire Professional Services. You cannot delete any of the other Sourcefire-provided detectors, though you can deactivate many of them.

IMPORTANT! While a detector is in use in an applied policy, you cannot deactivate or delete the detector.

To delete a detector:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Application Detectors**.
The Detectors page appears.
2. Select the check box next to the detector you want to delete and click **Delete**.
If the detector you want to delete is not on the first page, you can find it by paging through the detector list or applying one or more filters. For more information, see [Managing Detectors](#) on page 1673.
3. Click **OK** to confirm that you want to delete the detector.
The detector is deleted.

Importing Host Input Data

LICENSE: FireSIGHT

If your organization has the capability to write scripts or create command line import files to import network map data from third parties, you can import data to augment the information in the network map. You can also use the host input

feature by modifying operating system or application identities or deleting application protocols, protocols, host attributes, or clients using the web interface.

The system may reconcile data from multiple sources to determine the current identity of an operating system or application. For more information on how the system does this, see [Understanding Current Identities](#) on page 1646.

Note that all data except third-party vulnerabilities is discarded when the affected host is removed from the network map. For more information on setting up scripts or import files, see the *Sourcefire 3D System Host Input API Guide*.

To include imported data in impact correlations, you must map the data to the operating system and application definitions in the database. For more information, see the following sections:

- [Enabling the Use of Third-Party Data](#) on page 1681
- [Managing Third-Party Product Mappings](#) on page 1682
- [Mapping Third-Party Vulnerabilities](#) on page 1687
- [Managing Custom Product Mappings](#) on page 1688

Enabling the Use of Third-Party Data

LICENSE: FireSIGHT

You can import network map data from third-party systems on your network. However, to enable features where intrusion and discovery data are used together, such as FireSIGHT recommendations, adaptive profiles, or impact assessment, you should map as many elements of it as possible to corresponding definitions. Consider the following requirements for using third-party data:

- If you have a third-party system that has specific data on your network assets, you can import that data using the host input feature. However, because third parties may name the products differently, you must map the third-party vendor, product, and versions to the corresponding Sourcefire product definition. After you map the products, you must enable vulnerability mappings for impact assessment in the system policy to allow impact correlation. For versionless or vendorless application protocols, you need to map vulnerabilities for the application protocols in the system policy. For more information, see [Mapping Third-Party Products](#) on page 1682.
- If you import patch information from a third party and you want to mark all vulnerabilities fixed by that patch as invalid, you must map the third-party fix name to a fix definition in the database. All vulnerabilities addressed by the fix will then be removed from hosts where you add that fix. For more information, see [Mapping Third-Party Product Fixes](#) on page 1685.

- If you import operating system and application protocol vulnerabilities from a third party and you want to use them for impact correlation, you must map the third-party vulnerability identification string to vulnerabilities in the database. Note that although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot import and map third-party client vulnerabilities. After the vulnerabilities are mapped, you must enable third-party vulnerability mappings for impact assessment in the system policy. For more information, see [Mapping Third-Party Vulnerabilities](#) on page 1687. To cause application protocols without vendor or version information to map to vulnerabilities, an administrative user must also map vulnerabilities for the applications in the system policy. For more information, see [Mapping Vulnerabilities for Servers](#) on page 1996.
- If you import application data and you want to use that data for impact correlation, you must map the vendor string for each application protocol to the corresponding Sourcefire application protocol definition. For more information, see [Managing Custom Product Mappings](#) on page 1688.

Managing Third-Party Product Mappings

LICENSE: FireSIGHT

When you add data from third parties to the network map through the user input feature, you must map the vendor, product, and version names used by the third party to the Sourcefire product definitions. Mapping the products to Sourcefire definitions assigns vulnerabilities based on those definitions.

Similarly, if you are importing patch information from a third party, such as a patch management product, you must map the name for the fix to the appropriate vendor and product and the corresponding fix in the database.

For more information, see the following sections:

- [Mapping Third-Party Products](#) on page 1682
- [Mapping Third-Party Product Fixes](#) on page 1685

Mapping Third-Party Products

LICENSE: FireSIGHT

If you import data from a third party, you must map the Sourcefire product to the third-party name to assign vulnerabilities and perform impact correlation using that data. Mapping the product associates Sourcefire vulnerability information with the third-party product name, which allows the system to perform impact correlation using that data.

If you import data using the host input import feature, you can also use the `AddScanResult` function to map third-party products to operating system and application vulnerabilities during the import.

As an example, if you import data from a third party that lists Apache Tomcat as an application and you know it is version 6 of that product, you could add a third-party map where **Vendor Name** is set to `Apache`, **Product Name** is set to

Tomcat, **Apache** is selected from the **Vendor** drop-down list, **Tomcat** is selected from the **Product** drop-down list, and **6** is selected from the **Version** drop-down list. That mapping would cause any vulnerabilities for Apache Tomcat 6 to be assigned to hosts with an application listing for Apache Tomcat.

Note that for versionless or vendorless applications, you must map vulnerabilities for the application types in the system policy. For more information, see [Mapping Vulnerabilities for Servers](#) on page 1996. Note that although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot import and map third-party client vulnerabilities.

TIP! If you have already created a third-party mapping on another Defense Center, you can export it and then import it onto this Defense Center. You can then edit the imported mapping to suit your needs. For more information, see [Importing and Exporting Objects](#) on page 2225.

To map a third-party product to a Sourcefire product definition:

ACCESS: Admin

1. Select **Policies > Application Detectors**, then click **User Third-Party Mappings**.

The User Third-Party Mappings page appears.

2. You have two choices:

- To edit an existing map set, click **Edit** next to the map set.
- To create a new map set, click **Create Product Map Set**.

The Edit Third-Party Product Mappings page appears.

Mapping Set Name	<input type="text"/>	
Description	<input type="text"/>	
Product Maps + Add Product Map		
Vendor String	Product String	Version String
Fix Maps + Add Fix Map		
Fix String		
Save Cancel		

3. Type a name for the mapping set in the **Mapping Set Name** field.
4. Type a description in the **Description** field.

5. You have two choices:
 - To map a third-party product, click **Add Product Map**.
 - To edit an existing third-party product map, click **Edit** next to the map set.

The Add Product Map page appears.

Add Product Map
Specify the third-party strings and select the VDB product entry that the strings should map to.

Third-Party Strings

Vendor String	<input type="text"/>
Product String	<input type="text"/>
Version String	<input type="text"/>

Product Mapping

Vendor	Select a Vendor
Product	
Major	
Minor	
Revision	
Build	
Patch	
Extension	

6. Type the vendor string used by the third-party product in the **Vendor String** field.
7. Type the product string used by the third-party product in the **Product String** field.
8. Type the version string used by the third-party product in the **Version String** field.
9. In the Product Mappings section, select the operating system, product, and versions you want to use for vulnerability mapping from the following lists (if applicable):
 - **Vendor**
 - **Product**
 - **Major Version**
 - **Minor Version**
 - **Revision Version**
 - **Build**

- **Patch**
- **Extension**

For example, if you want a host running a product whose name consists of third-party strings to use the vulnerabilities from Red Hat Linux 9, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

10. Click **Save**.

Mapping Third-Party Product Fixes

LICENSE: FireSIGHT

If you map a fix name to a particular set of fixes in the database, you can then import data from a third-party patch management application and apply the fix to a set of hosts. When the fix name is imported to a host, the system marks all vulnerabilities addressed by the fix as invalid for that host.

To map third-party fixes to Sourcefire fix definitions:

ACCESS: Admin

1. Select **Policies > Application Detectors**, then click **User Third-Party Mappings**. The User Third-Party Mappings page appears.
2. You have two choices:
 - To edit an existing map set, click **Edit** next to the map set.
 - To create a new map set, click **Create Product Map Set**.

The Edit Third-Party Product Mappings page appears.

Mapping Set Name	<input type="text"/>	
Description	<input type="text"/>	
Product Maps + Add Product Map		
Vendor String	Product String	Version String
Fix Maps + Add Fix Map		
Fix String		
<input type="button" value="Save"/> <input type="button" value="Cancel"/>		

3. Type a name for the mapping set in the **Mapping Set Name** field.
4. Type a description in the **Description** field.

5. You have two choices:
 - To map a third-party product, click **Add Fix Map**.
 - To edit an existing third-party product map, click **Edit** next to it.

The Add Fix Map page appears.

Add Fix Map
Specify the third-party Fix Name (patch) and select the VDB product and fix entry that the string maps to.

Third-Party Fix Name

Fix Name

Product Mapping

Vendor	Select a Vendor
Product	
Major	
Minor	
Revision	
Build	
Patch	
Extension	

6. Type the name of the fix you want to map in the **Third-Party Fix Name** field.
7. In the Product Mappings section, select the operating system, product, and versions you want to use for fix mapping from the following lists (if applicable):
 - **Vendor**
 - **Product**
 - **Major Version**
 - **Minor Version**
 - **Revision Version**
 - **Build**
 - **Patch**
 - **Extension**

For example, if you want your mapping to assign the selected fixes from Red Hat Linux 9 to hosts where the patch is applied, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

8. Click **Save** to save the fix map.

Mapping Third-Party Vulnerabilities

LICENSE: FireSIGHT

To add vulnerability information from a third party to the VDB, you must map the third-party identification string for each imported vulnerability to any existing Sourcefire, Bugtraq, or Snort ID. After you create a mapping for the vulnerability, the mapping works for all vulnerabilities imported to hosts in your network map and allows impact correlation for those vulnerabilities.

Note that you must also enable impact correlation for third-party vulnerabilities to allow correlation to occur. For more information, see [Enabling Vulnerability Impact Assessment Mappings](#) on page 1291. For versionless or vendorless applications, you must also map vulnerabilities for the application types in the system policy. For more information, see [Mapping Vulnerabilities for Servers](#) on page 1996.

Also, although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot use third-party client vulnerabilities for impact assessment.

TIP! If you have already created a third-party mapping on another Defense Center, you can export it and then import it onto this Defense Center. You can then edit the imported mapping to suit your needs. For more information, see [Importing and Exporting Objects](#) on page 2225.

To map a third-party vulnerability to an existing vulnerability:

ACCESS: Admin

1. Select **Policies > Application Detectors**, then click **User Third-Party Mappings**.
The User Third-Party Mappings page appears.
2. You have two choices:
 - To edit an existing vulnerability set, click **Edit** next to the vulnerability set.
 - To create a new vulnerability set, click **Create Vulnerability Map Set**.

The Edit Third-Party Vulnerability Mappings page appears.

Vulnerability Set Name	<input type="text"/>
Description	<input type="text"/>

Vulnerability Maps ➕ Add Vulnerability Map

3. Click **Add Vulnerability Map**.

The Add Vulnerability Map pop-up window appears.

Add Vulnerability Map
Specify the third-party vulnerability name/ID and the applicable Sourcefire and Bugtraq references as comma separated lists of values.

Vulnerability

Vulnerability ID	<input type="text"/>
Vulnerability Description	<input type="text"/>
Snort Vulnerability ID Mappings	<input type="text"/>
Sourcefire Vulnerability ID Mappings	<input type="text"/>
Bugtraq Vulnerability ID Mappings	<input type="text"/>

4. Type the third-party identification for the vulnerability in the **Vulnerability ID** field.
5. Type a description in the **Vulnerability Description** field.
6. Optionally, enter a Signature ID in the **Snort Vulnerability ID Mappings** field.
7. Optionally, enter an Sourcefire vulnerability ID in the **Sourcefire Vulnerability ID Mappings** field.
8. Optionally, enter a Bugtraq identification number in the **Bugtraq Vulnerability ID Mappings** field.
9. Click **Add**.

Managing Custom Product Mappings

LICENSE: FireSIGHT

You can use product mappings to ensure that servers input by a third party are associated with the appropriate Sourcefire definitions. After you define and activate the product mapping, all servers or clients on hosts in your network map that have the mapped vendor strings use the custom product mappings. For this reason, you may want to map vulnerabilities for all servers in the network map with a particular vendor string instead of explicitly setting the vendor, product, and version for the server.

For more information, see the following:

- [Creating Custom Product Mappings](#) on page 1689
- [Editing Custom Product Mapping Lists](#) on page 1690
- [Managing Custom Product Mapping Activation State](#) on page 1691

Creating Custom Product Mappings

LICENSE: FireSIGHT

If the system cannot map a server in the network map to a vendor and product in the VDB, you can manually create the mapping for the system to use when identifying servers. When you activate a custom product mapping, the system maps vulnerabilities for the selected vendor and product to all servers in the network map where that vendor string occurs.

IMPORTANT! Custom product mappings apply to all occurrences of an application protocol, regardless of the source of the application data (such as Nmap, the host input feature, or the Sourcefire 3D System itself). However, if third-party vulnerability mappings for data imported using the host input feature conflicts with the mappings you set through a custom product mapping, the third-party vulnerability mapping overrides the custom product mapping and uses the third-party vulnerability mapping settings when the input occurs. For more information, see [Mapping Third-Party Vulnerabilities](#) on page 1687.

You create lists of product mappings and then enable or disable use of several mappings at once by activating or deactivating each list. When you select a vendor to map to, the system updates the list of products to include only those made by that vendor.

After you create a custom product mapping, you must activate the custom product mapping list. After you activate a list of custom product mappings, the system updates all servers with occurrences of the specified vendor strings. For data imported through the host input feature, vulnerabilities update unless you have already explicitly set the product mappings for this server.

If, for example, your company modifies the banner for your Apache Tomcat web servers to read **Internal web Server**, you can map the vendor string **Internal web Server** to the vendor **Apache** and the product **Tomcat**, then activate the list containing that mapping, all hosts where a server labelled **Internal web Server** occurs have the vulnerabilities for Apache Tomcat in the database.

TIP! You can use this feature to map vulnerabilities to local intrusion rules by mapping the SID for the rule to another vulnerability.

To create a custom product mapping:

ACCESS: Admin

1. Select **Policies > Application Detectors**, and click **Custom Product Mappings**. The Custom Product Mappings page appears.

2. Click **Create Custom Product Mapping List**.

The Edit Custom Product Mappings List page appears.

The screenshot shows a web interface for editing a custom product mapping list. At the top, there is a text input field labeled 'Custom Product Mapping List Name'. Below this is a table with three columns: 'Vendor String', 'Vendor', and 'Product'. To the right of the table is a button labeled '+ Add Vendor String'. At the bottom of the page are two buttons: 'Save' and 'Cancel'.

3. Type a name in the **Custom Product Mapping List Name** field.

4. Click **Add Vendor String**.

The Add Vendor String pop-up window appears.

The screenshot shows a pop-up window titled 'Add Vendor String'. It contains a text input field for 'Vendor String', a dropdown menu for 'Vendor' with the text 'Select a Vendor', and a dropdown menu for 'Product'. At the bottom are two buttons: 'Add' and 'Cancel'.

5. In the **Vendor String** field, type the vendor string that identifies the applications that should map to the selected vendor and product values.
6. Select the vendor you want to map to from the **Vendor** drop-down list.
7. Select the product you want to map to from the **Product** drop-down list.
8. Click **Add** to add the mapped vendor string to the list.
9. Optionally, repeat steps 4 to 8 as needed to add additional vendor string mappings to the list.
10. When you finish, click **Save**.

The Custom Product Mappings page appears again, with the list you added.

Editing Custom Product Mapping Lists

LICENSE: FireSIGHT

You can modify existing custom product mapping lists by adding or removing vendor strings or changing the list name.

To edit a custom product mapping:

ACCESS: Admin

1. Select **Policies > Application Detectors**, then click **Custom Product Mappings**.

The Custom Product Mappings page appears.

2. Click the edit icon (✎) next to the product mapping list to edit.
The Edit Custom Product Mappings List page appears.

Custom Product Mapping List Name		
Vendor String	Vendor	Product
+ Add Vendor String		
Save Cancel		

3. Make changes to the list as needed. For more information, see [Creating Custom Product Mappings](#) on page 1689.
4. When you finish, click **Save**.
The Custom Product Mappings page appears, with the list you updated.

Managing Custom Product Mapping Activation State

LICENSE: FireSIGHT

You can enable or disable use of an entire list of custom product mappings at once. After you activate a custom product mapping list, each mapping on that list applies to all applications on hosts in the network map with the specified vendor string, whether detected by managed devices or imported through the host input feature.

To activate or deactivate a custom product mapping list:

ACCESS: Admin

1. Select **Policies > Application Detectors**, then click **Custom Product Mappings**.
The Custom Product Mappings page appears.
2. Modify the state of custom product mapping lists:
 - To enable use of a custom product mapping list, click **Activate**.
 - To disable use of a custom product mapping list, click **Deactivate**.

CHAPTER 41

CONFIGURING ACTIVE SCANNING

The Sourcefire 3D System builds a network map through passive analysis of traffic on your network. However, you may sometimes need to actively scan a host to determine information about that host. For example, if a host has a server running on an open port but the server has not received or sent traffic during the time that the system has been monitoring your network, the system does not add information about that server to the network map. If you directly scan that host using an active scanner, however, you can detect the presence of the server.

When you actively scan a host, you send packets in an attempt to obtain information about the host. The Sourcefire 3D System integrates with Nmap™ 6.01, an open source active scanner for network exploration and security auditing that can be used to detect operating systems and servers running on a host. With an Nmap scan, you can check for detailed information about the operating system and servers running on the host and refine the system's vulnerability reporting based on those results.

IMPORTANT! Some scanning options (such as portscans) may place a significant load on networks with low bandwidths. You should always schedule scans like these to run during periods of low network use.

For more information, see the following sections:

- [Understanding Nmap Scans](#) on page 1693
- [Setting up Nmap Scans](#) on page 1702
- [Managing Nmap Scanning](#) on page 1710

- [Managing Scan Targets](#) on page 1714
- [Working with Active Scan Results](#) on page 1716

Understanding Nmap Scans

LICENSE: FireSIGHT

Nmap allows you to actively scan ports on hosts on your network to determine operating system and server data for the hosts, which allows you to enhance your network map and fine-tune the accuracy of the vulnerabilities mapped to scanned hosts. Note that a host must exist in the network map before Nmap can append its results to the host profile. You can also view scan results in a results file.

When you scan a host using Nmap, servers on previously undetected open ports are added to the Servers list in the host profile for that host. The host profile lists any servers detected on filtered or closed TCP ports or on UDP ports in the Scan Results section. By default, Nmap scans more than 1660 TCP ports.

Nmap compares the results of the scan to over 1500 known operating system fingerprints to determine the operating system and assigns scores to each. The operating system assigned to the host is the operating system fingerprint with the highest score.

If the system recognizes a server identified in an Nmap scan and has a corresponding server definition, the system maps vulnerabilities for that server to the host. The system maps the names Nmap uses for servers to the corresponding Sourcefire server definitions, and then uses the vulnerabilities mapped to each server in the system. Similarly, the system maps Nmap operating system names to Sourcefire operating system definitions. When Nmap detects an operating system for a host, the system assigns vulnerabilities from the corresponding Sourcefire operating system definition to the host.

For more information the underlying Nmap technology used to scan, refer to the Nmap documentation at <http://insecure.org>.

For more information on Nmap on your Sourcefire appliance, see the following topics:

- [Understanding Nmap Remediations](#) on page 1693
- [Creating an Nmap Scanning Strategy](#) on page 1697
- [Sample Nmap Scanning Profiles](#) on page 1699

Understanding Nmap Remediations

LICENSE: FireSIGHT

You can define the settings for an Nmap scan by creating an Nmap remediation. An Nmap remediation can be used as a response in a correlation policy, run on demand, or scheduled to run at a specific time. In order for the results of an Nmap scan to appear in the network map, the scanned host must already exist in the network map.

Note that Nmap-supplied server and operating system data remain static until you run another Nmap scan. If you plan to scan a host for operating system and server data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up-to-date. For more information, see [Automating Nmap Scans](#) on page 1936. Also note that if the host is deleted from the network map, any Nmap scan results for that host are discarded.

For more information about Nmap functionality, refer to the Nmap documentation at <http://insecure.org>. The following table explains the options configurable in Nmap remediations on a Sourcefire 3D System.

Nmap Remediation Options

OPTION	DESCRIPTION	CORRESPONDING NMAP OPTION
Scan Which Address(es) From Event?	When you use an Nmap scan as a response to a correlation rule, select an option to control which address in the event is scanned, that of the source host, the destination host, or both.	N/A
Scan Types	<p>Select how Nmap scans ports:</p> <ul style="list-style-type: none"> • The TCP Syn scan connects quickly to thousand of ports without using a complete TCP handshake. This options allows you to scan quickly in stealth mode on hosts where the admin account has raw packet access or where IPv6 is not running, by initiating TCP connections but not completing them. If a host acknowledges the Syn packet sent in a TCP Syn scan, Nmap resets the connection. • The TCP Connect scan uses the connect() system call to open connections through the operating system on the host. You can use the TCP Connect scan if the admin user on your Defense Center or managed device does not have raw packet privileges on a host or you are scanning IPv6 networks. In other words, use this option in situations where the TCP Syn scan cannot be used. • The TCP ACK scan sends an ACK packet to check whether ports are filtered or unfiltered. • The TCP Window scan works in the same way as a TCP ACK scan but can also determine whether a port is open or closed. • The TCP Maimon scan identifies BSD-derived systems using a FIN/ACK probe. 	<p>TCP Syn: -sS TCP Connect: -sT TCP ACK: -sA TCP Window: -sW TCP Maimon: -sM</p>
Scan for UDP ports	Enable to scan UDP ports in addition to TCP ports. Note that scanning UDP ports may be time-consuming, so avoid using this option if you want to scan quickly.	-sU

Nmap Remediation Options (Continued)

OPTION	DESCRIPTION	CORRESPONDING NMAP OPTION
Use Port From Event	If you plan to use the remediation as a response in a correlation policy, enable to cause the remediation to scan only the port specified in the event that triggers the correlation response. TIP! You can also control whether Nmap collects information about operating system and server information. Enable the Use Port From Event option to scan the port associated with the new server.	N/A
Scan from reporting detection engine	Enable to scan a host from the appliance where the detection engine that reported the host resides.	N/A
Fast Port Scan	Enable to scan only the TCP ports listed in the <code>nmap-services</code> file located in the <code>/var/sf/nmap/share/nmap/nmap-services</code> directory on the device that does the scanning, ignoring other port settings. Note that you cannot use this option with the Port Ranges and Scan Order option.	-F
Port Ranges and Scan Order	Set the specific ports you want to scan, using Nmap port specification syntax, and the order you want to scan them. Note that you cannot use this option with the Fast Port Scan option.	-p
Probe open ports for vendor and version information	Enable to detect server vendor and version information. If you probe open ports for server vendor and version information, Nmap obtains server data that it uses to identify servers. It then replaces the Sourcefire server data for that server.	-SV
Service Version Intensity	Select the intensity of Nmap probes for service versions. Higher service intensity numbers cause more probes to be used and result in higher accuracy, while lower intensity probes are faster but obtain less information.	--version-intensity <i><intensity></i>
Detect Operating System	Enable to detect operating system information for the host. If you configure detection of the operating system for a host, Nmap scans the host and uses the results to create a rating for each operating system that reflects the likelihood that the operating system is running on the host. For more information on when and how Nmap-identified identity data appears in the network map, see Understanding Current Identities on page 1646.	-o

Nmap Remediation Options (Continued)

OPTION	DESCRIPTION	CORRESPONDING NMAP OPTION
Treat All Hosts As Online	Enable to skip the host discovery process and run a port scan on every host in the target range. Note that when you enable this option, Nmap ignores settings for Host Discovery Method and Host Discovery Port List .	-PN
Host Discovery Method	<p>Select to perform host discovery for all hosts in the target range, over the ports listed in the Host Discovery Port List, or if no ports are listed, over the default ports for that host discovery method.</p> <p>Note that if you also enabled Treat All Hosts As Online, however, the Host Discovery Method option has no effect and host discovery is not performed.</p> <p>Select the method to be used when Nmap tests to see if a host is present and available:</p> <ul style="list-style-type: none"> • The TCP SYN option sends an empty TCP packet with the SYN flag set and recognizes the host as available if a response is received. TCP SYN scans port 80 by default. Note that TCP SYN scans are less likely to be blocked by a firewall with stateful firewall rules. • The TCP ACK option sends an empty TCP packet with the ACK flag set and recognizes the host as available if a response is received. TCP ACK also scans port 80 by default. Note that TCP ACK scans are less likely to be blocked by a firewall with stateless firewall rules. • The UDP option sends a UDP packet and assumes host availability if a port unreachable response comes back from a closed port. UDP scans port 40125 by default. 	<p>TCP SYN: -PS TCP ACK: -PA UDP: -PU</p>
Host Discovery Port List	Specify a customized list of ports, separated by commas, that you want to scan when doing host discovery.	port list for host discovery method
Default NSE Scripts	Enable to run the default set of Nmap scripts for host discovery and server and operating system and vulnerability detection. See http://nmap.org/nsedoc/categories/default.html for the list of default scripts.	-SC
Timing Template	Select the timing of the scan process; the higher the number you select, the faster and less comprehensive the scan.	<p>0: T0 (paranoid) 1: T1 (sneaky) 2: T2 (polite) 3: T3 (normal) 4: T4 (aggressive) 5: T5 (insane)</p>

Creating an Nmap Scanning Strategy

LICENSE: FireSIGHT

While active scanning can obtain valuable information, overuse of a tool such as Nmap may overload your network resources or even crash important hosts. When using any active scanner, you should create a scanning strategy to make sure that you are scanning only the hosts and ports that you need to scan.

For more information, see the following sections:

- [Selecting Appropriate Scan Targets](#) on page 1697
- [Selecting Appropriate Ports to Scan](#) on page 1698
- [Setting Host Discovery Options](#) on page 1698

Selecting Appropriate Scan Targets

LICENSE: FireSIGHT

When you configure Nmap, you can create scan targets that identify which hosts you want to scan. A scan target includes a single IP address, a CIDR block or octet range of IP addresses, an IP address range, or a list of IP addresses or ranges to scan, as well as the ports on the host or hosts.

You can specify targets in the following ways:

- For IPv6 hosts:
 - an exact IP address (for example, 192.168.1.101)
- For IPv4 hosts:
 - an exact IP address (for example, 192.168.1.101) or a list of IP addresses separated by commas or spaces
 - an IP address range using CIDR notation (for example, 192.168.1.0/24 scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive)
 - For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
 - an IP address range using octet range addressing (for example, 192.168.0-255.1-254 scans all addresses in the 192.168.x.x range, except those that end in .0 and or .255)
 - an IP address range using hyphenation (for example, 192.168.1.1 - 192.168.1.5 scans the six hosts between 192.168.1.1 and 192.168.1.5, inclusive)
 - a list of addresses or ranges separated by commas or spaces (for example, for example, 192.168.1.0/24, 194.168.1.0/24 scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive and the 254 hosts between 194.168.1.1 and 194.168.1.254, inclusive)

Ideal scan targets for Nmap scans include hosts with operating systems that the system is unable to identify, hosts with unidentified servers, or hosts recently

detected on your network. Remember that Nmap results cannot be added to the network map for hosts that do not exist in the network map.

WARNING! Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up to date. For more information, see [Automating Nmap Scans](#) on page 1936. Also note that if the host is deleted from the network map, any Nmap scan results are discarded. In addition, make sure you have permission to scan your targets. Using Nmap to scan hosts that do not belong to you or your company may be illegal.

Selecting Appropriate Ports to Scan

LICENSE: FireSIGHT

For each scan target you configure, you can select the ports you want to scan. You can designate individual port numbers, port ranges, or a series of port numbers and port ranges to identify the exact set of ports that should be scanned on each target.

By default, Nmap scans TCP ports 1 through 1024. If you plan to use the remediation as a response in a correlation policy, you can cause the remediation to scan only the port specified in the event that triggers the correlation response. If you run the remediation on demand or as a scheduled task, or if you do not use the port from the event, you can use other port options to determine which ports are scanned. You can choose to scan only the TCP ports listed in the `nmap-services` file, ignoring other port settings. You can also scan UDP ports in addition to TCP ports. Note that scanning for UDP ports may be time-consuming, so avoid using that option if you want to scan quickly. To select the specific ports or range of ports to scan, use Nmap port specification syntax to identify ports.

Setting Host Discovery Options

LICENSE: FireSIGHT

You can decide whether to perform host discovery before starting a port scan for a host, or you can assume that all the hosts you plan to scan are online. If you choose not to treat all hosts as online, you can choose what method of host discovery to use and, if needed, customize the list of ports scanned during host discovery. Host discovery does not probe the ports listed for operating system or server information; it uses the response over a particular port only to determine whether a host is active and available. If you perform host discovery and a host is not available, Nmap does not scan ports on that host.

Sample Nmap Scanning Profiles

LICENSE: FireSIGHT

The following scenarios provide examples of how Nmap might be used on your network:

- [Example: Resolving Unknown Operating Systems](#) on page 1699
- [Example: Responding to New Hosts](#) on page 1700

Example: Resolving Unknown Operating Systems

LICENSE: FireSIGHT

If the system cannot determine the operating system on a host on your network, you can use Nmap to actively scan the host. Nmap uses the information it obtains from the scan to rate the possible operating systems. It then uses the operating system that has the highest rating as the host operating system identification.

Using Nmap to challenge new hosts for operating system and server information deactivates the system's monitoring of that data for scanned hosts. If you use Nmap to discover host and server operating system for hosts the system marks as having unknown operating systems, you may be able to identify groups of hosts that are similar. You can then create a custom fingerprint based on one of them to cause the system to associate the fingerprint with the operating system you know is running on the host based on the Nmap scan. Whenever possible, create a custom fingerprint rather than inputting static data through a third-party source like Nmap because the custom fingerprint allows the system to continue to monitor the host operating system and update it as needed.

To discover operating systems with Nmap:

ACCESS: Admin/Discovery Admin

1. Configure a scan instance for an Nmap module.
For more information, see [Creating an Nmap Scan Instance](#) on page 1702.
2. Create an Nmap remediation using the following settings:
 - Enable **Use Port From Event** to scan the port associated with the new server.
 - Enable **Detect Operating System** to detect operating system information for the host.
 - Enable **Probe open ports for vendor and version information** to detect server vendor and version information.
 - Enable **Treat All Hosts as Online**, because you know the host exists.

For information on creating Nmap remediations, see [Creating an Nmap Remediation](#) on page 1705.

3. Create a correlation rule that triggers when the system detects a host with an unknown operating system.

The rule should trigger when **an discovery event occurs** and **the OS information for a host has changed** and it meets the following conditions: **OS Name is unknown**.

For information on creating correlation rules, see [Creating Rules for Correlation Policies](#) on page 1461.

4. Create a correlation policy that contains the correlation rule.
For more information on creating correlation policies, see [Creating Correlation Policies](#) on page 1513.
5. In the correlation policy, add the Nmap remediation you created in step 2 as a response to the rule you created in step 3.
6. Activate the correlation policy.
7. Purge the hosts on your network map to force network discovery to restart and rebuild the network map.
8. After a day or two, search for events generated by the correlation policy. Analyze the Nmap results for the operating systems detected on the hosts to see if there is a particular host configuration on your network that the system does not recognize.
For more information on analyzing Nmap results, see [Analyzing Scan Results](#) on page 1719.
9. If you find hosts with unknown operating systems whose Nmap results are identical, create a custom fingerprint for one of those hosts and use it to identify similar hosts in the future.

For more information, see [Fingerprinting Clients](#) on page 1650.

Example: Responding to New Hosts

LICENSE: FireSIGHT

When the system detects a new host in a subnet where intrusions may be likely, you may want to scan that host to make sure you have accurate vulnerability information for it.

You can accomplish this by creating and activating a correlation policy that detects when a new host appears in this subnet, and that launches a remediation that performs an Nmap scan on the host.

After you activate the policy, you can periodically check the remediation status view (**Policy & Response > Responses > Remediations > Status**) to see when the remediation launched. The remediation's dynamic scan target should include the IP addresses of the hosts it scanned as a result of the server detection. Check the host profile for those hosts to see if there are vulnerabilities that need to be

addressed for the host, based on the operating system and servers detected by Nmap.

WARNING! If you have a large or dynamic network, detection of a new host may be too frequent an occurrence to respond to using a scan. To prevent resource overload, avoid using Nmap scans as a response to events that occur frequently. In addition, note that using Nmap to challenge new hosts for operating system and server information deactivates Sourcefire monitoring of that data for scanned hosts.

To scan in response to the appearance of a new host:

ACCESS: Admin/Discovery Admin

1. Configure a scan instance for an Nmap module.
For more information, see [Creating an Nmap Scan Instance](#) on page 1702.
2. Create an Nmap remediation using the following settings:
 - Enable **Use Port From Event** to scan the port associated with the new server.
 - Enable **Detect Operating System** to detect operating system information for the host.
 - Enable **Probe open ports for vendor and version information** to detect server vendor and version information.
 - Enable **Treat All Hosts as Online**, because you know the host exists.For information on creating Nmap remediations, see [Creating an Nmap Remediation](#) on page 1705.
3. Create a correlation rule that triggers when the system detects a new host on a specific subnet.
The rule should trigger when **a discovery event occurs** and **a new host is detected**.
For information on creating correlation rules, see [Creating Rules for Correlation Policies](#) on page 1461.
4. Create a correlation policy that contains the correlation rule.
For more information on creating correlation policies, see [Creating Correlation Policies](#) on page 1513.
5. In the correlation policy, add the Nmap remediation you created in step 4 as a response to the rule you created in step 3.
6. Activate the correlation policy.
7. When you are notified of a new host, check the host profile to see the results of the Nmap scan and address any vulnerabilities that apply to the host.

Setting up Nmap Scans

LICENSE: FireSIGHT

To scan using Nmap, you must first configure a scan instance and a scan remediation. If you plan to schedule Nmap scans, you must also define a scan target.

For more information, see the following sections:

- [Creating an Nmap Scan Instance](#) on page 1702
- [Creating an Nmap Scan Target](#) on page 1704
- [Creating an Nmap Remediation](#) on page 1705

Creating an Nmap Scan Instance

LICENSE: FireSIGHT

You can set up a separate scan instance for each Nmap module that you want to use to scan your network for vulnerabilities. You can set up scan instances for the local Nmap module on your Defense Center and for any devices you want to use to run scans remotely. The results of each scan are always stored on the Defense Center where you configure the scan, even if you run the scan from a remote device. To prevent accidental or malicious scanning of mission-critical hosts, you can create a blacklist for the instance to indicate the hosts that should never be scanned with the instance.

Note that you cannot add a scan instance with the same name as any existing scan instance.

To create a scan instance:

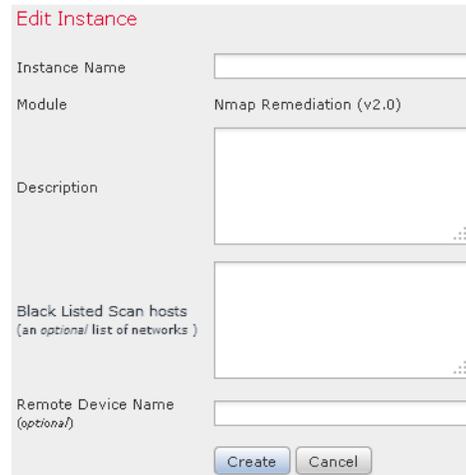
ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.

The Scanners page appears.

2. Click **Add Nmap Instance**.

The Instance Detail page appears.



The screenshot shows a web form titled "Edit Instance" for the "Nmap Remediation (v2.0)" module. The form contains the following fields:

- Instance Name:** A text input field.
- Module:** A dropdown menu currently showing "Nmap Remediation (v2.0)".
- Description:** A large text area with a "..." icon at the bottom right.
- Black Listed Scan hosts (an optional list of networks):** A text area with a "..." icon at the bottom right.
- Remote Device Name (optional):** A text input field.

At the bottom of the form are two buttons: "Create" and "Cancel".

3. In the **Instance Name** field, enter a name that includes 1 to 63 alphanumeric characters, with no spaces and no special characters other than underscore (_) and dash (-).
4. In the **Description** field, specify a description with 0 to 255 alphanumeric characters, which can include spaces and special characters.
5. Optionally, in the **Black Listed Scan hosts** field, specify any hosts or networks that should *never* be scanned with this scan instance, using the following syntax:
 - For IPv6 hosts, an exact IP address (for example, `2001:DB8::fedd:eeff`)
 - For IPv4 hosts, an exact IP address (for example, `192.168.1.101`) or an IP address range using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between `192.168.1.1` and `192.168.1.254`, inclusive)

If you specifically target a scan to a host that is in a blacklisted network, that scan will not run.

6. Optionally, to run the scan from a remote device instead of the Defense Center, specify the IP address or name of the device as it appears in the Information page for the device in the Defense Center web interface, in the **Remote Device Name** field.
7. Click **Create**.

The scan instance is created.

Creating an Nmap Scan Target

LICENSE: FireSIGHT

You can create and save scan targets that identify specific hosts and ports. Then, when you perform an on-demand scan or schedule a scan, you can use one of the saved scan targets.

For scans of targets with IPv4 addresses, you can use an IP address, a list of IP addresses, CIDR notation, or Nmap scan octets to select the hosts to scan. You can also specify a range of addresses using a hyphen. Separate addresses and ranges in a list with commas or spaces.

For scans of IPv6 addresses, use an IP address. Ranges are not supported.

Note that Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up to date. For more information, see [Automating Nmap Scans](#) on page 1936. Also note that if the host is deleted from the network map, any Nmap scan results for that host are discarded.

To create a scan target:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scanners page appears.
2. On the toolbar, click **Targets**.
The Scan Target List page appears.

Scan Targets

Name
Sample Scan Target

3. Click **Create Scan Target**.
The Scan Target page appears.

Target Information

Name

IP Range

Ports

4. In the **Name** field, type the name you want to use for this scan target.

5. In the **IP Range** text box, specify the host or hosts you want to scan, using the following syntax:
 - for IPv6 hosts, an exact IP address (for example, `2001:DB8::fedd:eeff`)
 - for IPv4 hosts, an exact IP address (for example, `192.168.1.101`) or comma-separated list of IP addresses
 - for IPv4 hosts, an IP address range using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive)
For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
 - for IPv4 hosts, an IP address range using octet range addressing (for example, `192.168.0-255.1-254` scans all addresses in the `192.168.x.x` range, except those that end in `.0` and or `.255`)
 - for IPv4 hosts, an IP address range using hyphenation (for example, `192.168.1.1 - 192.168.1.5` scans the 6 hosts between 192.168.1.1 and 192.168.1.5, inclusive)
 - for IPv4 hosts, a list of addresses or ranges separated by commas or spaces (for example, for example, `192.168.1.0/24, 194.168.1.0/24` scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive and the 254 hosts between 194.168.1.1 and 194.168.1.254, inclusive)

IMPORTANT! The **IP Range** text box accepts up to 255 characters. In addition, note that if you use a comma in a list of IP addresses or ranges in a scan target, the comma converts to a space when you save the target.

6. In the **Ports** field, specify the ports you want to scan.
You can enter any of the following, using values from 1 to 65535:
 - a port number
 - a list of ports separated by commas
 - a range of port numbers separated by a dash
 - ranges of port numbers separated by dashes, separated by commas
7. Click **Save**.
The scan target is created.

Creating an Nmap Remediation

LICENSE: FireSIGHT

You can define the settings for an Nmap scan by creating an Nmap remediation. An Nmap remediation can be used as a response in a correlation policy, run on demand, or scheduled to run at a specific time. In order for the results of an

Nmap scan to appear in the network map, the scanned host must already exist in the network map.

For more information on the specific settings in an Nmap remediation, see [Understanding Nmap Remediations](#) on page 1693.

Note that Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host for operating system and server data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up-to-date. For more information, see [Automating Nmap Scans](#) on page 1936. Also note that if the host is deleted from the network map, any Nmap scan results for that host are discarded.

For general information about Nmap functionality, refer to the Nmap documentation at <http://insecure.org>.

To create an Nmap remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.

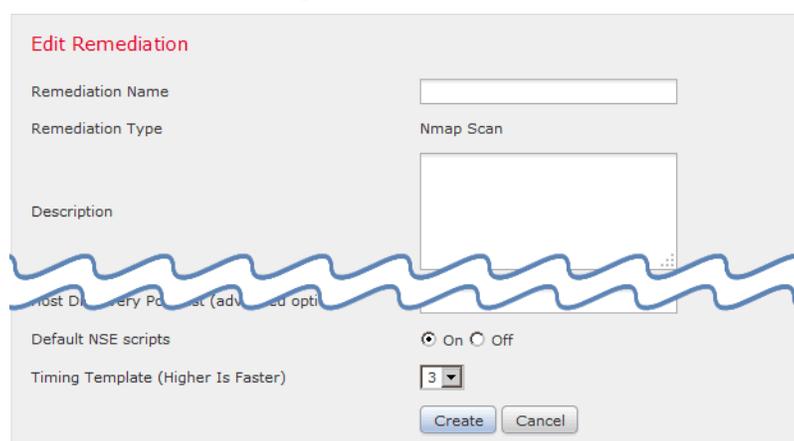
The Scanners page appears.

Nmap Scan Instances



2. Click **Add Remediation** next to the scan instance where you want to add a remediation.

The Edit Remediation page appears.



3. In the **Remediation Name** field, type a name for the remediation that includes 1 to 63 alphanumeric characters, with no spaces and no special characters other than underscore (_) and dash (-).

4. In the **Description** field, type a description for the remediation that includes 0 to 255 alphanumeric characters, including spaces and special characters.
5. If you plan to use this remediation in response to a correlation rule that triggers on an intrusion event, a connection event, or a user event, configure the **Scan Which Address(es) From Event?** option:
 - Select **Scan Source and Destination Addresses** to scan the hosts represented by the source IP address and the destination IP address in the event.
 - Select **Scan Source Address Only** to scan the host represented by the event's source IP address.
 - Select **Scan Destination Address Only** to scan the host represented by the event's destination IP address.

If you plan to use this remediation in response to a correlation rule that triggers on a discovery event or a host input event, by default the remediation scans the IP address of the host involved in the event; you do not need to configure this option.

IMPORTANT! Do **not** assign an Nmap remediation as a response to a correlation rule that triggers on a traffic profile change.

6. Configure the **Scan Type** option:
 - To scan quickly in stealth mode on hosts where the **admin** account has raw packet access or where IPv6 is not running, by initiating TCP connections but not completing them, select **TCP Syn Scan**.
 - To scan by using a system **connect()** call, which can be used on hosts where the **admin** account on your Defense Center does not have raw packet access or where IPv6 is running, select **TCP Connect Scan**.
 - To send an ACK packet to check whether ports are filtered or unfiltered, select **TCP ACK Scan**.
 - To send an ACK packet to check whether ports are filtered or unfiltered but also to determine whether a port is open or closed, select **TCP Window Scan**.
 - To identify BSD-derived systems using a FIN/ACK probe, select **TCP Maimon Scan**.
7. Optionally, to scan UDP ports in addition to TCP ports, select **On** for the **Scan for UDP ports** option.

TIP! A UDP portscan takes more time than a TCP portscan. To speed up your scans, leave this option disabled.

8. If you plan to use this remediation in response to correlation policy violations, configure the **Use Port From Event** option:
 - Select **On** to scan the port in the correlation event, rather than the ports you specify in step 11.

If you scan the port in the correlation event, note that the remediation scans the port on the IP addresses that you specified in step 5. These ports are also added to the remediation's dynamic scan target.
 - Select **Off** to scan only the ports you will specify in step 11.
9. If you plan to use this remediation in response to correlation policy violations and want to run the scan using the appliance running the detection engine that detected the event, configure the **Scan from reporting detection engine** option:
 - To scan from the appliance running the reporting detection engine, select **On**.
 - To scan from the appliance configured in the remediation, select **Off**.
10. Configure the **Fast Port Scan** option:
 - To scan only the ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the device that does the scanning, ignoring other port settings, select **On**.
 - To scan all TCP ports, select **Off**.
11. In the **Port Ranges and Scan Order** field, type the ports you want to scan by default, using Nmap syntax, in the order you want to scan those ports.

Specify values from 1 to 65535. Separate ports using commas or spaces. You can also use a hyphen to indicate a port range. When scanning for both TCP and UDP ports, preface the list of TCP ports you want to scan with a T and the list of UDP ports with a U. For example, to scan ports 53 and 111 for UDP traffic, then scan ports 21-25 for TCP traffic, enter `u:53,111,t:21-25`.

Note that the **Use Port From Event** option overrides this setting when the remediation is launched in response to a correlation policy violation, as described in step 8.
12. To probe open ports for server vendor and version information, configure **Probe open ports for vendor and version information**:
 - Select **On** to scan open ports on the host for server information to identify server vendors and versions.
 - Select **Off** to continue using Sourcefire server information for the host.
13. If you choose to probe open ports, set the number of probes used by selecting a number from the **Service Version Intensity** drop-down list:
 - To use more probes for higher accuracy with a longer scan, select a higher number.
 - To use fewer probes for less accuracy with a faster scan, select a lower number.

14. To scan for operating system information, configure **Detect Operating System** settings:
 - Select **On** to scan the host for information to identify the operating system.
 - Select **Off** to continue using Sourcefire operating system information for the host.
15. To determine whether host discovery occurs and whether port scans are only run against available hosts, configure **Treat All Hosts As Online**:
 - To skip the host discovery process and run a port scan on every host in the target range, select **On**.
 - To perform host discovery using the settings for **Host Discovery Method** and **Host Discovery Port List** and skip the port scan on any host that is not available, select **Off**.
16. Select the method you want Nmap to use when it tests for host availability:
 - To send an empty TCP packet with the SYN flag set and elicit an RST response on a closed port or a SYN/ACK response on an open port on available hosts, select **TCP SYN**.

Note that this option scans port 80 by default and that TCP SYN scans are less likely to be blocked by a firewall with stateful firewall rules.
 - To send an empty TCP packet with the ACK flag set and elicit an RST response on available hosts, select **TCP ACK**.

Note that this option scans port 80 by default and that TCP ACK scans are less likely to be blocked by a firewall with stateless firewall rules.
 - To send a UDP packet to elicit port unreachable responses from closed ports on available hosts, select **UDP**. This option scans port 40125 by default.
17. If you want to scan a custom list of ports during host discovery, type a list of ports appropriate for the host discovery method you selected, separated by commas, in the **Host Discovery Port List** field.
18. Configure the **Default NSE Scripts** option to control whether to use the default set of Nmap scripts for host discovery and server, operating system, and vulnerability discovery:
 - To run the default set of Nmap scripts, select **On**.
 - To skip the default set of Nmap scripts, select **Off**.See <http://nmap.org/nsedoc/categories/default.html> for the list of default scripts.
19. To set the timing of the scan process, select a timing template number; select a higher number for a faster, less comprehensive scan and a lower number for a slower, more comprehensive scan.

20. Click **Save**, then click **Done**.
The remediation is created.

Managing Nmap Scanning

LICENSE: FireSIGHT

You can modify or delete Nmap scan instances and remediations as needed. You can also run an on-demand Nmap scan. You can also view or download Nmap results for previous scans. For more information, see the following sections:

- [Managing Nmap Scan Instances](#) on page 1710
- [Managing Nmap Remediations](#) on page 1712
- [Running an On-Demand Nmap Scan](#) on page 1713

Managing Nmap Scan Instances

LICENSE: FireSIGHT

You can edit or delete Nmap scan instances. For more information, see the following sections:

- [Editing an Nmap Scan Instance](#) on page 1710
- [Deleting an Nmap Scan Instance](#) on page 1711

Editing an Nmap Scan Instance

LICENSE: FireSIGHT

Use the following procedure to modify scan instances. Note that you can view, add, and delete remediations associated with the instance when you modify it.

To edit a scan instance:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scanners page appears.

2. Click **View** next to the instance you want to edit.
The Instance Detail page appears.

Edit Instance

Instance Name

Module Nmap Remediation (v2.0)

Description

Black Listed Scan hosts
(an optional list of networks)

Remote Device Name
(optional)

3. Optionally, click **View** next to the remediation you want to view or edit.
For more information on editing remediations, see [Editing an Nmap Remediation](#) on page 1712.
4. Optionally, click **Delete** next to the remediation you want to delete.
For more information on deleting remediations, see [Deleting an Nmap Remediation](#) on page 1712.
5. Optionally, click **Add** to add a new remediation to this scan instance.
For more information on creating new remediations, see [Managing Nmap Remediations](#) on page 1712.
6. Optionally, make changes to the scan instance settings, then click **Save**.
7. Click **Done**.
The scan instance is modified.

Deleting an Nmap Scan Instance

LICENSE: FireSIGHT

Delete an Nmap scan instance when you no longer want to use the Nmap module profiled in the instance. Note that when you delete the scan instance, you also delete any remediations that use that instance.

To delete a scan instance:

ACCESS: Admin/Discovery Admin

1. Click **Policies > Actions > Scanners**.
The Scanners page appears.

2. Click **Delete** next to the scan instance you want to delete.
The instance is deleted.

Managing Nmap Remediations

LICENSE: FireSIGHT

You can edit or delete Nmap remediations. For more information, see the following sections:

- [Editing an Nmap Remediation](#) on page 1712
- [Deleting an Nmap Remediation](#) on page 1712

Editing an Nmap Remediation

LICENSE: FireSIGHT

Modifications you make to Nmap remediations do not affect scans in progress. The new settings take effect when the next scan starts.

To edit an Nmap remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scanners page appears.
2. Next to the remediation you want to edit, click **View**.
The Remediation Edit page appears.
3. Make modifications as necessary.
For information on the settings you can change, see [Creating an Nmap Remediation](#) on page 1705.
4. Click **Save**, then click **Done**.
The remediation is modified.

Deleting an Nmap Remediation

LICENSE: FireSIGHT

Delete an Nmap remediation if you no longer need it.

To delete an Nmap remediation:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scanners page appears.
2. Next to the remediation you want to delete, click **Delete**.
3. Confirm that you want to delete the remediation.
The remediation is deleted.

Running an On-Demand Nmap Scan

LICENSE: FireSIGHT

You can launch on-demand Nmap scans whenever needed. You can specify the target for an on-demand scan by entering the IP addresses and ports you want to scan or by selecting an existing scan target.

Note that Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up to date. For more information, see [Automating Nmap Scans](#) on page 1936. In addition, note that if the host is deleted from the network map, any Nmap scan results are discarded.

To run an on-demand Nmap scan:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.

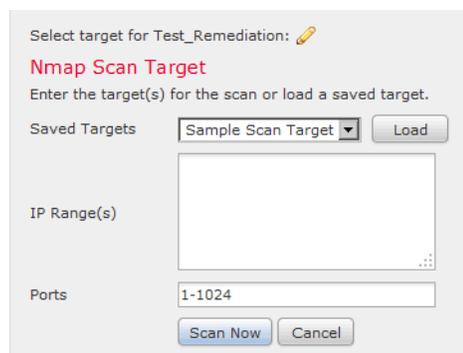
The Scanners page appears.

Nmap Scan Instances



2. Next to the Nmap remediation you want to use to perform the scan, click **Scan**.

The Nmap Scan Target dialog box appears.

A screenshot of the 'Nmap Scan Target' dialog box. The title is 'Select target for Test_Remediation:'. Below the title is the heading 'Nmap Scan Target' and the instruction 'Enter the target(s) for the scan or load a saved target.' There are three main sections: 'Saved Targets' with a dropdown menu showing 'Sample Scan Target' and a 'Load' button; 'IP Range(s)' with a large empty text area; and 'Ports' with a text input field containing '1-1024'. At the bottom are 'Scan Now' and 'Cancel' buttons.

3. Optionally, to scan using a saved scan target, select a target from the **Saved Targets** drop-down list and click **Load**.

The IP addresses and ports associated with the scan target populate the **IP Range(s)** and **Ports** fields.

TIP! To create a scan target, click **Edit/Add Targets**. For more information, see [Creating an Nmap Scan Target](#) on page 1704.

4. In the **IP Range(s)** field, specify the IP address for hosts you want to scan or modify the loaded list, up to 255 characters.
For hosts with IPv4 addresses, you can specify multiple IP addresses separated by commas or use CIDR notation. You can also negate IP addresses by preceding them with an exclamation point (!). For information on using CIDR notation in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
For hosts with IPv6 addresses, use an exact IP address. Ranges are not supported.
5. In the **Ports** field, specify the ports you want to scan or modify the loaded list.
You can enter a port number, a list of ports separated by commas, or a range of port numbers separated by a dash. For details on entering ports, see [Specifying Ports in Searches](#) on page 1777.
6. Click **Scan Now**.
The Nmap server performs the scan.
Note that Nmap validates IP address ranges and displays an error message if the range is invalid. If this occurs, correct the contents of the **IP Range(s)** field to indicate a valid IP address range.

Managing Scan Targets

LICENSE: FireSIGHT

When you configure an Nmap module, you can create and save scan targets that identify the hosts and ports you want to target when you perform an on-demand or a scheduled scan, so that you do not have to construct a new scan target every time. A scan target includes a single IP address or a block of IP addresses to scan, as well as the ports on the host or hosts. For Nmap targets, you can also use Nmap octet range addressing or IP address ranges. For more information on Nmap octet range addressing, refer to the Nmap documentation at <http://insecure.org>.

Note that scans for scan targets containing a large number of hosts can take an extended period of time. As a workaround, scan fewer hosts at a time.

After you create a scan target, you can modify or delete it.

For more information, see the following sections:

- [Creating an Nmap Scan Target](#) on page 1704
- [Editing a Scan Target](#) on page 1715
- [Deleting a Scan Target](#) on page 1716

Editing a Scan Target

LICENSE: FireSIGHT

You can modify scan targets you created.

TIP! You might want to edit a remediation's dynamic scan target if you do not want to use the remediation to scan a specific IP address, but the IP address was added to the target because the host was involved in a correlation policy violation that launched the remediation.

To edit an existing scan target:

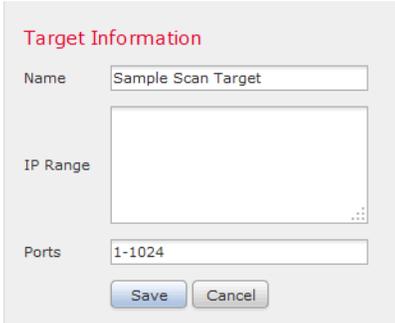
ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scanners page appears.
2. On the toolbar, click **Targets**.
The Scan Target List page appears.

Scan Targets

Name
Sample Scan Target  

3. Click **Edit** next to the scan target you want to edit.
The Scan Target page appears.



The dialog box titled "Target Information" contains the following fields and buttons:

- Name:** A text input field containing "Sample Scan Target".
- IP Range:** A large empty text area for entering IP addresses.
- Ports:** A text input field containing "1-1024".
- Buttons:** "Save" and "Cancel" buttons at the bottom.

4. Make modifications as necessary and click **Save**.
The scan target is updated.

Deleting a Scan Target

LICENSE: FireSIGHT

Delete a scan target if you no longer want to scan the hosts listed in it.

To delete a scan target:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scanners page appears.
2. On the toolbar, click **Targets**.
The Scan Target List page appears.

Scan Targets

Name
Sample Scan Target

3. Next to the scan target you want to delete, click **Delete**.
The scan target is deleted.

Working with Active Scan Results

LICENSE: FireSIGHT

For information on how to monitor Nmap scans in progress, import results from scans previously performed through the Sourcefire 3D System or results performed outside the Sourcefire 3D System, and view and analyze scan results, see the following sections:

- [Viewing Scan Results](#) on page 1716
- [Understanding the Scan Results Table](#) on page 1718
- [Analyzing Scan Results](#) on page 1719
- [Monitoring Scans](#) on page 1719
- [Importing Scan Results](#) on page 1720
- [Searching for Scan Results](#) on page 1721

Viewing Scan Results

LICENSE: FireSIGHT

You can view a table of scan results, and then manipulate the event view depending on the information you are looking for.

The page you see when you access scan results differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of scan results.

You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The [Scan Results Table Functions](#) table below describes some of the specific actions you can perform on a scan results workflow page.

Scan Results Table Functions

To...	YOU CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Scan Results Table on page 1718.
modify the time and date range for the scan result	click the time range link. For more information, see Setting Event Time Constraints on page 1820.
sort scan results	click the column title. Click the column title again to reverse the sort order.
constrain the columns that appear	<p>click the close icon (✕) in the column heading that you want to hide. In the pop-up window that appears, click Apply.</p> <p>TIP! To hide or show other columns, select or clear the appropriate check boxes before you click Apply. To add a disabled column back to the view,</p> <p>Click the expand arrow (▶) to expand the search constraints, then click the column name under Disabled Columns.</p>
drill down to the next page in the workflow, constraining on a specific value	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • on a drill-down page that you created in a custom workflow, click a value within a row. Note that clicking a value within a row in a table view constrains the table view and does not drill down to the next page. • To drill down to the next workflow page constraining on some users, select the check boxes next to the users you want to view on the next workflow page, then click View. • To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include "Table View" in the page name.</p> <p>For more information, see Constraining Events on page 1830.</p>

Scan Results Table Functions (Continued)

To...	YOU CAN...
configure scan instances and remediations	Click Scanners in the toolbar. For more information, see Setting up Nmap Scans on page 1702.
navigate within and between workflow pages	find more information in Using Workflow Pages on page 1812.
navigate to other event views to view associated events	click the name of the event view you want to see. For more information, see Navigating Between Workflows on page 1836.
search for scan results	click Search . For more information, see Searching for Scan Results on page 1721.

To view scan results:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
2. Click **Scan Results**.

The first page of the default scan results workflow appears. To use a different workflow, including a custom workflow, click (**switch workflows**) by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

Understanding the Scan Results Table

LICENSE: FireSIGHT

When you run an Nmap scan, the Defense Center collects the scan results in a database. The fields in the scan results table are described in the [Scan Results Fields](#) table.

Scan Results Fields

FIELD	DESCRIPTION
Start Time	The date and time that the scan that produced the results started.
End Time	The date and time that the scan that produced the results ended.

Scan Results Fields (Continued)

FIELD	DESCRIPTION
Scan Target	The IP address (or host name, if DNS resolution is enabled) of the scan target for the scan that produced the results.
Scan Type	Either Nmap or the name of the third-party scanner to indicate the type of the scan that produced the results.
Scan Mode	The mode of the scan that produced the results: <ul style="list-style-type: none">• On Demand — results from scans run on demand.• Imported — results from scans on a different system and imported onto the Defense Center.• Scheduled — results from scans run as a scheduled task.

Analyzing Scan Results

LICENSE: FireSIGHT

You can view scan results that you create using the local Nmap module as a rendered page in a pop-up window. You can also download the Nmap results file in raw XML format.

You can also view operating system and server information detected by Nmap in host profiles and in the network map. If a scan of a host produces server information for servers on filtered or closed ports, or if a scan collects information that cannot be included in the operating system information or the servers section, the host profile includes those results in an Nmap Scan Results section. For more information, see [Viewing Host Profiles](#) on page 1337.

Monitoring Scans

LICENSE: FireSIGHT

You can check the progress of an Nmap scan and cancel scan jobs currently in progress. Scan results provide the start time and end time of each scan. Also, after a scan is completed, you can also view the scan results as a rendered page in a pop-up window. Nmap results you can download and view using the Nmap Version 1.01 DTD, available at <http://insecure.org>. You can also clear scan results.

To monitor a scan:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
2. Click **Scan Results**.

The first page of the default scan results workflow appears. To use a different workflow, including a custom workflow, click **(switch workflows)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

TIP! If you are using a custom workflow that does not include the table view of scan results, click **(switch workflows)** by the workflow title, then select **Scan Results**.

3. You can perform the following actions:
 - To view the scan results as a rendered page in a pop-up window, click **View** next to the scan job.
 - To save a copy of the scan results file so that you can view the raw XML code in any text editor, click **Download** next to the scan job.

Importing Scan Results

LICENSE: FireSIGHT

You can import XML results files created by an Nmap scan performed outside of the Sourcefire 3D System. You can also import XML results files that you previously downloaded from the Sourcefire 3D System. To import Nmap scan results, the results file must be in XML format and adhere to the Nmap Version 1.01 DTD. For more information on creating Nmap results and on the Nmap DTD, refer to the Nmap documentation at <http://insecure.org>. For information on downloading XML results from the Sourcefire 3D System, see [Monitoring Scans](#) on page 1719.

To import results:

ACCESS: Admin/Discovery Admin

1. Select **Policies > Actions > Scanners**.
The Scan Instances page appears.

- On the toolbar, click **Import Results**.
The Import Results page appears.

Import Scan Results

- Click **Browse** to navigate to the results file.
- After you return to the Import Results page, click **Import** to import the results.
The results file is imported.

Searching for Scan Results

LICENSE: FireSIGHT

You can search for Nmap or third-party scan results for any scans run on an appliance or managed appliance in your Sourcefire 3D System.

Scan Results Search Criteria

FIELD	SEARCH CRITERIA RULES
Start Time	Type the date and time that the scan that produced the results started. See Specifying Time Constraints in Searches on page 1775 for the syntax for entering time.
End Time	Type the date and time that the scan that produced the results ended. See Specifying Time Constraints in Searches on page 1775 for the syntax for entering time.
Scan Target	Type the IP address (or host name, if DNS resolution is enabled) of the scan target for the scan that produced the results. Use a specific IP address or CIDR notation to specify a range of IP addresses. See Specifying IP Addresses in Searches on page 1776 for a full description of the syntax allowed for IP addresses.
Scan Type	Type Nmap or a third-party scanner ID to indicate the type of the scan that produced the results.
Scan Mode	Type the mode of the scan that produced the results: <ul style="list-style-type: none"> Type On Demand to retrieve results from scans run on demand. Type Imported to retrieve results from scans on a different system and imported onto the Defense Center. Type Scheduled to retrieve results from scans run as a scheduled task.

For more information on searching, including how to load and delete saved searches, see [Searching for Events](#) on page 1770.

To search for scan results:

ACCESS: Admin/Discovery Admin

1. Select **Analysis > Search**, then select **Scan Results** from the **Table** drop-down list.

The Scan Results search page appears.

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: Scan Results

Name: Search 1, My Search

Save As Private:

Constraint

Start Time: > 2009-07-16 13:00:31, < today at 4:30pm

End Time: > 2009-07-16 13:00:31, < today at 4:30pm

Target: 192.168.1.2, 2001:db8:85a3::1370

Scan Type: Nmap

Scan Mode: On Demand, Scheduled

Search Save As New Search

TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

2. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, the Defense Center automatically creates one when you save the search.
3. Enter your search criteria in the appropriate fields, as described in the [Scan Results Search Criteria](#) table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria.
4. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search so that only you can use it.

TIP! If you want to save a search as a restriction for custom user roles with restricted privileges (or for converted Restricted Event Analysts from pre-4.10.1 versions), you **must** save it as a private search.

5. You have the following options:
 - Click **Search** to start the search.
Your search results appear.
 - Click **Save** if you are modifying an existing search and want to save your changes.
6. Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

CHAPTER 42

WORKING WITH REPORTS

The Sourcefire 3D System provides a flexible reporting system that allows you to quickly and easily generate multi-section reports with the event views or dashboards that appear on your Defense Center. You can also design your own custom reports from scratch. Reporting is available only on Defense Centers.

A report is a document file formatted in PDF, HTML, or CSV with the content you want to communicate. A report template specifies the data searches and formats for the report and its sections. The Sourcefire 3D System includes a powerful report designer that automates the design of report templates. You can replicate the content of any event view table or dashboard graphic displayed in the web interface.

You can build as many report templates as you need. Each report template defines the individual sections in the report and specifies the database search that creates the report's content, as well as the presentation format (table, chart, detail view, and so on) and the time frame. Your template also specifies document attributes, such as the cover page and table of contents and whether the document pages have headers and footers (available only for reports in PDF format). You can export a report template in a single object package file and import it for reuse on another Defense Center.

You can include input parameters in a template to expand its usefulness. Input parameters allow you to produce tailored variations of the same report. When you generate a report with input parameters, the generation process prompts you to enter a value for each input parameter. The values you type constrain the report contents on a one-time basis. For example, you can place an input parameter in the destination IP field of the search that produces an intrusion event report; at report generation time, you can specify a department's network segment when

prompted for the destination IP address. The generated report then contains only information concerning that particular department.

See the following sections for more information on reports and report templates:

- [Generating Reports](#) on page 1725
- [Using Report Templates](#) on page 1736
- [Managing Report Templates and Report Files](#) on page 1766

Generating Reports

LICENSE: Any

The Sourcefire 3D System's reporting feature allows you to quickly capture the content of any event view, dashboard, or workflow from your Defense Center and present it in report format.

For more information, see the following sections:

- [Creating a Report Template from an Event View](#) on page 1725
- [Creating a Report Template by Importing a Dashboard or Workflow](#) on page 1727
- [Generating Reports from a Report Template](#) on page 1729
- [Using Report Generation Options](#) on page 1732

Creating a Report Template from an Event View

LICENSE: Any

Before you generate a report, the reporting system creates a report template that you can modify to meet your needs. You can add additional sections, modify automatically included sections, and delete sections.

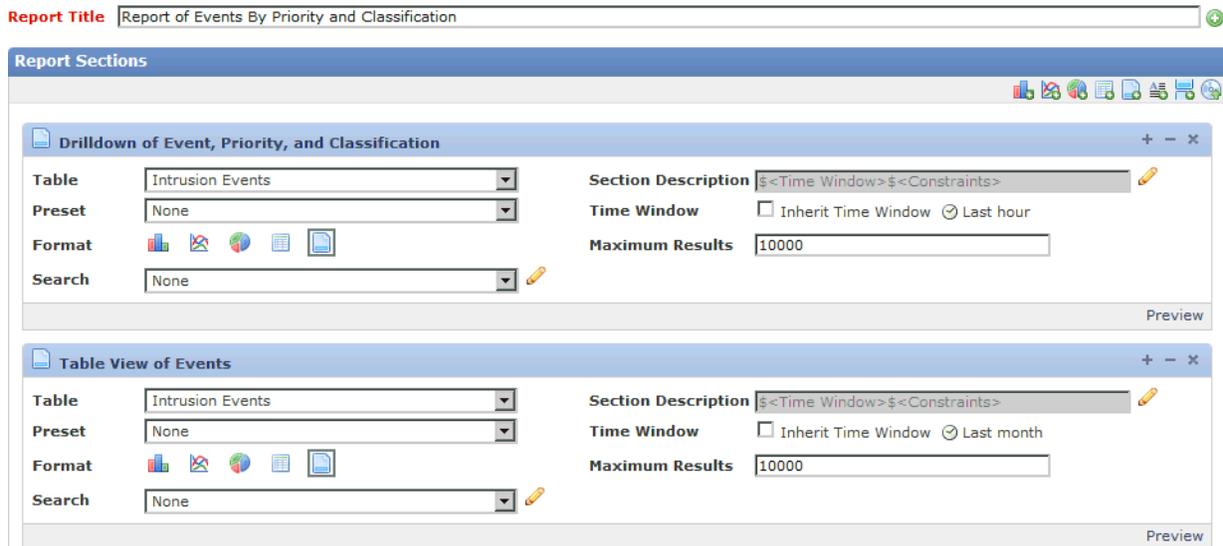
To create a report template from an event view:

ACCESS: Admin/Any Security Analyst

1. Populate an event view with the events you want in the report. You can do this in various ways:
 - Use an event search to define the events you want to view. For details on using the event search, see [Searching for Events](#) on page 1770.
 - Drill down through a workflow until you have the appropriate events in your event view. For details on workflows and how to constrain events within a workflow, see [Understanding and Using Workflows](#) on page 1791.

2. From the event view page, click **Report Designer**.

The Report Sections page appears with a section for each view in the captured workflow. The following graphic shows the first two sections of a report template built from an intrusion event workflow.



3. Optionally, type a new name in the **Report Title** field and click **Save**.
4. Optionally, delete any template sections that you want to exclude from the report by clicking the delete icon (**X**) in the section's title bar, and confirm the deletion.

The deleted sections disappear.

IMPORTANT! The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the Defense Center.

5. Optionally, adjust the settings of the fields in your report sections. For details on configuring the fields in a report section, see [Editing the Sections of a Report Template](#) on page 1743.

TIP! To view the current column layout or chart formatting for a section, click the section's **Preview** link.

6. Optionally, change the title of any section by clicking the section title in the title bar. The Set Section Title pop-up window appears. Type the section title and click **OK**.

7. Optionally, add page breaks. Click the add page break icon ().
A new page break object appears at the bottom of the template. Drag it in front of the section that should start the new page. For information on using page breaks, see [Editing the Sections of a Report Template](#) on page 1743.
8. Optionally, add text sections. Click the add text section icon ().
A new text section appears at the bottom of the template. Drag it to the place where it should appear in the report template. For information about editing a text section, see [Editing the Sections of a Report Template](#) on page 1743.

TIP! Text sections, which support rich text (bold, italics, variable font size, and so on) as well as imported images, are useful for introductions to your report or your report sections.

9. Optionally, click **Advanced Settings** to add a cover page, table of contents, starting page number, or header and footer text. For more information, see [Editing Document Attributes in a Report Template](#) on page 1756.
10. When the report template is correct, click **Save**.
The report template is saved and an entry for the report template appears on the Report Templates page.

Creating a Report Template by Importing a Dashboard or Workflow

LICENSE: Any

You can quickly create a new report by importing dashboards, workflows, and statistics summaries. The import creates a section for each widget graphic in your dashboard and each event view in your workflow. You can delete any unnecessary sections to focus on the most important information. The [Data Source Options on Import Report Sections Window](#) table describes the import options.

Data Source Options on Import Report Sections Window

SELECT THIS OPTION...	TO IMPORT...
Import Dashboard	any custom analysis widget on the selected dashboard.

Data Source Options on Import Report Sections Window (Continued)

SELECT THIS OPTION...	TO IMPORT...
Import Workflow	any predefined or custom workflow. TIP! Selections have the format: Table - workflow name For example, Connection Events - Traffic by Port imports the views in the Traffic by Port workflow generated from the Connection Events table.
Import Summary Sections	any of the following generic summaries: <ul style="list-style-type: none"> • Intrusion Detailed Summary • Intrusion Short Summary • Discovery Detailed Summary • Discovery Short Summary

To create a report template from a dashboard, workflow, or statistics summary:

ACCESS: Admin/Any Security Analyst

1. Identify the dashboard, workflow, or summary you want to replicate in your report.
2. Select **Overview > Reporting**.
3. Click the **Report Templates** tab.
The Report Templates page appears.
4. Click **Create Report Template**.
The Report Sections page appears.
5. Type a name for your new report template in the **Report Title** field.
6. Click **Save** to save the report template under the new name.
The report template is saved and an entry for the report template appears on the Report Templates page.
7. Click the import sections from dashboard, summaries and workflows icon ().
The Import Report Sections pop-up window appears. You can choose any of the data sources described in the [Data Source Options on Import Report Sections Window](#) table.
8. Select a dashboard, workflow or summary from the drop-down menus.

9. For the data sources you want to add, click **Import**.
The Report Sections page for your template reappears with a section for each element of the selected data source. For dashboards, each widget graphic will have its own section; for workflows, each event view will have its own section.
10. Make changes to the content of your sections as needed.
For information on editing a report template, see [Editing the Sections of a Report Template](#) on page 1743.

IMPORTANT! The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the Defense Center.

11. When the report template is correct, click **Save**.
The report template is saved and an entry for the report template appears on the Report Templates page.

Generating Reports from a Report Template

LICENSE: Any

After you define your report template, you are ready to generate the report itself. The generation process lets you select the report's format (HTML, PDF, or CSV). You can also adjust the report's global time window, which applies a consistent time frame to all sections except those you exempt. For information on setting the report time window, see [Setting the Time Window for a Template and Its Sections](#) on page 1747.

If the report template includes user input parameters in its search specification, the generation process prompts you to enter values, which tailor this run of the report to a subset of the data. For information on input parameters, see [Using Input Parameters](#) on page 1750.

The Reports tab lists all previously generated reports, with report name, date and time of generation, generating user, and whether the report is stored locally or remotely. A status column indicates whether the report is already generated, is in

the generation queue (for example, for scheduled tasks), or failed to generate (for example, due to lack of disk space).

Reports		Report Templates				
<input type="checkbox"/>	Name	Time Requested	Time Completed	User	Location	Status
<input type="checkbox"/>	Report_of_Events_By_Priority-20111027192647.pdf Reports	2011-10-27 15:26:47	2011-10-27 15:26:47	admin	Local	Successfully Processed
<input type="checkbox"/>	Report_of_Events_By_Priority-20111027184716.zip Reports	2011-10-27 14:47:16	N/A	admin	Local	In queue
<input type="checkbox"/>	Report_of_Events_By_Priority-20111027184204.pdf Reports	2011-10-27 14:42:04	2011-10-27 14:42:05	admin	Local	Successfully Processed
<input type="checkbox"/>	Report_of_Events_By_Priority-20111027184204.zip Reports	2011-10-27 14:42:04	2011-10-27 14:42:05	admin	Local	Successfully Processed

Storage Location: /var/sf/reports/ (Disk Usage: 1%)

The Reports tab page shows all locally stored reports. It shows remotely stored reports as well, if remote storage is currently configured. The location of your currently configured report storage appears at the bottom of the page, with disk usage for local, NFS, and SMB storage. If you access remote storage using SSH, disk usage data is not available. For information on setting up remote storage, see [Using Remote Storage for Reports](#) on page 1765.

IMPORTANT! If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Use the following procedure to generate and view reports. Note that users with Administrator access can view all reports; other users can view only the reports they generated.

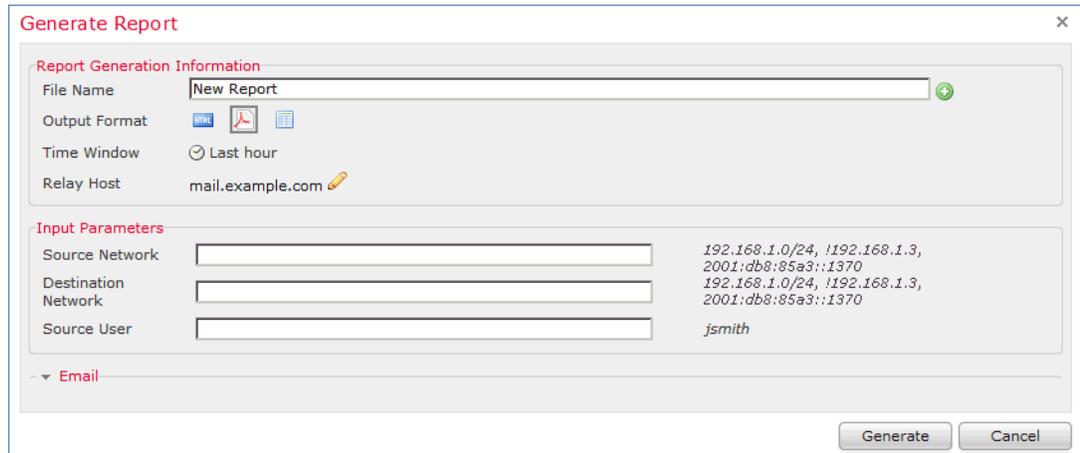
To generate a report from a report template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Click the **Report Templates** tab.

The Report Templates page appears.

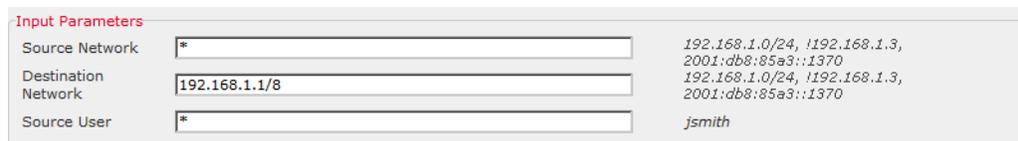
- Click the generate report icon () for the template you want to use. The Generate Report pop-up dialog appears.



- Optionally, type a new name in the **File Name** field. This sets the name of the generated report file. You may also use the input parameter icon () to add one or more input parameters to the file name. For information on input parameters, see [Using Input Parameters](#) on page 1750.
- Select the output format for the report by clicking the corresponding icon: HTML, PDF, or CSV.
- Optionally, change the global time window by clicking the time window icon (). The Events Time Window pop-up window appears. For information on setting the events time window, see [Setting Event Time Constraints](#) on page 1820.

IMPORTANT! Setting the global time window affects the content of individual report sections only if they are configured to inherit the global setting. For information on report section inheritance of the global time window, see [Setting the Time Window for a Template and Its Sections](#) on page 1747.

- Type values for any fields that appear in the **Input Parameters** section.



TIP! You can ignore user parameters by typing the * wildcard character in the field. This eliminates the user parameter's constraint on the search.

8. Optionally, if an email relay host is configured in your system policy, click **Email** to automate email delivery of the report when it generates. For details about email delivery features, see [Distributing Reports by Email at Generation Time](#) on page 1763.
9. Click **OK** and confirm when prompted.
The Report Generation Complete pop-up window appears with a link to view your report.



10. Click either:
 - the report link, which opens a new window to display the report, or
 - **OK** to return to the Report Section page, where you can modify your report design.

You can review completed reports after initial generation, as well.

To view a generated report:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Click the **Reports** tab.
The Reports page appears.
3. Click the name of the report.
The default program on your local host opens the report in a new window.
4. When finished viewing the document, use your browser to return to the **Reports** tab.

Using Report Generation Options

LICENSE: Any

You have the following options when generating reports:

- schedule your reports to be automatically generated — see [Generating Reports Using the Scheduler](#) on page 1763
- automatically distribute new reports by email — see [Distributing Reports by Email at Generation Time](#) on page 1763
- store your report files on a remote host — see [Using Remote Storage for Reports](#) on page 1765

Managing Reports

LICENSE: Any

You have the following options for managing the report files you have generated:

- saving report files to your local computer — see [Downloading Reports](#) on page 1768
- deleting report files — see [Deleting Reports](#) on page 1769

Understanding Report Templates

LICENSE: Any

You use report templates to define the content and format of the data in each of the report's sections, as well as the document attributes of the report file (cover page, table of contents, and page headers and footers). After you generate a report, the template stays available for reuse until you delete it.

Your reports contain one or more information sections. You choose the format (text, table, or chart) for each section individually. The format you select for a section may constrain the data that can be included. For example, you cannot show time-based information in certain tables using a pie chart format. You can change the data criteria or format of a section at any time to obtain optimum presentation.

You can base a report's initial design on a predefined event view, or you can start your design by importing content from any defined dashboard, workflow, or summary. You can also start with an empty template shell, adding sections and defining their attributes one by one.

All sections in a report template have a title bar and various attribute fields that control the section's contents and appearance. For more information, see the following:

- the [Report Section Title Bar Elements table](#) on page 1734
- the [Report Section Fields table](#) on page 1734

The [Report Section Title Bar Elements](#) table describes the controls on the title bar for each template section.

Report Section Title Bar Elements

ATTRIBUTE	DEFINITION
title bar name	Contains the name of the section as it appears in the report. Change it by clicking it and typing a new name.
title bar icons	Click the duplicate icon (+) to add a duplicate of the section to the report template. Click the minimize icon (-) to minimize the section. Click the delete icon (✕) to delete the section, after confirmation.

The [Report Section Fields](#) table defines the fields in each section of a report template.

Report Section Fields

FIELD NAME	DEFINITION
Table	Presents a drop-down menu that allows you to select the table from which the section data is extracted.
Preset	Presents a drop-down menu of predefined searches. You can select an appropriate preset to initialize the search criteria when you define a new search.

Report Section Fields (Continued)

FIELD NAME	DEFINITION
Format	<p>Presents icons that allow you to select the format of the section data. Options include:</p> <p> Bar chart: Compares quantities of the selected variables.</p> <p> Line chart: Shows trends/changes over time of a selected variable. Available only for time-based tables.</p> <p> Pie chart: Shows each selected variable as a percentage of the whole. Variables with quantities of zero are dropped from the chart. Very small quantities are clustered into a category labeled Other.</p> <p> Table view: Shows values of attributes for each record. Not available for summary or statistical data.</p> <p> Detail view: Shows complex object data associated with certain events, such as packets (for intrusion events) and host profiles (for host events). Format is available only for certain event types that involve such objects. Output may degrade performance if large numbers are requested.</p>
Search or Filter	<p>Presents a drop-down menu of searches or application filters.</p> <p>For most tables, you can constrain a report using a predefined or saved Search. You can also create a new search by clicking the edit icon (); see Working with Searches in Report Template Sections on page 1749.</p> <p>For the Application Statistics table, you use a user-defined application Filter to constrain a report; for information on creating filters, see Working with Application Filters on page 176.</p>
X-Axis	<p>Presents a drop-down menu of available data columns for the X-axis of the selected chart. Appears only when you select a chart format. For line charts, the X-axis value is always Time. For bar and pie charts, you cannot select Time as the X-axis value.</p>
Y-Axis	<p>Presents a drop-down menu of available data columns for the Y-axis of the selected chart.</p>

Report Section Fields (Continued)

FIELD NAME	DEFINITION
Section Description	Defines the descriptive text that precedes the search data in the section. Enter a combination of text and input parameters. Default for a new section is the following set of two input parameters: \$<Time Window> and \$<Constraints>. For more information on input parameters, see Using Input Parameters on page 1750.
Time Window	Defines the time window for the data that appears in the section. If the section searches time-based tables, you can select the check box to inherit the report's global time window. Alternatively, you can set a specific time window for the section. For information about setting the time window, see Editing the Sections of a Report Template on page 1743.
Results	Select either Top or Bottom and enter the maximum number of records to be included in the section.
Color	Defines the colors for graphed data in the section. Select one or more colors, as applicable.

Using Report Templates

LICENSE: Any

You can build a new report template in any of the following ways:

- [Creating a Report Template from an Event View](#) on page 1725
- [Creating a Report Template by Importing a Dashboard or Workflow](#) on page 1727
- [Creating Report Templates from Existing Templates](#) on page 1737
- [Creating New Report Templates](#) on page 1740

After you have populated your template with sections, you can edit the document attributes and section attributes as necessary. For information, see:

- [Editing Document Attributes in a Report Template](#) on page 1756
- [Editing the Sections of a Report Template](#) on page 1743

The key to successful reports is defining the searches that populate the report's sections. For information on defining these searches, see [Working with Searches in Report Template Sections](#) on page 1749.

Creating Report Templates from Existing Templates

LICENSE: Any

You can build a new report template by copying any existing template. If you identify a good model among your existing templates, you can copy it and edit its attributes as necessary.

TIP! Sourcefire provides a set of predefined report templates. You can see their names in the list of templates on the **Report Templates** tab. For a description of their attributes, see [Using Sourcefire Predefined Report Templates](#) on page 1737.

To create a report template from an existing template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Click the **Report Templates** tab.
The Report Templates page appears. For information about Sourcefire-provided report templates, see [Using Sourcefire Predefined Report Templates](#) on page 1737.
3. Click the copy icon () next to the report template you want to copy as a model.
The copied template appears as a new report template.
4. In the **Report Title** field, type a name for your new report template.
5. Click **Save**.
The report template is saved and an entry for the new report template appears on the Report Templates page.
6. Make changes to the template as needed.
For information on defining the sections of the template and the document attributes, see:
 - [Editing the Sections of a Report Template](#) on page 1743
 - [Editing Document Attributes in a Report Template](#) on page 1756

Using Sourcefire Predefined Report Templates

LICENSE: Any

The following Sourcefire-predefined report templates are included with the Sourcefire 3D System:

- [Host Report: \\$<Host>](#)
- [User Report: \\$<User>](#)
- [Attack Report: Attack \\$<Attack SID>](#)
- [Malware Report](#)

- [Sourcefire FireSIGHT Report: \\$<Customer Name>](#)
- [Files Report](#)

You can use these templates as is, edit them, or use them as the basis for your own templates.

Host Report: \$<Host>

The Host Report: \$<Host> report template provides information about a specific host on the network. This report template contains the following sections:

- Server Applications
- Client Applications
- Intrusion Events Originating from This Host
- Intrusion Events Destined to This Host
- Connections Originating from This Host
- Connections Destined to This Host
- Users of This Host
- White List Violations by This Host

User Report: \$<User>

The User Report: \$<User> report template provides information about a specific user on the network. This report template contains the following sections:

- Client Applications Used by This User
- Web Applications Used by This User
- Application Protocols Used by This User
- Comprehensive List of Applications Used by This User
- Intrusion Events Originated By This User's Machines
- Intrusion Events Destined to This User's Machines
- Connections Originating from This User's Machines
- Connections Destined to This User's Machines
- Hosts for This User

Attack Report: Attack \$<Attack SID>

The Attack Report: Attack \$<Attack SID> report template provides information about a specific attack on the network. This report template contains the following sections:

- General Information About This Attack
- Number of Attacks
- Number of Machines Initiating Attack
- Number of Machines Being Attacked

- Sources of This Attack
- Destinations of This Attack
- Traffic Patterns of This Attack

Malware Report

The Malware Report report template provides information about network and endpoint-based malware events. This report template contains the following sections:

- Malware Threats
- Threat Detections over Time
- Application Protocols Transferring Malware
- Hosts Receiving Malware
- Hosts Sending Malware
- Users Affected by Malware
- Malware Intrusions
- File Types Infected with Malware
- Applications Introducing Malware
- Table View of Malware Events

Note that neither Series 2 devices nor the DC500 Defense Center support network-based malware protection, which can affect the data detected and displayed. For example, a Series 3 Defense Center managing only Series 2 devices can display only endpoint-based malware events.

Sourcefire FireSIGHT Report: \$<Customer Name>

The Sourcefire FireSIGHT Report: \$<Customer Name> report template provides overall information about an organization's network. This report template contains the following sections:

- Summary of Application Traffic by Risk
- Risky Applications with Low Business Relevance
- Users of Risky Applications
- Anonymizers and Proxies
- Typically High Bandwidth Applications
- Applications by Total Bandwidth
- Hosts Accessing Sensitive Network
- Users Accessing Sensitive Network
- Applications on Sensitive Network
- Ports and Protocols Related to Sensitive Network
- Hosts Visiting Malicious URLs

- Users Visiting Malicious URLs
- Granular Application Usage
- Web Applications
- Client Applications
- Application Protocols
- Web Browser Versions
- Operating System Versions
- Overall User Activity
- Intrusion Events by Impact
- Intrusion Events by Impact (After Blocking)
- Intrusion Events by Application
- Top Intrusion Events
- Comprehensive Application List

Files Report

The Files Report report template provides information about files detected in network traffic by managed devices. This report template contains the following sections:

- File Transfers over Time
- Application Protocols Used by File Transfers
- File Dispositions
- File Actions
- Hosts Receiving Files
- Hosts Sending Files
- Users Transferring Files
- File Categories
- File Types
- File Names
- Table View of File Events

Creating New Report Templates

LICENSE: Any

If you do not want to copy an existing report template, you can create an entirely new template. First, you create a default template shell. Then, in the order you

prefer, you design the individual template sections and set attributes for the report document. For information on these steps, see the following sections:

- [Creating a Template Shell](#) on page 1741
- [Configuring the Content of the Template Sections](#) on page 1742
- [Setting Attributes for PDF and HTML Report Documents](#) on page 1742

Creating a Template Shell

LICENSE: Any

A report template is a framework of sections, each independently built from its own database query. The first step in creating a template is to generate the framework shell that allows you to add and format the sections.

To create a template shell:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Click the **Report Templates** tab.
The Report Templates page appears.
3. Click **Create Report Template**.
The Report Sections page appears with the default template name, **New Report**, in the **Report Title** field.
4. Optionally, type a name for your new template in the **Report Title** field and click **Save**. The report title can contain any combination of alphanumeric characters and spaces.
An entry with the new template name appears on the Report Templates page list.
5. The report title can also contain input parameters. To add an input parameter, place your cursor at the spot in the title where the value of the parameter should appear, then click the insert input parameter icon (⊕). For information on input parameters, see [Using Input Parameters](#) on page 1750.
The added input parameters appear in the **Report Title** field.
6. Use the set of add icons under the Report Sections title bar to insert section shells as necessary. For information on section formatting, see the [Report Section Fields table](#) on page 1734.
Each added section appears at the bottom of the template. Drag it to the correct location.
7. Click the section title on the section title bar and type a name for the section.
8. Click **Save** to save the template.
Your template is saved.

Configuring the Content of the Template Sections

LICENSE: Any

Each template section consists of a dataset generated by a search or filter, and has a format specification (table, pie chart, and so on) that determines the mode of presentation. You further determine section content by selecting the fields in the data records you want to include in the output, as well as the time frame and number of records to show.

To configure report template sections:

ACCESS: Admin/Any Security Analyst

1. Edit the section attributes as described in the [Editing the Sections of a Report Template table](#) on page 1743.
2. Optionally, click **Preview** at the bottom of the section window to view the column layout or graphic format you selected.

IMPORTANT! Use the section preview utility to check the column selection and output characteristics such as pie chart colors. It is not a reliable indicator of the correctness of your configured search.

Setting Attributes for PDF and HTML Report Documents

LICENSE: Any

The report you generate from the template has several document attributes that span all sections and control features, such as the cover page, headers and footers, page numbering, and so on.

To set attributes for the report documents:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Click the **Report Templates** tab.
The Report Templates page appears.
3. Click **Edit** for the report template you want to use to generate the report.
The Report Sections page for your template appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears.
5. For documents in PDF or HTML format, perform the tasks described in [Editing Document Attributes in a Report Template](#) on page 1756.
If you selected CSV as your document format, you have no document attributes to set.

Editing the Sections of a Report Template

LICENSE: Any

You can modify a variety of report section attributes to adjust the content of the section and its data presentation. For information, see the following sections:

- [Setting the Table and Data Format for a Template Section](#) on page 1743
- [Specifying the Search or Filter for a Template Section](#) on page 1744
- [Setting the Search Fields that Appear in Table Format Sections](#) on page 1745
- [Adding a Text Section to a Report Template](#) on page 1745
- [Adding a Page Break to a Report Template](#) on page 1746
- [Setting the Time Window for a Template and Its Sections](#) on page 1747
- [Renaming a Template Section](#) on page 1748
- [Previewing a Template Section](#) on page 1748

IMPORTANT! Security Analysts can edit only report templates they created.

Setting the Table and Data Format for a Template Section

LICENSE: Any

Each section in a report template queries a database table to generate content for that section. Changing the section's data format uses the same data query, but modifies the fields that appear in the section according to the analytical purpose of the format type. For example, the table view of intrusion events populates the section with a large number of data fields per event record, while a pie chart section shows the portion of all matching records that each selected attribute represents, with no details about individual events. Bar chart sections compare the total counts of matching records that have specific attributes. Line charts summarize changes in the matching records over time with respect to a single attribute. Line charts are available only for data that is time-based, not for information about hosts, users, third-party vulnerabilities, and so on.

For information on the different available formats, see the [Report Section Fields table](#) on page 1734.

To select the table and output format for a template section:

ACCESS: Admin/Any Security Analyst

1. Use the **Table** drop-down menu to select the table to query in this section. Icons appear in the **Format** field for each of the output formats available for the selected table.
2. Select the applicable output format icon for the section. For information about these formats, see the [Report Section Title Bar Elements table](#) on page 1734. The fields included in the output appear.

3. To change the search constraints, click the edit icon () next to the **Search** or **Filter** field.

The Search Editor pop-up window appears with options for constraining the search. For information on using this window, see [Working with Searches in Report Template Sections](#) on page 1749.

4. For graphic output formats (pie chart, bar chart, and so on), adjust the **X-Axis** and **Y-Axis** parameters using the drop-down menus.

When you select a value for the X-axis, only compatible values appear in the Yaxis drop-down menu, and vice versa.

5. For table output, select the columns, order of appearance, and sort order in your output. For detailed information, see [Setting the Search Fields that Appear in Table Format Sections](#) on page 1745.

6. Click **Save** to save the template.

Your template is saved.

Specifying the Search or Filter for a Template Section

LICENSE: Any

The search or filter in a report section specifies the database query on which the section content is based. For most tables, you can constrain a report using a predefined or saved search, or you can create a new search on the fly:

- Sourcefire predefined searches serve as examples for searching certain event tables and can provide quick access to important information about your network that you may want to include in reports.
- Saved event searches include all public event searches that you or others have created, plus all your saved private event searches. For information on defining, naming, and using saved event searches, see [Searching for Events](#) on page 1770.
- Saved searches for the current report template are accessible only in the report template itself. The search names of saved report template searches end with the string "Custom Search." Users create these searches while designing reports.

For the Application Statistics table, you use a user-defined application filter to constrain a report; for information on creating filters, see [Working with Application Filters](#) on page 176.

To specify a search or filter for a template section:

ACCESS: Admin/Any Security Analyst

1. Select the database table to query from the **Table** drop-down menu:
 - For most tables, the **Search** drop-down list appears.
 - For the Application Statistics table, the **Filter** drop-down list appears.

2. Select the search or filter you want to use to constrain the report.
You can view the search criteria or create a new search by clicking the edit icon (✎). For more information, see [Working with Searches in Report Template Sections](#) on page 1749.

Setting the Search Fields that Appear in Table Format Sections

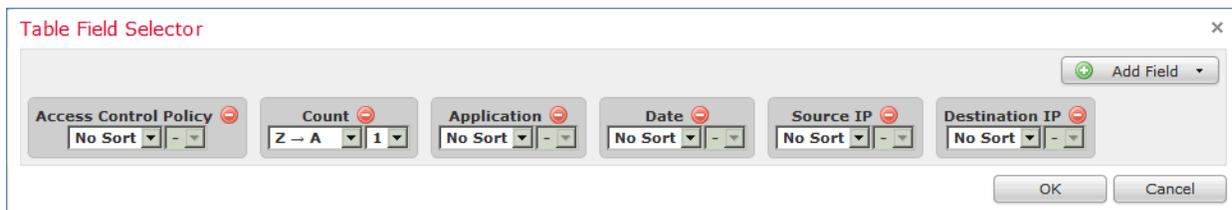
LICENSE: Any

If you include table data in a section, you can choose which fields in the data record to show. All fields in the table are available for inclusion or exclusion. You select fields that accomplish the purpose of the report, then order and sort them accordingly.

To add and delete the fields in a table format section:

ACCESS: Admin/Any Security Analyst

1. For table format sections, click the edit icon (✎) next to the **Fields** parameter. The Table Field Selector window appears.



2. Optionally, add and delete fields, and drag the field icons into the column order you want.
3. Optionally, change the sort order of any column. Use the drop-down lists on each field icon to set the sort order and priority.
4. When the fields are in the right order and have the necessary sort characteristics, click **OK**.
The Report Sections page appears.

Adding a Text Section to a Report Template

LICENSE: Any

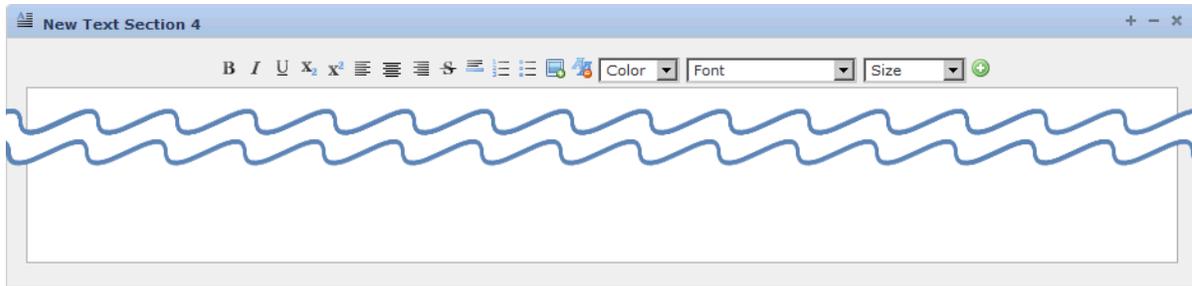
You can add text sections to your templates to provide custom text, such as an introduction, for the whole report or for individual sections. Text sections can have rich text with multiple font sizes and styles (bold, italic, and so on) as well as input parameters and imported images. For information on input parameters, see [Using Input Parameters](#) on page 1750.

To add a text section to a report template:

ACCESS: Admin/Any Security Analyst

1. Click the add text section icon ().

A text section appears at the bottom of the template.



2. Drag the new text section to its intended position in the report template.
3. Optionally, add page breaks before and after the text section. For information on page breaks, see [Adding a Page Break to a Report Template](#) on page 1746.
4. Optionally, click the text section's generic name in the title bar to type a new name.
5. Add formatted text and images to the body of the text section. You can include input parameters that dynamically update when you generate the report.
6. Click **Save** when finished.
Your template is saved.

Adding a Page Break to a Report Template

LICENSE: Any

You can add page breaks before or after any section in the template. This feature is particularly helpful for multi-section reports with text pages that introduce the various sections.

To add a page break:

ACCESS: Admin/Any Security Analyst

1. Click the add page break icon ().

A page break appears at the bottom of the template.

2. Drag the page break to its intended location, before or after a section.
3. Repeat the process for all page breaks you add to the template.

Setting the Time Window for a Template and Its Sections

LICENSE: Any

A report template's time window defines the template's reporting period. Report templates with time-based data (such as intrusion or discovery events) have a global time window, which the time-based sections in the template inherit by default when created. Changing the global time window changes the local time window for the sections that are configured to inherit the global time window. You can disable time window inheritance for an individual section by clearing its **Inherit Time Window** check box. You can then edit the local time window.

IMPORTANT! Global time window inheritance applies only to report sections with data from time-based tables, such as intrusion events and discovery events. For sections that report on network assets (hosts and devices) and related information (such as vulnerabilities), you must set each time window individually.

To change a report template's global time window:

ACCESS: Admin/Any Security Analyst

1. On the Report Templates page, click the edit icon (✎) next to the report template you want to edit.
The Report Sections page appears.
2. Click **Generate**.
The Generate Report pop-up window appears.
3. To modify the global time window, click the time window icon (🕒).
The Events Time Window page appears in a new window. For information about using this page, see [Setting Event Time Constraints](#) on page 1820.
4. When you are finished, click **Apply** on the Events Time Window.
The Generate Report pop-up window reappears with the new time window.
5. Click **Cancel** to return to the Report Sections page, or **OK** to generate the report.

Your report can have different time ranges per section. For example, your first section could be a summary for the month, and the remaining sections could drill down into details at the week level. In such cases, you set the section-level time windows individually.

To configure a section's local time window:

ACCESS: Admin/Any Security Analyst

1. On the Report Sections page of a template, clear the **Inherit Time Window** check box for the section if it is present.
The local section time window icon appears.

2. To change the section's local time window, click the time window icon (🕒). The Events Time Window page appears. For information about using this page, see [Setting Event Time Constraints](#) on page 1820.

IMPORTANT! Sections with data from statistics tables can have only sliding time windows.

3. When you have set a new local time window, click **Apply** on the Events Time Window.
4. Click **Save**.
The Report Sections page appears for further editing.

Renaming a Template Section

LICENSE: Any

When you create a new template, the sections you add receive generic section names and should be renamed to indicate their content.

To rename a template section:

ACCESS: Admin/Any Security Analyst

1. Click the current section name in the section header.
The Set Section Title pop-up window appears.
2. Type a new name for the section and click **OK**.
The name in the section title bar is changed.

Previewing a Template Section

LICENSE: Any

The preview function shows the field layout and sort order for table views and important legibility characteristics of graphics, such as pie chart colors.

To preview a template section:

ACCESS: Admin/Any Security Analyst

1. At any time while editing a section, click **Preview** for the section.
The Preview pop-up window appears.
2. Close the preview by clicking **OK** at the bottom of the window.
The Report Sections page appears.

Working with Searches in Report Template Sections

LICENSE: Any

The Sourcefire 3D System provides a search editor to view the searches available in your report templates and to define new custom searches.

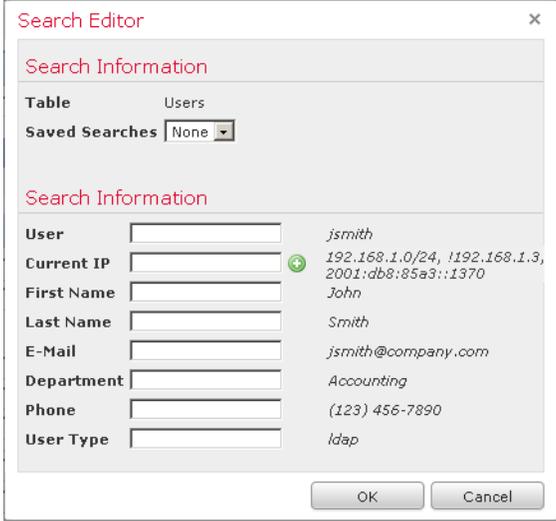
TIP! The custom searches you make in a report template are specific to the template where you create them. You can make searches that are reusable across all report templates in the event viewer. When you save a custom search in the event viewer, it appears in the **Search** drop-down menu of all report templates. For details on using the event viewer to create and save custom searches, see [Searching for Events](#) on page 1770.

To create a custom search:

ACCESS: Admin/Any Security Analyst

1. From the relevant section in the report template, click the edit icon () next to the **Search** field.

The Search Editor page appears with the table to be searched selected.



Search Information	
Table	Users
Saved Searches	None
Search Information	
User	jsmith
Current IP	192.168.1.0/24, 1192.168.1.3, 2001:db8:85a3::1370
First Name	John
Last Name	Smith
E-Mail	jsmith@company.com
Department	Accounting
Phone	(123) 456-7890
User Type	ldap

2. Optionally, from the **Saved Searches** drop-down menu, select a predefined search.

The drop-down presents all available predefined searches for this table, including system-wide and report-specific predefined searches.

3. Edit the search criteria in the appropriate fields. For certain fields, your constraints can include the same operators (<, >, and so on) as event searches. For the syntax of search criteria, see [Searching for Events](#) on page 1770.

If you enter multiple criteria, the search returns only the records that match all the criteria.

4. Optionally, where the input parameter icon (⊕) appears, you can insert an input parameter from the drop-down menu instead of typing a constraint value. For information on using input parameters in report designs, see [Using Input Parameters](#) on page 1750.

For some search fields, the drop-down menu may contain user-defined managed objects instead of, or with, input parameters. Managed objects, which have distinctive icons depending on their type, are system configuration variables you can use as values in constraining searches. However, they do not produce the generation-time query for user input that occurs with input parameters. For information on managed objects, see [Using Objects and Security Zones](#) on page 158.

IMPORTANT! When you edit the constraints of a reporting search, the system saves your edited search under the following name: *section custom search*, where *section* is the name in the section title bar followed by the string *custom search*. To have meaningful names for your saved custom searches, be sure you change the section name before you save the edited search. You cannot rename a saved reporting search.

5. When finished modifying the fields in the search editor, click **OK**.
The Report Sections page reappears and a new predefined search appears in the section's **Search** drop-down menu.

Using Input Parameters

LICENSE: Any

You can use input parameters in a report template that the report can dynamically update at generation time. The input parameter icon (⊕) indicates the fields that can process them. There are two kinds of input parameters:

- Predefined — see the [Predefined Input Parameters](#) on page 1751
- User-defined — see the [User-Defined Input Parameter Types table](#) on page 1753

Predefined Input Parameters

LICENSE: Any

Predefined input parameters are resolved by internal system functions or configuration information. For example, at report generation time, the system

replaces the `<Time>` parameter with the current date and time. The [Predefined Input Parameters](#) table defines the parameters available for use. You might, for example, include `<Month>` in the title of a monthly summary report that generates automatically under scheduler control. Your report title then automatically updates with the correct month.

Predefined Input Parameters

INSERT THIS PARAMETER...	...TO INCLUDE THIS INFORMATION IN YOUR TEMPLATE:
<code><Logo></code>	The selected uploaded logo
<code><Report Title></code>	The report title
<code><Time></code>	The date and time of day the report ran, with one-second granularity
<code><Month></code>	The current month
<code><Year></code>	The current year
<code><System Name></code>	The name of the Defense Center
<code><Model Number></code>	The model number of the Defense Center
<code><Time Window></code>	The time window currently applied to the report section
<code><Constraints></code>	The search constraints currently applied to the report section

The [Predefined Input Parameter Usage](#) table lists the valid input parameters that can be used in different areas within the Report Templates page.

Predefined Input Parameter Usage

PARAMETER	REPORT TEMPLATE COVER PAGE	REPORT TEMPLATE REPORT TITLE	REPORT TEMPLATE SECTION DESCRIPTION	REPORT TEMPLATE TEXT SECTION	GENERATE REPORT FILE NAME	GENERATE REPORT EMAIL SUBJECT, BODY
<code><Logo></code>	X					
<code><Report Title></code>	X		X	X	X	X
<code><Time></code>	X	X	X	X	X	X

Predefined Input Parameter Usage (Continued)

PARAMETER	REPORT TEMPLATE COVER PAGE	REPORT TEMPLATE REPORT TITLE	REPORT TEMPLATE SECTION DESCRIPTION	REPORT TEMPLATE TEXT SECTION	GENERATE REPORT FILE NAME	GENERATE REPORT EMAIL SUBJECT, BODY
\$<Month>	X	X	X	X	X	X
\$<Year>	X	X	X	X	X	X
\$<System Name>	X	X	X	X	X	X
\$<Model Number>	X	X	X	X	X	X
\$<Time Window>			X			
\$<Constraints>			X			

User-Defined Input Parameters

LICENSE: Any

You can create your own input parameters to supply as constraints in section searches. Constraining a search with an input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically tailor a report at generation time to show a particular subset of data without changing the template. For example, you can provide an input parameter for the **Destination IP** field of a report section's search. Then, when you generate the report, you can type the IP network segment for a particular department to get data for that department only.

TIP! You can also type * in an input parameter field, with the effect of ignoring the constraint.

You can also define string-type input parameters to add dynamic text in certain fields of your report, such as in emails (subject or body), report file names, and text sections. You can personalize reports for different departments, with customized report file names, email addresses, and email messages, using the same template for all.

Each input parameter you define has a name and a type. The [User-Defined Input Parameter Types](#) table describes the parameter types.

User-Defined Input Parameter Types

USE THIS PARAMETER TYPE...	WITH FIELDS WITH THIS DATA...
Network/IP	any IP address or network segment in CIDR format
Application	name of an application protocol, client application, or web application
Event Message	any event view message
Device	3D appliance (Defense Center or Sourcefire managed device)
Username	user identification such as initiator user and responder user
Number (VLAN ID, Snort ID, Vuln ID)	any VLAN ID, Snort ID, or vulnerability ID
String	text fields such as application or OS version, notes, or descriptions

An input parameter’s type determines the search fields where you can use it. You can use a given type only in appropriate fields, as described in the [User-Defined Input Parameter Types](#) table. For example, a user parameter you define as a string type is available for insertion in text fields but not in fields that take an IP address. In the graphic below, the input parameters **Source Network** and **Destination Network** are both of type **Network/IP** and therefore appear as options for search fields, such as **Source IP** and **Destination IP**, that take IP addresses.



To create user-defined input parameters for a report template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the edit icon (✎) for the template you want to edit.
The Report Sections page appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears.
5. Click the Add Input Parameter icon (➕).
The Add Input Parameter pop-up window appears.
6. Type the parameter name in the **Name** field and use the **Type** drop-down menu to select the type, then click **OK**.
The new parameter appears in the **Input Parameters** menu.
7. Repeat the steps above until you have defined all the parameters you need.
8. Click **OK**.
Your new input parameters are saved for this template and the Report Sections page reappears.

If you reuse a report template, you can change the name and type for any input parameters to better reflect the purpose of the new report.

To edit user-defined input parameters for a report template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the edit icon (✎) for the template you want to edit.
The Report Sections page appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears. The **Input Parameters** section lists all available user-defined parameters for the report template.
5. Click the edit icon (✎).
The Edit Input Parameter pop-up window appears.

6. Change the parameter name in the **Name** field and the parameter type using the **Type** drop-down menu, then click **OK**.

The changed parameter appears in the **Input Parameters** section.

7. Repeat the steps above until you have defined all the parameters you need. Click **OK**.

Your changes are saved and the Report Sections page reappears.

To delete user-defined input parameters for a report template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.

2. Select the **Report Templates** tab.

The Report Templates page appears.

3. Click the edit icon () for the template you want to edit.

The Report Sections page appears.

4. Click **Advanced**.

The Advanced Settings pop-up window appears. The **Input Parameters** section lists all available user-defined parameters for the report template.

5. Click the delete icon () next to the input parameter and confirm.

6. Click **OK**.

The input parameter is deleted and the Report Sections page reappears.

You use input parameters to expand the usefulness of your searches. The input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically constrain a report at generation time to show a particular subset of data without changing the search. For example, you can provide an input parameter for the **Destination IP** field of a report section that drills down on security events at a department level. When you generate the report, you can type the IP network segment for a particular department to get data for that department only.

To constrain the search in a report template with user-defined input parameters:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.

2. Select the **Report Templates** tab.

The Report Templates page appears.

3. Click the edit icon () for the template you want to edit.

The Report Sections page appears.

4. Click the edit icon (✎) next to the **Search** field within the section.
The Search Editor pop-up window appears. Fields that can take an input parameter are marked with the input parameter icon (⊕).
5. Click the input parameter icon (⊕) next to the field, then select the input parameter from the drop-down menu. User-defined input parameters are marked with the icon (🔗).
The input parameter appears in the field.

IMPORTANT! Input parameters you define are available only for search fields that match their parameter type. For example, a parameter of type **Network/IP** is available only for fields that accept IP addresses or network segments in CIDR format.

6. Click **OK** when you have added all necessary input parameters.
The Report Sections page appears with your changes.

Editing Document Attributes in a Report Template

LICENSE: Any

Before you generate your report, you can set document attributes that affect the report's appearance. These attributes include the optional cover page and table of contents. Support for some attributes depends on the selected report format: PDF, HTML, or CSV. The [Document Attribute Support](#) table provides further details on attribute support by format.

Document Attribute Support

ATTRIBUTE	PDF SUPPORT?	HTML SUPPORT?	CSV SUPPORT?
Cover page	yes, with optional logo and custom appearance	yes, with optional logo and custom appearance	no
Table of contents	yes	yes	no
Page headers and footers	yes, with optional text or logo in any field	no	no

Document Attribute Support (Continued)

ATTRIBUTE	PDF SUPPORT?	HTML SUPPORT?	CSV SUPPORT?
Custom starting page number	yes	no	no
Option to suppress numbering of first page	yes	no	no

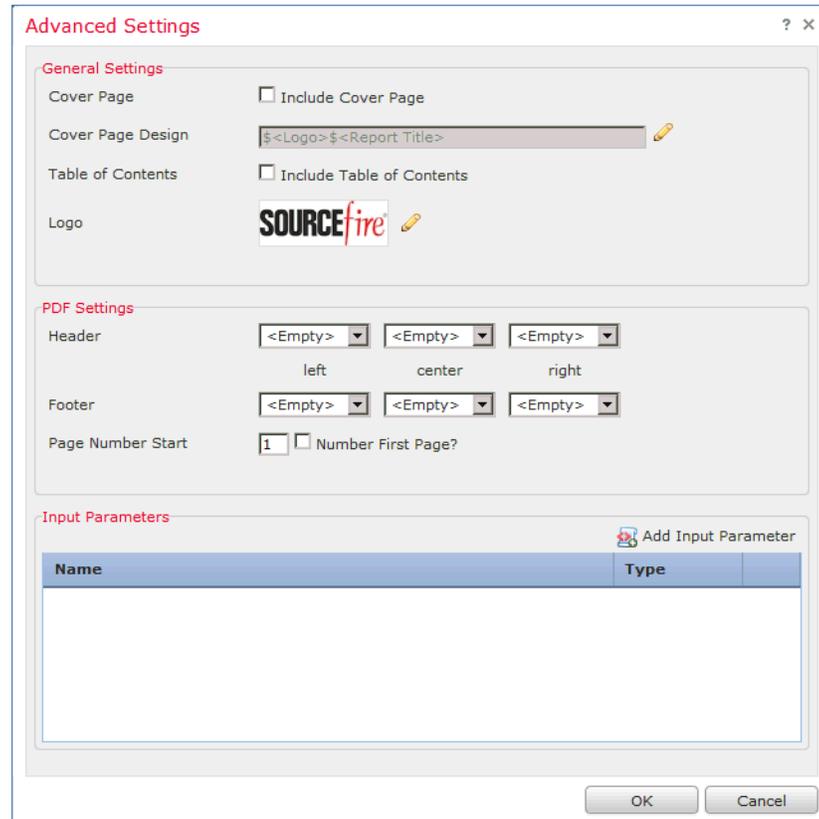
To set the document attributes for PDF and HTML reports:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the edit icon () for the report template you want to edit.
The Report Sections page appears.

4. Click **Advanced**.

The Advanced Settings pop-up window appears.



5. Select **Include Cover Page** to add a cover page.
6. Click the edit icon (✎) next to the **Cover Page Design** field to edit the cover page design. For more information, see [Customizing a Cover Page](#) on page 1759.
7. Select **Include Table of Contents** to add a table of contents.
8. Configure the header and footer using the drop-downs of the three **Header** and **Footer** fields. You select header and footer content from the drop-down menus: logo, date, page number, and so on.
If you select **Logo**, the default logo image appears in the selected field. To change the default logo image, see [Managing Logos](#) on page 1760.
9. In the **Page Number Start** field, select the page number of the report's first page.

Select **Number First Page?** to show the page number on the first page following the cover page. If selected, the cover page is not numbered.

10. Click **OK**.

The document attributes are saved and the Report Sections page reappears.

Customizing a Cover Page

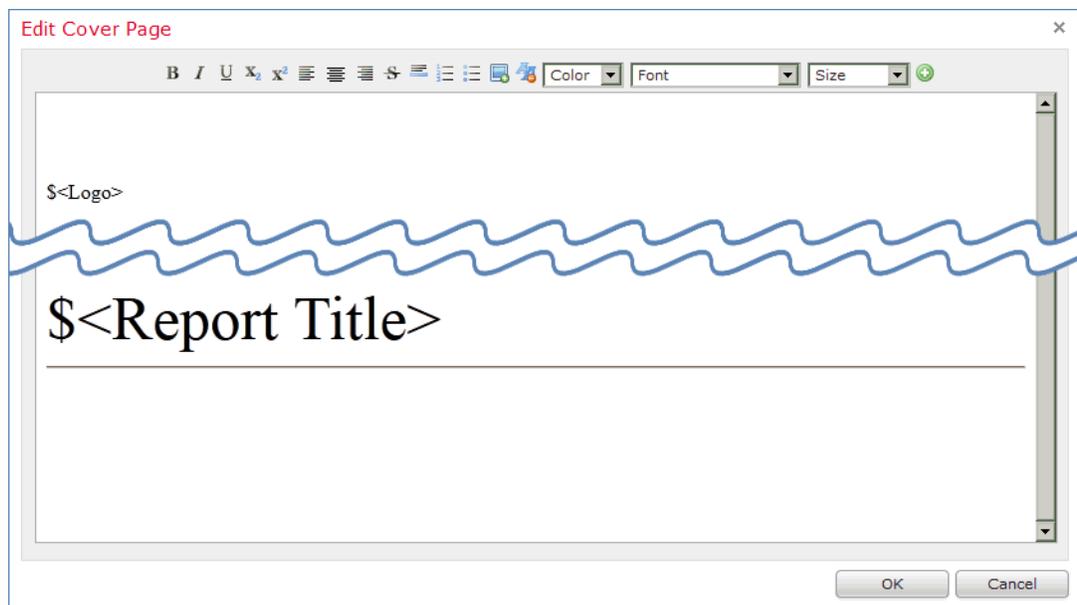
LICENSE: Any

You can customize a report template's cover page. Cover pages can have rich text with multiple font sizes and styles (bold, italic, and so on) as well as input parameters and imported images. For information on input parameters, see [Using Input Parameters](#) on page 1750.

To customize a report template cover page:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears and displays the list of templates.
3. Click the edit icon (✎) for a report template.
The Report Sections page appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears.
5. Click the edit icon (✎) next to **Cover Page Design**.
The Edit Cover Page window appears, displaying the default cover page design.



6. Edit the cover page design within the rich text editor.
7. Click **OK**.
The cover page design is saved and the Advanced Settings window reappears.

Managing Logos

LICENSE: Any

You can store multiple logos on the Defense Center and associate them with different report templates. You set the logo association when you design the template. If you export the template, the export package contains the logo.

For information on where you can insert a logo in reports, see [Editing Document Attributes in a Report Template](#) on page 1756.

See the following related procedures for more information:

- [Adding a New Logo](#) on page 1760
- [Changing the Logo for a Report Template](#) on page 1761
- [Deleting a Logo](#) on page 1762

Adding a New Logo

LICENSE: Any

Logos uploaded to your Defense Center are available for all report templates on that Defense Center. Logo images must be in JPG format.

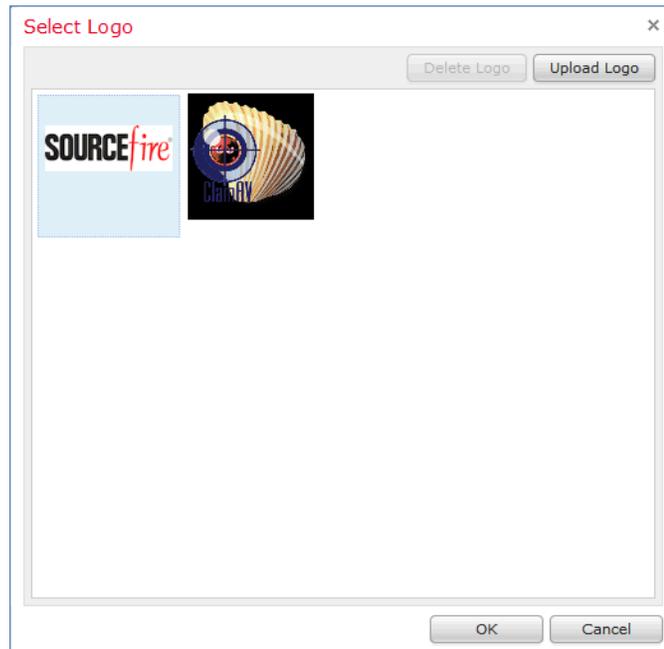
To add a logo to a Defense Center:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the edit icon (✎) for the report template you want to edit.
The Report Sections page appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears. The logo currently associated with the template appears under **Logo** in **General Settings**.

5. Click the edit icon (✎) for the logo.

The Select Logo pop-up window appears with images of currently uploaded logos.



6. Click **Upload Logo**.

The Upload Logo pop-up window appears.

7. Select the logo file to upload by doing one of the following:
 - type the location of the logo file
 - click the **Browse** button and browse to the file's location
8. Click **Upload**.

The image is uploaded to the Defense Center and appears in the Select Logo pop-up window.

9. Optionally, associate the new logo with the current template by selecting it and clicking **OK**.

The Advanced Settings window reappears with the associated logo image.

Changing the Logo for a Report Template

LICENSE: Any

You can change the logo in a report to any JPG image uploaded to your Defense Center. For example, if you reuse a template, you can associate a logo for a different organization with the report.

To change the logo for a report template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the edit icon (✎) for the report template you want to edit.
The Report Sections page appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears. The logo currently associated with the template appears under **Logo** in **General Settings**.
5. Click the edit icon (✎) for the logo.
The Select Logo pop-up window appears with images of currently uploaded logos.
6. Select the logo to associate with the report template.
The selected logo is highlighted.
7. Click **OK**.
The Advanced Settings window reappears with the associated logo image.

Deleting a Logo

LICENSE: Any

You can delete logos from your Defense Center. Deleting a logo removes it from all templates where it is used. The deletion cannot be undone.

Note that you cannot delete the predefined Sourcefire logo.

To delete a logo from a Defense Center:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the edit icon (✎) for the report template you want to edit.
The Report Sections page appears.
4. Click **Advanced**.
The Advanced Settings pop-up window appears. The logo currently associated with the template appears under **Logo** in **General Settings**.

5. Click the edit icon () for the logo.
The Select Logo pop-up window appears with images of currently uploaded logos.
6. Select the logo you want to delete.
The selected logo is highlighted.
7. Click **Delete Logo**.
The deleted logo disappears from the Select Logo pop-up window.
8. Click **OK**.
Your changes are saved and the Advanced Settings window reappears.

Using Report Generation Options

LICENSE: Any

You have several additional options when generating reports. You can automatically schedule report generation, send reports via email, and store generated reports remotely. For more information, see the following sections:

- [Generating Reports Using the Scheduler](#) on page 1763
- [Distributing Reports by Email at Generation Time](#) on page 1763
- [Using Remote Storage for Reports](#) on page 1765

Generating Reports Using the Scheduler

LICENSE: Any

You can use the Sourcefire 3D System scheduler to automate report generation. You can customize the schedule on a full range of time frames such as daily, weekly, monthly, and so on. For more information, see [Automating Reports](#) on page 1940.

Distributing Reports by Email at Generation Time

LICENSE: Any

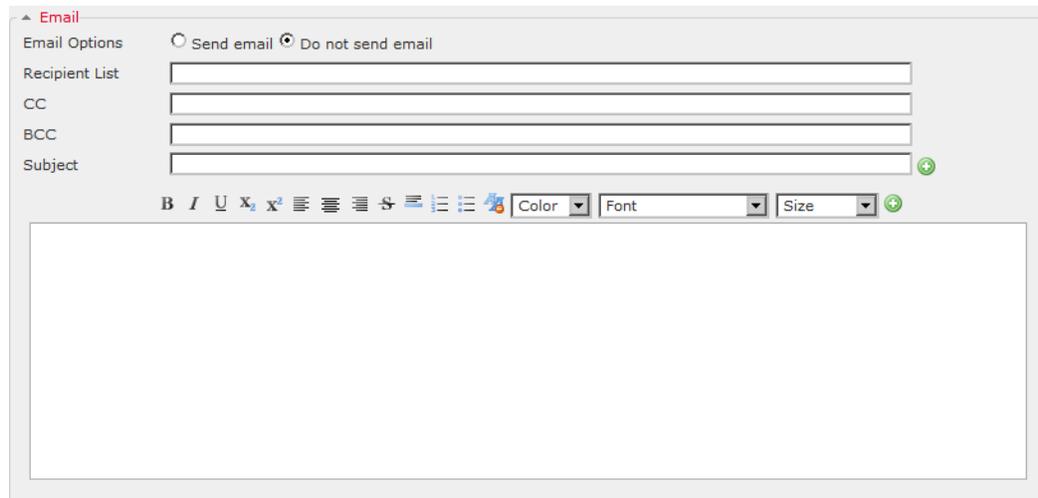
When you generate a report from its template, you can choose to automatically send the report as an email attachment to a list of recipients.

IMPORTANT! You must have a properly configured mail relay host to deliver a report by email. If you have not previously set up a mail host, see [Configuring a Mail Relay Host and Notification Address](#) on page 1983.

To email a report at generation time:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Click the generate report icon (📄) for the template you want to generate from.
The Generate Report pop-up window appears.
4. Expand the **Email** section of the window.



5. In the **Email Options** field, select **Send Email**.
6. In the **Recipient List**, **CC**, and **BCC** fields, type recipients' email addresses in comma-separated lists.
7. In the **Subject** field, type a subject for your email.

TIP! You can provide input parameters in the **Subject** field and the message body to dynamically generate information in the email, such as a timestamp or the name of the Defense Center. For further information, see [Using Input Parameters](#) on page 1750.

8. Type a cover letter in the email body as necessary. The available rich text features include a wide range of fonts, numbered and bullet lists, and so on.

9. When all fields in the Generate Report window are correct, click **OK** and confirm.

The system distributes the generated report by email. You can configure the email's From address under **Email Notification** in the system policy. For more information, see [Managing System Policies](#) on page 1961.

Using Remote Storage for Reports

LICENSE: Any

You can configure the reporting system to place newly generated report files in your configured remote storage location. You can also move any locally stored report to your remote storage location.

IMPORTANT! You cannot move reports in remote storage back to local storage.

To use remote storage, you must first configure a remote storage location. When configured, the remote storage location appears at the bottom of the report list. The location includes current disk usage for NFS and SMB mounted storage, but not for SSH. For configuration information, see [Managing Remote Storage](#) on page 2017.

To store reports remotely as they are generated:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Reports** tab.
The Reports page appears.
3. Select the **Enable Remote Storage of Reports** check box at the bottom of the page.

The Defense Center stores newly generated reports in the remote location indicated at the bottom of the page. The **Location** column data for these reports is **Remote**.

You can move your reports in local storage to a remote storage location in batch mode or singly.

To move generated reports from local to remote storage:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Reports** tab.
The Reports page appears.

3. Select the check boxes next to the reports you want to move, then click **Move**.

TIP! Select the check box at the top left of the page to move all reports on the page. If you have multiple pages of reports, a second check box appears that you can select to move all reports on all pages.

4. Confirm that you want to move the reports.
The reports are moved.

Managing Report Templates and Report Files

LICENSE: Any

In addition to creating and editing templates, you can perform the following template management tasks:

- [Exporting and Importing Report Templates](#) on page 1766
- [Deleting Report Templates](#) on page 1768

You can also perform the following management tasks for your generated report files:

- [Downloading Reports](#) on page 1768
- [Deleting Reports](#) on page 1769

Exporting and Importing Report Templates

LICENSE: Any

The file that you generate when you export a report template contains all necessary data to create the same report on another Defense Center. The export file, which is in a proprietary SFO format, includes:

- the report template, with all section design elements and document attributes
- all saved searches used in the report
- all images used in the report
- all custom tables used in the report

The only configuration that may be required after you import the template on another Defense Center is automatic report generation scheduling.

IMPORTANT! Importing and exporting report templates requires both Defense Centers to be at the same software version level.

To export a report template:

ACCESS: Admin

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. For the template you want to export, click the export icon (📄).
The system produces a Sourcefire object package file with **.sfo** extension and opens an Opening Object pop-up window that displays the package's file name.
4. Select **Save file** and **OK** to save the file to your local computer.
5. You can change the name of the **.sfo** package to a more descriptive one for your convenience. When you import the package, regardless of its name, the importing Defense Center will give the template the same name it had on the source Defense Center.

The SFO files exported from a Defense Center contain all elements necessary to add the report template to another Defense Center. The import process therefore requires only uploading the package to the second Defense Center and running the import process.

To import a report template:

ACCESS: Admin

1. Select **System > Tools > Import/Export**.
The Import/Export page appears, including a list of the report templates on the Defense Center.
2. Click **Upload Package**.
The Package Name page appears.
3. You have two options:
 - Type the path to the package you want to upload.
 - Click **Browse** to locate the package.
4. Click **Upload**.
The **Report Template** section of the object list appears, showing the template to be imported.



5. Select the check box next to the template and click **Import**.
The template appears in the list of objects on the destination Defense Center.

Deleting Report Templates

LICENSE: Any

Report templates remain listed on the Report Templates tab for reuse until you delete them. Note that you cannot delete Sourcefire-provided report templates.

IMPORTANT! Security Analysts can delete only report templates they created.

To delete a report template:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Report Templates** tab.
The Report Templates page appears.
3. Next to the template you want to delete, click the delete icon () and confirm.
The template name disappears from the list.

Downloading Reports

LICENSE: Any

You can download any report file to your local computer. From there, you can email it or distribute it electronically by other available means. For information on distributing reports automatically by email at generation time, see [Distributing Reports by Email at Generation Time](#) on page 1763.

To download reports:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Reports** tab.
The Reports page appears.
3. Select the check boxes next to the reports you want to download, then click **Download**.

TIP! Select the check box at the top left of the page to download all reports on the page. If you have multiple pages of reports, a second check box appears that you can select to download all reports on all pages.

4. Follow your browser's prompts to download the reports.
If you select multiple reports, they are downloaded in a single **.zip** file.

Deleting Reports

LICENSE: Any

You can delete your report files at any time. The procedure completely removes the files, and no recovery is possible. Although you still have the report template that generated the report, it may be difficult to regenerate a particular report file if the time window was expanding or sliding. For information on the time window, see [Editing the Sections of a Report Template](#) on page 1743. Regeneration may also be difficult if your template uses input parameters. For information on using input parameters, see [Using Input Parameters](#) on page 1750.

To delete reports:

ACCESS: Admin/Any Security Analyst

1. Select **Overview > Reporting**.
2. Select the **Reports** tab.
The Reports page appears.
3. Select the check boxes next to the reports you want to delete, then click **Delete**.

TIP! Select the check box at the top left of the page to delete all reports on the page. If you have multiple pages of reports, a second check box appears that you can select to delete all reports on all pages.

4. Confirm the deletion.
The reports are deleted.

CHAPTER 43

SEARCHING FOR EVENTS

Sourcefire appliances generate information that is stored as events in database tables. Events contain multiple fields that describe the activity that caused the appliance to generate the event.

The Sourcefire 3D System provides predefined searches that serve as examples and can provide quick access to important information about your network. You can modify fields within the predefined searches for your network environment, then save the searches to reuse later. You can also use your own search criteria.

The search criteria you can use depends on the type of search, but the mechanics are the same. See the following sections for more information on how to perform a search and on the correct syntax to use in search fields:

- [Performing and Saving Searches](#) on page 1770
- [Using Wildcards and Symbols in Searches](#) on page 1775
- [Using Objects and Application Filters in Searches](#) on page 1775
- [Specifying Time Constraints in Searches](#) on page 1775
- [Specifying IP Addresses in Searches](#) on page 1776
- [Specifying Ports in Searches](#) on page 1777

Performing and Saving Searches

LICENSE: Any

You can create and save searches for any of the different event types. When you create a search you give it a name and specify whether the search will be available to you alone or to all users of the appliance. If you want to use the

search as a data restriction for a custom user role, you **must** save it as a private search.

For more information, see the following sections:

- [Performing a Search](#) on page 1771
- [Loading a Saved Search](#) on page 1774
- [Deleting a Saved Search](#) on page 1774

IMPORTANT! To search a custom table, follow a slightly different procedure; see [Searching Custom Tables](#) on page 1788.

Performing a Search

LICENSE: Any

For some event types, the Sourcefire 3D System provides predefined searches that serve as examples and can provide quick access to important information about your network. You can modify fields within the predefined searches for your network environment, then save the searches to reuse later. You can also use your own search criteria.

To perform a search:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.

2. From the **Table** drop-down list, select the type of event or data you want to search for

The page reloads with the appropriate search constraints. The following graphic shows the search page for the audit log.

The screenshot shows a web form titled "Search Information". At the top, a note states: "Note: If a search name is not specified, an automatically generated name will be used." Below this, the "Table" dropdown menu is set to "Audit Log Events". The "Name" field is empty, and the text "Search 1, My Search" is displayed to its right. The "Save As Private" checkbox is checked. Under the "Constraint" section, there are several input fields: "User" (with "username" to its right), "Subsystem" (with "subsystem" to its right), "Message" (with "message" to its right), "Time" (with "> 2009-07-16 13:00:31, < today at 4:30pm" to its right), "Source IP" (with "192.168.1.3, 2001:db8:85a3::1370" to its right), and "Configuration Change" (with "yes, no" to its right). At the bottom of the form are two buttons: "Search" and "Save As New Search".

3. Optionally, if you want to save the search, enter a name for it in the **Name** field.
If you do not enter a name, a name is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields.
 - All fields accept negation (!).
 - All fields accept comma-separated lists. If you enter multiple criteria, the search returns only the records that match all the criteria.
 - Many fields accept one or more asterisks (*) as wild cards.
 - Specify n/a in any field to identify events where information is not available for that field; use !n/a to identify the events where that field is populated.
 - Click the add object icon (🟢+) that appears next to a search field to use an object as a search criterion.
5. See the following sections for detailed information on the search criteria you can use:
 - [Searching Audit Records](#) on page 2197
 - [Searching for Applications](#) on page 1427
 - [Searching for Application Details](#) on page 1432
 - [Searching for Connection Data](#) on page 575
 - [Searching for Correlation Events](#) on page 1526
 - [Searching for Discovery Events](#) on page 1399

- [Searching for File Events](#) on page 1227
 - [Searching for Health Events](#) on page 2184
 - [Searching for Host Attributes](#) on page 1416
 - [Searching for Hosts](#) on page 1408
 - [Searching for Intrusion Events](#) on page 644
 - [Searching for Malware Events](#) on page 1239
 - [Searching for Scan Results](#) on page 1721
 - [Searching for Servers](#) on page 1422
 - [Searching for Sourcefire Vulnerabilities](#) on page 1439
 - [Searching the Rule Update Import Log](#) on page 2089
 - [Searching for Remediation Status Events](#) on page 1638
 - [Searching for Third-Party Vulnerabilities](#) on page 1444
 - [Searching for Users](#) on page 1451
 - [Searching for User Activity](#) on page 1456
 - [Searching for Compliance White List Events](#) on page 1576
 - [Searching for White List Violations](#) on page 1582
6. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

7. You have the following options:
- Click **Search** to start the search.
Your search results appear in the default workflow for the table you are searching, constrained by time (if applicable). To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. Note that you **cannot** use a different workflow for scan results.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save As New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**) so that you can run it at a later time.

Loading a Saved Search

LICENSE: Any

If you previously saved a search, you can load it, make any necessary modifications, and then start the search.

To load a saved search:

ACCESS: Admin/Any Security Analyst

1. You have two options:
 - From any page on a workflow, click **Search**.
 - Select **Analysis > Search**, then select the type of events you want to search for.

The Search page appears.

2. From the list of saved searches on the left of the page, select the search you want to load and click **Load**.

Settings from the saved search populates the search constraints fields.

3. Optionally, change the search constraints.
4. Click **Search**.

The events that match your search constraints appear.

Deleting a Saved Search

LICENSE: Any

If you have saved searches, you can delete them from the Search page.

To delete a saved search:

ACCESS: Admin/Any Security Analyst

1. You have two options:
 - From any page on a workflow, click **Search**.
 - Select **Analysis > Search**, then select the event type for the search that you want to delete.

The Search page appears.

2. From the list of saved searches, select the search you want to delete and click **Delete**.

The search is deleted.

Using Wildcards and Symbols in Searches

LICENSE: Any

Many text fields on search pages allow you to use an asterisk (*) to match characters in a string. For example, specifying `net*` matches `network`, `netware`, `netscape`, and so on.

If you want to search for non-alphanumeric characters (including the asterisk character), enclose the search string in quotation marks. For example, to search for the string:

Find an asterisk (*)

enter:

“Find an asterisk (*)”

Note that in text fields that allow a wildcard, you **must** use the wildcard if you want to match a partial string. For example, if you are searching the audit log for all audit records that involve page views (that is, the message is Page View), searching for `Page` returns no results. Instead, specify `Page*`.

Using Objects and Application Filters in Searches

LICENSE: Any

The Sourcefire 3D System allows you to create named objects, object groups, and application filters that can be used as part of your network configuration. You can use these objects, groups, and filters as search criteria when performing or saving searches.

When you perform a search, objects, object groups, and application filters appear in the format, `${object_name}`. For example, a network object with the object name `ten_ten_network` appears as `${ten_ten_network}` in a search.

You can click the add object icon () that appears next to a search field where you can use an object as a search criterion.

Specifying Time Constraints in Searches

LICENSE: Any

You can use a number of formats for specifying time search constraints. You can enter a time you want to match, and, optionally, a less than (<) or greater than (>) operator to match times before or after the time you enter.

The formats accepted by search criteria fields that take a time value are shown in the following table.

Time Specification in Search Fields

TIME FORMATS	EXAMPLE
today [at HH:MMam pm]	today today at 12:45pm
YYYY-MM-DD HH:MM:SS	2006-03-22 14:22:59

You can precede a time value with one of the following operators/keyword.

Time Specification Operators

OPERATOR	EXAMPLE	EXPLANATION
<	< 2006-03-22 14:22:59	Returns events with a timestamp before 2:23 PM, March 22, 2006.
>	> today at 2:45pm	Returns events with a timestamp later than today at 2:45 PM.

Specifying IP Addresses in Searches

LICENSE: Any

When specifying IP addresses in searches, you can enter an individual IP address, a comma-separated list of addresses, an address block, or a range of IP addresses separated with a hyphen (-). You can also use negation.

For searches that support IPv6 (such as intrusion event, connection data, and correlation event searches) you can enter IPv4 and IPv6 addresses and CIDR/prefix length address blocks in any combination.

When you use CIDR or prefix length notation to specify a block of IP addresses, the Sourcefire 3D System 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 Sourcefire 3D System uses 10.0.0.0/8.

The following table contains examples of valid ways to enter IP addresses. Because IP addresses can be represented by network objects, you can also click the add network object icon (⊕) that appears next to an IP address search field

to use a network object as an IP address search criterion. For more information, see [Using Objects and Application Filters in Searches](#) on page 1775.

Acceptable IP Address Syntax

TO SPECIFY...	TYPE...	FOR EXAMPLE...
a single IP address	the IP address.	192.168.1.1 2001:db8::abcd
multiple IP addresses using a list	a comma-separated list of IP addresses. Do not add a space before or after the commas.	192.168.1.1,192.168.1.2 2001:db8::b3ff, 2001:db8::0202
a range of IP addresses that can be specified with a CIDR block or prefix length	the IP address block in IPv4 CIDR or IPv6 prefix length notation.	192.168.1.0/24 This specifies any IP 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 on page 57.
a range of IP addresses that cannot be specified with a CIDR block or prefix	the IP address range using a hyphen. Do not add a space before or after the hyphen.	192.168.1.1-192.168.1.5 2001:db8::0202-2001:db8::8329
negation of any of the other ways to specify IP addresses or ranges of IP addresses	an exclamation point in front of the IP address, block, or range.	192.168.0.0/32, !192.168.1.10 !2001:db8::/32 !192.168.1.10, !2001:db8::/32
on a Defense Center, all IP addresses in the networks you are monitoring	local	local
on a Defense Center, all IP addresses that are not in the networks you are monitoring	remote	remote

Specifying Ports in Searches

LICENSE: Any

The Sourcefire 3D System accepts specific syntax for port numbers in searches. You can enter:

- a single port number
- a comma-separated list of port numbers

- two port numbers separated by a dash to represent a range of port numbers
- a port number followed by a protocol abbreviation, separated by a forward slash (only when searching for intrusion events)
- a port number or range of port numbers preceded by an exclamation mark to indicate a negation of the specified ports

IMPORTANT! Do **not** use spaces when specifying port numbers or ranges.

The [Port Syntax Examples](#) table contains examples of valid ways to enter ports as search constraints.

Port Syntax Examples

EXAMPLE	DESCRIPTION
21	Returns all events on port 21, including TCP and UDP events.
!23	Returns all events except those on port 23.
25/tcp	Returns all TCP-related intrusion events on port 25.
21/tcp,25/tcp	Returns all TCP-related intrusion events on ports 21 and 25
21-25	Returns all events on ports 21 through 25.

CHAPTER 44

USING CUSTOM TABLES

As the Sourcefire 3D System collects information about your network, the Defense Center stores it in a series of database tables. When you use a workflow to view the resulting information, the Defense Center pulls the data from one of these tables. For example, the columns on each page of the Network Applications by Count workflow are taken from the fields in the Applications table.

If you determine that your analysis of the activity on your network would be enhanced by combining fields from different tables, you can create a custom table. For example, you could combine the host criticality information from the predefined Host Attributes table with the fields from the predefined Connection Data table and then examine connection data in a new context.

Note that you can create custom workflows for either predefined or custom tables. For more information on creating custom workflows, see [Creating Custom Workflows](#) on page 1841.

The following sections describe how to create and use your own custom tables:

- [Understanding Custom Tables](#) on page 1780
- [Creating a Custom Table](#) on page 1783
- [Modifying a Custom Table](#) on page 1786
- [Deleting a Custom Table](#) on page 1786
- [Viewing a Workflow Based on a Custom Table](#) on page 1787
- [Searching Custom Tables](#) on page 1788

Understanding Custom Tables

LICENSE: FireSIGHT

Custom tables contain fields from two or more predefined tables. The Sourcefire 3D System is delivered with a number of Sourcefire-defined custom tables, but you can create additional custom tables that contain only information that matches your specific needs.

For example, the Sourcefire 3D System is delivered with Sourcefire-defined custom tables that correlate intrusion event data with host data, so you can search for events that impact critical systems and view the results of that search in one workflow. The [Sourcefire-defined Custom Tables](#) table describes the custom tables provided with the system.

Sourcefire-defined Custom Tables

TABLE	DESCRIPTION
Hosts with Servers	Includes fields from the Hosts and Servers tables, providing you with information about the detected applications running on your network, as well as basic operating system information about the hosts running those applications.
Intrusion Events with Destination Criticality	Includes fields from the Intrusion Events table and the Hosts table, providing you with information on the intrusion events, as well as the host criticality of the destination host involved in each intrusion event. TIP! Use this table to search for intrusion events involving destination hosts with high host criticality.
Intrusion Events with Source Criticality	Includes fields from the Intrusion Events table and the Hosts table, providing you with information on the intrusion events and the host criticality of the source host involved in each intrusion event. TIP! Use this table to search for intrusion events involving source hosts with high host criticality.

Understanding Possible Table Combinations

LICENSE: FireSIGHT + Protection

When you create a custom table, you can combine fields from predefined tables that have related data. The [Custom Table Combinations](#) table lists the predefined tables you can combine to create a new custom table. Keep in mind that you can

create a custom table that combines fields from more than two predefined custom tables.

Custom Table Combinations

YOU CAN COMBINE FIELDS FROM...	WITH FIELDS FROM...
Applications	<ul style="list-style-type: none"> • Correlation Events • Intrusion Events • Connection Summary Data • Host Attributes • Application Details • Discovery Events • Connection Events • Hosts • Servers • White List Events
Correlation Events	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts
Intrusion Events	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts • Servers
Connection Summary Data	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts • Servers
Host Attributes	<ul style="list-style-type: none"> • Applications • Correlation Events • Intrusion Events • Connection Summary Data • Application Details • Discovery Events • Connection Events • Hosts • Servers • White List Events
Application Details	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts

Custom Table Combinations (Continued)

YOU CAN COMBINE FIELDS FROM...	WITH FIELDS FROM...
Discovery Events	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts
Connection Events	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts • Servers
Hosts	<ul style="list-style-type: none"> • Applications • Correlation Events • Intrusion Events • Connection Summary Data • Host Attributes • Application Details • Discovery Events • Connection Events • Servers • White List Events
Servers	<ul style="list-style-type: none"> • Applications • Intrusion Events • Connection Summary Data • Host Attributes • Connection Events • Hosts
White List Events	<ul style="list-style-type: none"> • Applications • Host Attributes • Hosts

Sometimes a field in one table maps to more than one field in another table. For example, the predefined **Intrusion Events with Destination Criticality** custom table combines fields from the Intrusion Events table and the Hosts table. Each event in the Intrusion Events table has two IP addresses associated with it—a source IP address and a destination IP address. However, the “events” in the Hosts table each represent a single host IP address (hosts may have multiple IP addresses). Therefore, when you create a custom table based on the Intrusion Events table and the Hosts table, you must choose whether the data you display from the Hosts table applies to the host source IP address or the host destination IP address in the Intrusion Events table.

When you create a new custom table, a default workflow that displays all the columns in the table is automatically created. Also, just as with predefined tables,

you can search custom tables for data that you want to use in your network analysis. You can also generate reports based on custom tables, as you can with predefined tables.

For more information on creating custom tables, see:

- [Creating a Custom Table](#) on page 1783
- [Modifying a Custom Table](#) on page 1786
- [Deleting a Custom Table](#) on page 1786
- [Viewing a Workflow Based on a Custom Table](#) on page 1787
- [Searching Custom Tables](#) on page 1788

Creating a Custom Table

LICENSE: FireSIGHT

If you determine that your analysis of the activity on your network would be enhanced by combining fields from different tables, you can create a custom table.

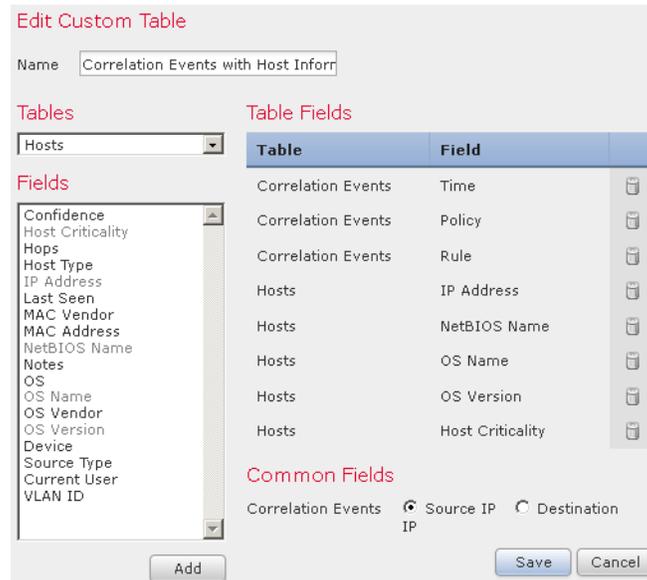
TIP! Instead of creating a new custom table, you can export a custom table from another Defense Center, then import it onto your Defense Center. You can then edit the imported custom table to suit your needs. For more information, see [Importing and Exporting Objects](#) on page 2225.

To create a custom table, decide which predefined tables delivered with the Sourcefire 3D System contain the fields you want to include in your custom table. You can then choose which fields you want to include and, if necessary, configure field mappings for any common fields.

TIP! Data involving the Hosts table allows you to view data associated with all IP addresses from one host, rather than one specific IP address.

For example, consider a custom table that combines fields from the Correlation Events table and the Hosts table. You can use this custom table to get detailed information about the hosts involved in violations of any of your correlation policies. Note that you must decide whether to display data from the Hosts table

that matches the source IP address or the destination IP address in the Correlation Events table.



If you view the table view of events for this custom table, it displays correlation events, one per row. The following information is included:

- the date and time the event was generated
- the name of the correlation policy that was violated
- the name of the rule that triggered the violation
- the IP address associated with the source, or initiating, host involved in the correlation event
- the source host's NetBIOS name
- the operating system and version the source host is running
- the source host criticality

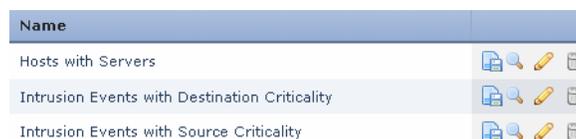
TIP! You could create a similar custom table that displays the same information for destination, or responding, hosts.

To build the custom table in the previous example:

ACCESS: Admin

1. Select **Analysis > Custom > Custom Tables**.

The Custom Tables page appears.



Name	
Hosts with Servers	 
Intrusion Events with Destination Criticality	 
Intrusion Events with Source Criticality	 

2. Click **Create Custom Table**.

The Create Custom Table page appears.



Create Custom Table

Name

Tables

Table Fields

Table	Field
-------	-------

Fields

- Type
- Application
- Category
- IP Address
- Business Relevance
- Risk
- Tag
- Current User

3. In the **Name** field, type a name for the custom table, such as **Correlation Events with Host Information (Src IP)**.
4. From the **Tables** drop-down list, select **Correlation Events**.
The fields in the Correlation Events table appear in the **Fields** list.
5. Under **Fields**, select **Time** and click **Add** to add the date and time when a correlation event was generated.
6. Repeat step 5 to add the **Policy** and **Rule** fields.

TIP! You can use Ctrl or Shift while clicking to select multiple fields. You can also click and drag to select multiple adjacent values. However, if you want to specify the order the fields appear in the table view of events associated with the table, add the fields one at a time.

7. From the **Tables** drop-down list, select **Hosts**.
The fields in the Hosts table appear in the **Fields** list. For more information on these fields, see [Understanding the Hosts Table](#) on page 1403.
8. Add the **IP Address**, **NetBIOS Name**, **OS Name**, **OS Version**, and **Host Criticality** fields to the custom table.

9. Under **Common Fields**, next to **Correlation Events**, select **Source IP**.
Your custom table is configured to display the host information you chose in step 8 for the source, or initiating, hosts involved in correlation events.

TIP! You could create a custom table that displays detailed host information for the destination, or responding, hosts involved in a correlation event by following this procedure but selecting **Destination IP** instead of **Source IP**.

10. Click **Save**.
The custom table is saved.

Modifying a Custom Table

LICENSE: FireSIGHT

You can add or delete fields in a custom table as your needs change.

To modify a custom table:

ACCESS: Any/Admin

1. Select **Analysis > Custom > Custom Tables**.
The Custom Tables page appears.
2. Click the edit icon (✎) next to the table you want to edit.
The Edit Custom Table page appears. See [Creating a Custom Table](#) on page 1783 for information on the various configurations you can change.
3. Optionally, remove fields from the table by clicking the delete icon (🗑) next to the fields you want to remove.
4. Make other changes as needed and click **Save**.
Your custom table is updated.

Deleting a Custom Table

LICENSE: FireSIGHT

You can delete a custom table that you no longer need.

To delete a custom table:

ACCESS: Any/Admin

1. Select **Analysis > Custom > Custom Tables**.
The Custom Tables page appears.

2. Click the delete icon () next to the custom table you want to delete.
The table is deleted.

Viewing a Workflow Based on a Custom Table

LICENSE: FireSIGHT

When you create a custom table, the system automatically creates a default workflow for it. The first page of this workflow displays a table view of events. If you include intrusion events in your custom table, the second page of the workflow is the packet view. Otherwise, the second page of the workflow is a hosts page. You can also create your own custom workflows based on your custom table.

TIP! If you create a custom workflow based on a custom table, you can specify it as the default workflow for that table. For more information, see [Configuring Event View Settings](#) on page 2218.

You can use the same techniques to view events in your custom table that you use for event views based on predefined tables. See [Using Workflow Pages](#) on page 1812 for more information.

To view a workflow based on a custom table:

ACCESS: Any/Admin

1. Select **Analysis > Custom > Custom Tables**.
The Custom Tables page appears.
2. Click the view icon () next to the custom table on which the workflow you want to see is based.

The first page of the default workflow for the custom table appears. To use a different workflow, click (**switch workflow**) by the workflow title. For information on how to specify a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

Searching Custom Tables

LICENSE: FireSIGHT

You can create and save searches for a custom table. You may want to create searches customized for your network environment, then save them to reuse later.

TIP! Click the object icon (🟢) next to a search field to use an object as a search criterion. For more information on searches, including information on special search syntax, using objects in searches, and saving and loading searches, see [Performing and Saving Searches](#) on page 1770.

To perform a search on a custom table:

ACCESS: Any/Admin

1. Select **Analysis > Custom > Custom Tables**.

The Custom Tables page appears.

2. Click the view icon (🔍) next to the custom table you want to search.

The first page of the default workflow for the custom table appears. To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

3. Click **Search**.

The custom table's search page appears. The following graphic shows the search page for the Hosts with Servers predefined custom table.

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: Hosts with Servers

Name: Search 1, My Search

Save As Favorite:

OS Vendor: Microsoft

OS Name: Windows, unknown

IP Address: 192.168.1.0/24, 192.168.1.3, 2001:0db8:85a3::137

Buttons: Search, Save As New Search

TIP! To search the database for a different kind of event or data, select it from the **Table** drop-down list.

4. Optionally, if you want to save the search, enter a name for it in the **Name** field.

If you do not enter a name, one is created automatically when you save the search.

5. Enter your search criteria in the appropriate fields.

If you enter multiple criteria, the search returns only the records that match all criteria.

The search criteria you can use are the same as the criteria for the predefined tables you used to build your custom table. See the sections listed in the following table for detailed information on the search criteria you can use.

FOR SEARCH CRITERIA FOR...	SEE...
Audit Events	Searching Audit Records on page 2197
Application Details	Searching for Application Details on page 1432
Correlation Events	Searching for Correlation Events on page 1526
Connection Data	Searching for Connection Data on page 575
Hosts	Searching for Hosts on page 1408
Host Attributes	Searching for Host Attributes on page 1416
Hosts with Applications	Searching for Hosts on page 1408 and Searching for Servers on page 1422
Intrusion Events	Searching for Intrusion Events on page 644
Intrusion Events with Destination Criticality	Searching for Intrusion Events on page 644 and Searching for Hosts on page 1408
Intrusion Events with Source Criticality	Searching for Intrusion Events on page 644 and Searching for Hosts on page 1408
Status Events	Searching for Remediation Status Events on page 1638
Discovery Events	Searching for Discovery Events on page 1399
User Events	Searching for User Activity on page 1456

FOR SEARCH CRITERIA FOR...	SEE...
Rule Update Import Log	Searching the Rule Update Import Log on page 2089
Applications	Searching for Applications on page 1427
Users	Searching for Users on page 1451
Vulnerabilities	Searching for Sourcefire Vulnerabilities on page 1439
White List Events	Searching for Compliance White List Events on page 1576
White List Violations	Searching for White List Violations on page 1582

6. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

7. You have the following options:
- Click **Search** to start the search.
Your search results appear in the default workflow for the custom table, constrained by the current time range (if applicable). To use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

CHAPTER 45

UNDERSTANDING AND USING WORKFLOWS

A workflow is a tailored series of data pages on the Defense Center web interface that analysts can use to evaluate events generated by the system. The Defense Center provides three types of workflows:

- *Predefined workflows*, which are preset workflows installed on the system that you cannot modify or delete.
- *Saved custom workflows*, which are predefined custom workflows that you can modify or delete.
- *Custom workflows*, which are workflows that you create and customize for your specific needs.

For example, when you analyze intrusion events, you can choose from several predefined workflows specifically created for the task.

Note that the data displayed in a workflow often depends on such factors as how you license and deploy your managed devices, whether you configure features that provide the data and, in the case of Series 2 appliances, whether the appliance supports a feature that provides the data. For example, because neither the DC500 Defense Center nor Series 2 devices support URL filtering by category and reputation, the DC500 Defense Center does not display data for this feature and Series 2 devices do not detect this data.

See the following sections for more information about using predefined and custom workflows:

- [Components of a Workflow](#) on page 1792
- [Using Workflows](#) on page 1808
- [Using Custom Workflows](#) on page 1840

TIP! You can also use custom workflows as the basis for event reports. See [Working with Reports](#) on page 1724 for more information.

Components of a Workflow

LICENSE: Any

Workflows can include several types of pages, as described in the following sections.

Table Views

Table views include a column for each of the fields in the database on which your workflow is based.

For example, the table view of discovery events includes the Time, Event, IP Address, User, MAC Address, MAC Vendor, Port, Description, and Device columns.

By contrast, the table view of servers includes the Last Used, IP Address, Port, Protocol, Application Protocol, Vendor, Version, Web Application, Application Risk, Business Relevance, Hits, Source Type, Device, and Current User columns.

Drill-Down Pages

Drill-down pages contain a subset of columns that are available in the database.

For example, a drill-down page for discovery events might include only the IP Address, MAC Address, and Time columns. A drill-down page for intrusion events, on the other hand, might include the Priority, Impact Flag, Inline Result, and Message columns.

Generally, drill-down pages are intermediate pages that you use to narrow your investigation to a few events before moving to a table view page.

Graphs

Workflows based on connection data can include graph pages, also called *connection graphs*.

For example, a connection graph might display a line graph that shows the number of connections detected by the system over time. Generally, connection graphs are, like drill-down pages, intermediate pages that you use to narrow your

investigation. For more information, see [Working with Connection Graphs](#) on page 555.

Final Pages

The final page of a workflow depends on the type of event on which the workflow is based:

- The host view is the final page for workflows based on applications, application details, discovery events, hosts, servers, or any type of vulnerabilities. Viewing host profiles from this page allows you to easily view data on all IP addresses associated with hosts that have multiple addresses. For more information, see [Using Host Profiles](#) on page 1334.
- The user detail view is the final page for workflows based on users and user activity. For more information, see [Understanding User Details and Host History](#) on page 1449.
- The vulnerability detail view is the final page for workflows based on Sourcefire vulnerabilities. For more information, see [Viewing Vulnerability Details](#) on page 1365.
- The packet view is the final page for workflows based on intrusion events. For more information, see [Using the Packet View](#) on page 622.

Workflows based on other kinds of events (for example, audit log events or malware events) do not have final pages.

See the following sections for more information on workflows:

- [Comparing Predefined and Custom Workflows](#) on page 1794
- [Comparing Workflows for Predefined and Custom Tables](#) on page 1794
- [Predefined Intrusion Event Workflows](#) on page 1794
- [Predefined Malware Workflows](#) on page 1797
- [Predefined File Workflows](#) on page 1798
- [Predefined Host Workflows](#) on page 1800
- [Predefined Applications Workflows](#) on page 1801
- [Predefined Application Details Workflows](#) on page 1802
- [Predefined Servers Workflows](#) on page 1803
- [Predefined Host Attributes Workflows](#) on page 1803
- [Predefined Discovery Events Workflows](#) on page 1804
- [Predefined Vulnerabilities Workflows](#) on page 1804
- [Predefined Third-Party Vulnerabilities Workflows](#) on page 1805
- [Predefined User Workflows](#) on page 1804
- [Predefined Correlation and White List Workflows](#) on page 1805
- [Predefined System Workflows](#) on page 1806
- [Saved Custom Workflows](#) on page 1807

Comparing Predefined and Custom Workflows

LICENSE: Any

The Sourcefire 3D System is delivered with a set of *predefined* workflows (described in the sections that follow) that you can use to analyze the events and other data it collects.

Custom workflows are workflows that you create to meet the unique needs of your organization. When you create a custom workflow, you choose the kind of event (or database table) on which the workflow is based. On the Defense Center, you can base a custom workflow on a custom table. You can also choose the pages a custom workflow contains; custom workflows can contain drill-down, table view, and host or packet view pages.

The Defense Center is delivered with several *saved custom workflows*, which are based on the saved custom tables that are also delivered with the Defense Center. The differences between workflows based on predefined and custom tables is described in the next section, [Comparing Workflows for Predefined and Custom Tables](#).

Comparing Workflows for Predefined and Custom Tables

LICENSE: FireSIGHT

You can use the custom tables feature to create tables that use the data from two or more types of events. This is useful because you can, for example, create tables and workflows that correlate intrusion event data with discovery data to allow simple searches for events that affect critical systems. See [Using Custom Tables](#) on page 1779 for information about creating custom tables.

Each custom table has, by default, a workflow that you can use to view the events associated with the table. The features in the workflow differ depending on which type of table you use. For example, custom table workflows based on the intrusion event table always end with the packet view. However, custom table workflows based on discovery events end with the host view.

Unlike workflows based on the predefined event tables, workflows based on custom tables do not have links to other types of workflows.

Predefined Intrusion Event Workflows

LICENSE: Protection

The [Predefined Intrusion Event Workflows](#) table describes the predefined intrusion event workflows included with the Sourcefire 3D System. For

information on accessing these workflows, see [Viewing Intrusion Events](#) on page 602 and [Reviewing Intrusion Events](#) on page 611.

Predefined Intrusion Event Workflows

WORKFLOW NAME	DESCRIPTION
Destination Port	<p>Because destination ports are usually tied to an application, this workflow can help you detect applications that are experiencing an uncommonly high volume of alerts. The Destination Port column can also help you identify applications that should not be present on your network.</p> <p>This workflow begins with a page showing the destination ports associated with the intrusion events, followed by a page showing the event types that were generated. You can then see a tabular view of event information, called the table view of events, followed by a packet view that shows the decoded contents of the packets associated with each event.</p>
Event-Specific	<p>This workflow provides two useful features. Events that occur frequently may indicate:</p> <ul style="list-style-type: none">• false positives• a worm• a badly misconfigured network <p>Events that occur infrequently are most likely evidence of a targeted attack and warrant special attention.</p> <p>This workflow begins with a page showing the event types that were generated. You can then view a page with two tables, one listing the source IP addresses associated with the events, the other showing the destination IP addresses associated with the events. The last pages in the workflow are the table view of events and the packet view.</p>
Events by Priority and Classification	<p>This workflow lists events and their type in order of event priority, along with a count showing how many times each event has occurred.</p> <p>This workflow begins with a drill-down page that contains the priority level, classification and count of each listed event. The last pages in the workflow are the table view of events and the packet view.</p>
Events to Destinations	<p>This workflow provides a high-level view of which host IP addresses are being attacked and the nature of the attack; where available, you can also see information about the countries involved in attacks.</p> <p>This workflow begins with a page of paired event types and destination IP addresses that you can use to investigate what types of events are directed towards specific IP addresses. The last pages in the workflow are the table view of events and the packet view.</p>

Predefined Intrusion Event Workflows (Continued)

WORKFLOW NAME	DESCRIPTION
IP-Specific	<p>This workflow shows which host IP addresses are generating the most alerts. Hosts with the greatest number of events are either public-facing and receiving worm-type traffic (indicating a good place to look for tuning) or require further investigation to determine the cause of the alerts. Hosts with the lowest counts also warrant investigation as they could be the subject of a targeted attack. Low counts may also indicate that a host may not belong on the network.</p> <p>This workflow begins with a page showing two tables, one each for the source and destination IP addresses that are associated with the events. The next page shows the event types that were generated. The last pages in the workflow are the table view of events and the packet view.</p>
Impact and Priority	<p>This workflow lets you find high-impact recurring events quickly. The reported impact flag level is shown with the number of times the event has occurred. Using this information, you can identify the high-impact events that recur most often, which might be an indicator of a widespread attack on your network.</p> <p>This workflow begins with a page showing the impact level, priority, and count associated with each event. Next, a drill-down page appears with the source and destination IP addresses for each event. Events on the second page are sorted by count. The last pages in the workflow are the table view of events and the packet view.</p>
Impact and Source	<p>This workflow can help you identify the source of an attack in progress. The reported impact flag level is shown with the associated source IP address for the event. If, for example, events with a level 1 impact flag are coming from the same source IP repeatedly, they may indicate an attacker who has identified vulnerable systems and is targeting them.</p> <p>This workflow begins with a page showing the impact level, source IP address, priority, and count associated with each event. Within each event level, events are sorted by count, then priority. Next, a drill-down page appears with the source and destination IP addresses for each event. Events on the second page are sorted by count. The last pages in the workflow are the table view of events and the packet view.</p>

Predefined Intrusion Event Workflows (Continued)

WORKFLOW NAME	DESCRIPTION
Impact to Destination	<p>You can use this workflow to identify events repeatedly occurring on vulnerable computers, so you can address the vulnerabilities on those systems and stop any attacks in progress.</p> <p>This workflow begins with a page showing the impact level, inline result (whether the packet was or would have been dropped), destination IP address, priority, and count associated with each event. Within each event level, events are sorted by count, then priority. Next, a drill-down page appears with the source and destination IP addresses for each event. Events on the second page are sorted by count. The last pages in the workflow are the table view of events and the packet view.</p>
Source Port	<p>This workflow indicates which servers are generating the most alerts. You can use this information to identify areas that require tuning, and to decide which servers require attention.</p> <p>This workflow begins with a page showing the source ports associated with the intrusion events, followed by a page showing the types of events that were generated. The last pages in the workflow are the table view of events and the packet view.</p>
Source and Destination	<p>This workflow identifies host IP addresses sharing high levels of alerts. Pairs at the top of the list could be false positives, and may identify areas that require tuning. You can check pairs at the bottom of the list for targeted attacks, for users accessing resources they should not be accessing, or for hosts that do not belong on the network.</p> <p>This workflow begins with a page showing the source and destination IP addresses for each event, followed by a page showing the types of events that were generated. The last pages in the workflow are the table view of events and the packet view.</p>

Predefined Malware Workflows

LICENSE: Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEVICES: feature dependent

The [Predefined Malware Workflows](#) table describes the predefined malware workflows included on the Defense Center. All predefined malware workflows use the table view of malware events.

Note that because neither the DC500 Series 2 Defense Center nor Series 2 devices support network-based advanced malware protection, the DC500 Defense Center does not display data for this feature and Series 2 devices do not detect this data.

For information on accessing malware events, see [Working with Malware Events](#) on page 1229.

Predefined Malware Workflows

WORKFLOW NAME	DESCRIPTION
Malware Summary	This workflow provides a list of the malware detected in network traffic or by endpoint-based FireAMP Connectors, grouped by individual threat.
Malware Event Summary	This workflow provides a quick breakdown of the different malware event types and subtypes.
Hosts Receiving Malware	This workflow provides a list of host IP addresses that have received malware, grouped by the malware files' associated dispositions.
Hosts Sending Malware	This workflow provides a list of host IP addresses that have sent malware, grouped by the malware files' associated dispositions.
Applications Introducing Malware	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.

Predefined File Workflows

LICENSE: Protection

The [Predefined File Workflows](#) table describes the predefined file event workflows included on the Defense Center. All the predefined file event workflows use the table view of file events. For information on accessing file events, see [Working with File Events](#) on page 1222.

Predefined File Workflows

WORKFLOW NAME	DESCRIPTION
File Summary	This workflow provides a quick breakdown of the different file event categories and types, along with any associated malware dispositions.
Hosts Receiving Files	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.
Hosts Sending Files	This workflow provides a list of host IP addresses that have sent files, grouped by the associated malware dispositions for those files.

Predefined Connection Data Workflows

LICENSE: FireSIGHT

The [Predefined Connection Data Workflows](#) table describes the predefined connection data workflows included on the Defense Center. All the predefined

connection data workflows use the table view of connection data. For information on accessing connection data, see [Viewing Connection Data Graphs and Tables](#) on page 555.

Predefined Connection Data Workflows

WORKFLOW NAME	DESCRIPTION
Connection Events	This workflow provides a summary view of basic connection and detected application information, which you can then use to drill down to the table view of events.
Connections by Application	This workflow contains a graph of the 10 most active applications on the monitored network segment, based on the number of detected connections.
Connections by Initiator	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the number of connections where the host initiated the connection transaction.
Connections by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment, based on the number of detected connections.
Connections by Responder	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the number of connections where the host IP was the responder in the connection transaction.
Connections over Time	This workflow contains a graph of the total number of connections on the monitored network segment over time.
Traffic by Application	This workflow contains a graph of the 10 most active applications on the monitored network segment, based on the number of kilobytes transmitted.
Traffic by Initiator	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the total number of kilobytes transmitted from each address.
Traffic by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment, based on the number of kilobytes transmitted.
Traffic by Responder	This workflow contains a graph of the 10 most active host IP addresses on the monitored network segment, based on the total number of kilobytes received by each address.

Predefined Connection Data Workflows (Continued)

WORKFLOW NAME	DESCRIPTION
Traffic over Time	This workflow contains a graph of the total kilobytes transmitted on the monitored network segment over time.
Unique Initiators by Responder	This workflow contains a graph of the 10 most active responding host IP addresses on the monitored network segment, based on the number of unique initiators that contacted each address.
Unique Responders by Initiator	This workflow contains a graph of the 10 most active initiating host IP addresses on the monitored network segment, based on the number of unique responders that the addresses contacted.

Predefined Host Workflows

LICENSE: FireSIGHT

The [Predefined Host Workflows](#) table describes the predefined workflows that you can use with host data.

Predefined Host Workflows

WORKFLOW NAME	DESCRIPTION
Hosts	This workflow contains a table view of hosts followed by the host view. Workflow views based on the Hosts table allow you to easily view data on all IP addresses associated with a host. See Viewing Hosts on page 1402 for more information.
Operating System Summary	You can use this workflow to analyze the operating systems in use on your network. This workflow provides a series of pages that start with a list of the operating systems and operating system vendors on your network, continuing with the number of hosts running each version of that operating system. The next page lists hosts by criticality, IP address, and NetBIOS name, with their associated operating systems and operating system vendors. The workflow finishes with a table view of hosts, followed by the host view. See Viewing Hosts on page 1402 for more information.

Predefined Applications Workflows

LICENSE: FireSIGHT

The [Predefined Applications Workflows](#) table describes the predefined workflows that you can use with application data.

Predefined Applications Workflows

WORKFLOW NAME	DESCRIPTION
Application Business Relevance	You can use this workflow to analyze running applications of each estimated business relevance level on your network, so you can monitor appropriate use of your network resources. This workflow begins with a count of hosts running applications of each relevance level, followed by a table of individual applications with their business relevance levels and host counts, a table view of applications, and the host view. See Viewing Applications on page 1425 for more information.
Application Category	You can use this workflow to analyze running applications of each category (such as email, search engine, or social networking) on your network, so you can monitor appropriate use of your network resources. This workflow begins with a count of hosts running applications of each category, followed by a count of hosts running individual applications, a table view of applications, and the host view. See Viewing Applications on page 1425 for more information.
Application Risk	You can use this workflow to analyze running applications of each estimated security risk level on your network, so you can estimate the potential risk of users' activity and take appropriate action. This workflow begins with a count of hosts running applications of each risk level, followed by a table of individual applications with their business relevance levels and host counts, a table view of applications, and the host view. See Viewing Applications on page 1425 for more information.
Application Summary	You can use this workflow to obtain detailed information about the applications and associated hosts on your network, so you can closely examine host application activity. This workflow begins with a list of individual host IP addresses running applications, followed by a table view of applications and the host view.
Applications	You can use this workflow to analyze running applications on your network, so you can gain an overview of how the network is being used. This workflow begins with a count of hosts running individual applications, followed by a table view of applications and the host view. See Viewing Applications on page 1425 for more information.

Predefined Application Details Workflows

LICENSE: FireSIGHT

The [Predefined Application Details Workflows](#) table describes the predefined workflows that you can use with application detail and client data.

Predefined Application Details Workflows

WORKFLOW NAME	DESCRIPTION
Application Details	You can use this workflow to analyze the client applications on your network in more detail. This workflow contains a series of pages that begin with a list of the client applications and application products on your network and a count of the number of hosts running each application. You can then view the number of hosts running each version of that application. The next page lets you identify which applications have been accessed most frequently on specific hosts. The workflow then provides a table view of client applications, followed by the host view. See Viewing Application Details on page 1429 for more information.
Clients	This workflow contains a table view of client applications, followed by the host view. See Viewing Application Details on page 1429 for more information.

Predefined Servers Workflows

LICENSE: FireSIGHT

The [Predefined Servers Workflows](#) table describes the predefined workflows that you can use with server data.

Predefined Servers Workflows

WORKFLOW NAME	DESCRIPTION
Network Applications by Count	You can use this workflow to analyze the most frequently used applications on your network. This workflow contains a series of pages that show applications with a count of hosts where each application occurs, then add the vendor and version of each application. The workflow then concludes with a table view listing the applications per host, followed by the host view. See Viewing Servers on page 1418 for more information.
Network Applications by Hit	You can use this workflow to analyze the most active applications on your network. This workflow contains a series of pages that show applications with a count of how often each application is accessed, then add the vendor and version information for each application. The workflow finishes with a page containing a table view listing the applications per host, followed by the host view. See Viewing Servers on page 1418 for more information.
Server Details	You can use this workflow to analyze the vendors and versions of detected server application protocols in detail. The workflow contains a list of servers associated with their vendors, then a list of servers correlated with both vendor and version, finishing with a table view of servers and the host view.
Servers	This workflow contains a table view of applications followed by the host view. See Viewing Servers on page 1418 for more information.

Predefined Host Attributes Workflows

LICENSE: FireSIGHT

The [Predefined Host Attributes Workflows](#) table describes the predefined workflow that you can use with host attribute data.

Predefined Host Attributes Workflows

WORKFLOW NAME	DESCRIPTION
Attributes	You can use this workflow to monitor IP addresses of hosts on your network and the hosts' status. This workflow begins with a table view of host attributes that lists individual IP addresses with current user, host criticality, notes, and white list compliance. It finishes with the host view. For more information, see Viewing Host Attributes on page 1412.

Predefined Discovery Events Workflows

LICENSE: FireSIGHT

The [Predefined Discovery Event Workflows](#) table describes the predefined workflow that you can use with discovery event data.

Predefined Discovery Event Workflows

WORKFLOW NAME	DESCRIPTION
Discovery Events	This workflow provides a detailed list, in table view form, of discovery events, followed by the host view. For more information, see Understanding the Discovery Events Table on page 1397.

Predefined User Workflows

LICENSE: FireSIGHT

The [Predefined User Workflows](#) table describes the predefined user workflows included on the Defense Center.

Predefined User Workflows

WORKFLOW NAME	DESCRIPTION
Users	This workflow provides a list of user information collected from user events or from the LDAP server connection. For details about the user identity workflow, see Viewing Users on page 1447.

Predefined Vulnerabilities Workflows

LICENSE: FireSIGHT

The [Predefined Vulnerabilities Workflows](#) table describes the predefined vulnerabilities workflow included on the Defense Center.

Predefined Vulnerabilities Workflows

WORKFLOW NAME	DESCRIPTION
Vulnerabilities	You can use this workflow to review s a table view of vulnerabilities showing all the vulnerabilities in the database, followed by a table view of only those active vulnerabilities that apply to the detected hosts on your network. The workflow ends in a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints. For more information, see Viewing Sourcefire Vulnerabilities on page 1435.

Predefined Third-Party Vulnerabilities Workflows

LICENSE: FireSIGHT

The [Predefined Third-Party Vulnerabilities Workflows](#) table describes the predefined third-party vulnerabilities workflows included on the Defense Center.

Predefined Third-Party Vulnerabilities Workflows

WORKFLOW NAME	DESCRIPTION
Vulnerabilities by IP Address	You can use this workflow to see quickly how many third-party vulnerabilities you have detected per host IP address on your monitored network. The workflow concludes with a table view of third-party vulnerabilities, followed by the host view. For more information, see Viewing Third-Party Vulnerabilities on page 1441.
Vulnerabilities by Source	You can use this workflow to see quickly how many third-party vulnerabilities you have detected per third-party vulnerability source, such as the QualysGuard Scanner. This workflow provides some details about those vulnerabilities on an intermediate drill-down page, then concludes with a table view of third-party vulnerabilities and the host view. For more information, see Viewing Third-Party Vulnerabilities on page 1441.

Predefined Correlation and White List Workflows

LICENSE: FireSIGHT

There is a predefined workflow for each type of correlation data, white list events, white list violations, and remediation status events.

Predefined Correlation Workflows

WORKFLOW NAME	DESCRIPTION
Correlation Events	This workflow contains a table view of correlation events. See Working with Correlation Events on page 1521 for more information.
White List Events	This workflow contains a table view of white list events. See Working with White List Events on page 1572 for more information.

Predefined Correlation Workflows (Continued)

WORKFLOW NAME	DESCRIPTION
Host Violation Count	This workflow provides a series of pages that list all the host IP addresses that violate at least one white list. The first page sorts the addresses based on the number of violations per address, with the IP addresses with the most number of violations at the top of the list. If a host IP address violates more than one white list, there is a separate row for each violated white list. The workflow also contains a table view of white list violations that lists all violations, with the most recently detected violation at the top of the list. Each row in the table contains a single detected violation. See Working with White List Violations on page 1579 for more information.
White List Violations	This workflow includes a table view of white list violations that lists all violations with the most recently detected violation at the top of the list. Each row in the table contains a single detected violation. See Working with White List Violations on page 1579 for more information.
Status	This workflow contains a table view of remediation status, which includes the name of the policy that was violated and the name and status of the remediation that was applied. See Working with Remediation Status Events on page 1633 for more information.

Predefined System Workflows

LICENSE: Any

The Sourcefire 3D System is delivered with some additional workflows, including system events such as audit events and health events, as well as workflows that list results from rule update imports and active scans.

Additional Predefined Workflows

WORKFLOW NAME	DESCRIPTION
Audit Log	This workflow contains a table view of the audit log that lists audit events. See Viewing Audit Records on page 2188 for more information.
Health Events	This workflow displays events triggered by the health monitoring policy. See Working with the Health Events Table View on page 2178 for more information.
Rule Update Import Log	This workflow contains a table view listing information about both successful and failed rule update imports. For more information, see Importing Rule Updates and Local Rule Files on page 2072.
Scan Results	This workflow contains a table view listing each completed scan. For more information, see Working with Active Scan Results on page 1716.

Saved Custom Workflows

LICENSE: Protection + FireSIGHT

In addition to predefined workflows, which cannot be modified, your Defense Center includes several saved custom workflows. Each of these workflows is based on a custom table and can be modified. For information on accessing these workflows, see [Viewing a Workflow Based on a Custom Table](#) on page 1787.

Saved Custom Workflows

WORKFLOW NAME	DESCRIPTION
Events by Impact, Priority, and Host Criticality	<p>You can use this workflow to quickly pick out and focus in on hosts that are important to your network, currently vulnerable, and possibly currently under attack.</p> <p>By default, this workflow starts with a summary of events sorted by impact flag level, then by host criticality, and then by the number of occurrences of the event. You can use the second page of the workflow to drill down and view the source and destination addresses where specific events occur. The workflow concludes with a table view of Intrusion Events with Destination Criticality, then the packet view. This workflow is based on the Intrusion Events with Destination Criticality custom table. For more information, see Understanding Custom Tables on page 1780.</p>
Events by Priority and Classification	<p>This workflow lists events and their type in order of event priority, along with a count showing how many times each event has occurred.</p> <p>This workflow begins with a drill-down page that contains the priority level, classification and count of each listed event. The last pages in the workflow are the table view of events and the packet view. This workflow is based on the Intrusion Events custom table. For more information, see Understanding Custom Tables on page 1780.</p>
Events with Destination, Impact, and Host Criticality	<p>You can use this workflow to find the most recent attacks on hosts that are important to your network and currently vulnerable.</p> <p>By default, this workflow starts with a list of the most recent events, sorted by impact level. The next page of the workflow provides a table view of Intrusion Events with Destination Criticality, followed by the packet view. This workflow is based on the Intrusion Events with Destination Criticality custom table. For more information, see Understanding Custom Tables on page 1780.</p>
Hosts with Servers Default Workflow	<p>You can use this workflow to quickly view the basic information in the Hosts with Servers custom table.</p> <p>By default, this workflow begins with a table view of hosts with servers, followed by the host view. This workflow is based on the Hosts with Servers custom table. For more information, see Understanding Custom Tables on page 1780.</p>

Saved Custom Workflows (Continued)

WORKFLOW NAME	DESCRIPTION
Intrusion Events with Destination Criticality Default Workflow	<p>You can use this workflow to quickly view the basic information in the Intrusion Events with Destination Criticality custom table.</p> <p>By default, this workflow starts with a table view of Intrusion Events with Destination Criticality, followed by the packet view. This workflow is based on the Intrusion Events with Destination Criticality custom table. For more information, see Understanding Custom Tables on page 1780.</p>
Intrusion Events with Source Criticality Default Workflow	<p>You can use this workflow to quickly view the basic information in the Intrusion Events with Source Criticality custom table.</p> <p>By default, this workflow starts with a table view of Intrusion Events with Source Criticality, followed by the packet view. This workflow is based on the Intrusion Events with Source Criticality custom table. For more information, see Understanding Custom Tables on page 1780.</p>
Server and Host Details	<p>You can use this workflow to determine what servers are most frequently used on your network and which hosts are running those servers.</p> <p>By default, this workflow begins with a summary of servers with the frequency of access for each service. The next page lists servers by operating system vendor and version. The workflow concludes with a table view of hosts with servers, followed by the host view. This workflow is based on the Hosts with Servers custom table. For more information, see Understanding Custom Tables on page 1780.</p>

Using Workflows

LICENSE: Any

The drill-down and table view pages in workflows allow you to quickly narrow your view of the data so you can zero in on events that are significant to your analysis. Although the data in each type of workflow is different, all workflows share a common set of features. The following sections describe these features and explain how to use them:

- [Selecting Workflows](#) on page 1809 describes the workflow selection page and how to select a workflow to use.
- [Understanding the Workflow Toolbar](#) on page 1812 describes the toolbar options available in workflows.
- [Using Workflow Pages](#) on page 1812 describes the features that appear on all workflow pages and explains how to use them.
- [Setting Event Time Constraints](#) on page 1820 describes how to set the time range for event-based workflows. The workflow includes events generated in the specified time range.

- [Constraining Events](#) on page 1830 describes features that are used in workflows to constrain, or narrow, the view of data in workflows and to advance through workflow pages.
- [Using Compound Constraints](#) on page 1832 explains how compound constraints can be used and provides examples.
- [Sorting Drill-Down Workflow Pages](#) on page 1835 describes features for sorting the data displayed in workflows, and for removing and restoring table columns to view.
- [Selecting Rows on a Workflow Page](#) on page 1835 describes how to select data rows in the displayed table that you want to analyze or on which you want to perform some other action.
- [Navigating to Other Pages in the Workflow](#) on page 1836 describes how to open other workflows using the constraints, including any selected events, from the current workflow.
- [Navigating Between Workflows](#) on page 1836 describes the links that appear above the column headings of a workflow page on the Defense Center and explains how you can use them to apply the current constraints to a different workflow.
- [Searching for Events](#) on page 1770 provides information about the feature used to search event data.
- [Using Bookmarks](#) on page 1838 describes how to create, manage, and use bookmarks.

Selecting Workflows

LICENSE: Any

The Sourcefire 3D System provides predefined workflows for the types of data listed in the [Features Using Workflows](#) table.

Features Using Workflows

FEATURE	MENU PATH	OPTION
Intrusion events	Analysis > Intrusion	Events Reviewed Events Clipboard Incidents
Malware events	Analysis > Malware	n/a
File events	Analysis > Files	n/a
Connection events	Analysis > Connections	n/a

Features Using Workflows (Continued)

FEATURE	MENU PATH	OPTION
Host events	Analysis > Hosts	Network Map Hosts Applications Application Details Servers Host Attributes Discovery Events
User events	Analysis > Users	User Activity Users
Vulnerability events	Analysis > Vulnerabilities	Vulnerabilities Third-Party Vulnerabilities
Correlation events	Analysis > Correlation	Correlation Events White List Events White List Violations Status
Audit events	System > Monitoring	Audit
Health events	Health > Health Events	n/a
Rule Update Import Log	System > Updates	n/a
Scan Results	Policies > Actions > Scanners	n/a

When you view any of the kinds of data described in the above table, events appear on the first page of the default workflow for that data.

Also note that workflow access depends on your user role (see [Configuring User Roles](#) on page 1905), as follows:

- Administrator users can access any workflow, and are the only users who can access the audit log, scan results, and the rule update import log.
- Maintenance Users can access health events.
- Security Analyst and Security Analyst (Read Only) users can access intrusion, malware, file, connection, discovery, vulnerability, correlation, and health workflows.

To view the data using a workflow other than the default:

ACCESS: Admin/Any Security Analyst

1. Select the appropriate menu path and option as described in the [Features Using Workflows](#) table.

The first page of the default workflow for that data type appears. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.

2. Optionally, use a different workflow. Click (**switch workflow**) next to the workflow title, then select the workflow you want to use.

For example, the following graphic shows the workflows you can use to view intrusion events.



3. The first page of your selected workflow appears.

Understanding the Workflow Toolbar

LICENSE: Any

Each page in a workflow includes a toolbar that offers quick access to related features. The [Workflow Toolbar Links](#) table describes each of the links on the toolbar.

Bookmark This Page Report Designer Dashboard View Bookmarks Search ▼

Workflow Toolbar Links

FEATURE	DESCRIPTION
Bookmark This Page	Bookmarks the current page so you can return to it later. Bookmarking captures the constraints in effect on the page you are viewing so you can return to the same data (assuming the data still exists) at a later time. See Using Bookmarks on page 1838 for information about creating bookmarks.
Report Designer	Opens the report designer with the currently constrained workflow as the selection criteria. See Creating a Report Template from an Event View on page 1725 for information about creating reports.
Dashboard	Opens a dashboard relevant to your current workflow. For example, Connection Events workflows link to the Connection Summary dashboard. See Using Dashboards on page 66 for information about using dashboards.
View Bookmarks	Displays a list of saved bookmarks from which you can select. See Using Bookmarks on page 1838 for information about creating and managing bookmarks.
Search	Displays a Search page where you can perform advanced searches on data in the workflow. You can also click the down arrow icon to select and use a saved search. See Searching for Events on page 1770 for information about searching workflows.

Using Workflow Pages

LICENSE: Any

The actions you can perform on a workflow page depend on the type of page. Table view pages and drill-down pages contain many features you can use to constrain the set of events you want to view or to navigate the workflow. For more information on the features available on each type of page, see the following sections:

- [Using Common Table View or Drill-Down Page Functionality](#) on page 1813
- [Using Geolocation](#) on page 1816
- [Using Table View Pages](#) on page 1818

- [Using Drill-Down Pages](#) on page 1819
- [Using the Host View, Packet View, or Vulnerability Detail Pages](#) on page 1819

Using Common Table View or Drill-Down Page Functionality

LICENSE: Any

Table view and drill-down workflow pages provide a set of icons and other features in the table header and table rows that you can use to perform actions on the displayed data.

Discovery Events 2011-11-30 09:04:29 - 2011-11-30 10:15:09 Expanding

Table View of Events ▸ Hosts

No Search Constraints (Edit Search)

Connection Events	Intrusion Events	FireAMP Events	Hosts	Applications	Application Details	Servers	Host Attributes	More ▸ Discovery
<input type="checkbox"/>	Time ×	Event ×	IP Address ×	User ×	MAC Address ×	MAC Vendor ×	Port ×	Description ×
<input type="checkbox"/>	2011-11-30 10:13:08	Client Timeout	10.10.10.4					SSH OpenSSH C
<input type="checkbox"/>	2011-11-30 10:13:08	TCP Port Timeout	10.10.10.10					
<input type="checkbox"/>	2011-11-30 10:13:08	TCP Server Connection Update	10.10.10.7			VMware, Inc.		SSH
<input type="checkbox"/>	2011-11-30 10:10:55	New Host	10.10.10.7			Cisco Systems		
<input type="checkbox"/>	2011-11-30 10:10:55	TCP Server Information Update	10.10.10.9			VMware, Inc.		HTTP Apache
<input type="checkbox"/>	2011-11-30 10:10:55	TCP Server Information Update	10.10.10.8			Cisco Systems		HTTP Apache

« Page 1 of 80 » Displaying rows 1-25 of 1995 rows

View Delete
View All Delete All

The features are described in the [Table View and Drill-Down Page Features](#) table.

Table View and Drill-Down Page Features

FEATURE	DESCRIPTION
	<p>Click the blue down-arrow icon to display the corresponding row in the next page of the workflow.</p>
	<p>Click the network file trajectory icon, which appears in file name and SHA-256 hash value columns, to view the file's trajectory map in a new window. For more information, see Analyzing Network File Trajectory on page 1244.</p> <p>Note that because neither the DC500 Defense Center nor Series 2 devices support network-based malware protection, you cannot view network file trajectory for network-based malware and file events on these appliances.</p>
  (blacklisted)  (blacklisted, set to monitor)	<p>Click the host profile icon, which appears in IP address columns, to display the host profile associated with that IP address in a pop-up window. For more information, see Using Host Profiles on page 1334.</p> <p>If the host profile icon is grayed out, you cannot view the host profile because that host cannot be in the network map (for example, 0.0.0.0).</p> <p>If you are performing traffic filtering based on Security Intelligence data, host icons next to blacklisted and monitored IP addresses in the connection event view look slightly different. This helps you identify which host in a connection was blacklisted. Note that neither the DC500 Defense Center nor Series 2 devices support Security Intelligence data.</p>
	<p>Click the user icon, which appears in user identity columns, to view user profile information. For more information, see Understanding User Details and Host History on page 1449.</p> <p>If the user icon is grayed out, you cannot view the user profile because that user cannot be in the database (FireAMP Connector user).</p>
	<p>Click the vulnerability icon, which appears in third-party vulnerability ID columns, to view vulnerability details for third-party vulnerabilities. For more information, see Viewing Vulnerability Details on page 1365.</p>
<p>Check boxes</p>	<p>Select the check boxes by two or more rows on a page to indicate which rows you want to affect, then click one of the buttons at the bottom of the page (for example, the View button). You can also select the check box at the top of the row to select all the rows on the page.</p>

Table View and Drill-Down Page Features (Continued)

FEATURE	DESCRIPTION
Country flags and codes	<p>In some workflow pages, such as those for connection events, intrusion events, file events, and malware events, routable IP addresses include information about the associated country. When this <i>geolocation</i> information is available, the country's flag and ISO code appear in the appropriate column (such as Source Country). Hover your pointer over the flag to view the country name. When viewing individual (rather than aggregated) data points, you can click the flag icon to view further geolocation details. See Using Geolocation on page 1816 for more information.</p> <p>Note that the DC500 Defense Center does not support geolocation data.</p>
Search Constraints	<p>Lists the values, if present, constraining the data view. Click the expand arrow (►) to display the active constraints and disabled columns list or the collapse arrow (▼) to hide the list from view. By default, this list is collapsed, which is useful when the list of constraints is long and takes up too much of the screen.</p> <p>To remove a single constraint, click it. To remove a compound constraint, click Compound Constraints.</p> <p>Click Edit Search or Save Search to open a search page pre-populated with the current single constraints. See Constraining Events on page 1830 for more information.</p> <p>IMPORTANT! Compound constraints are constraints created based on rows with multiple non-count values. You cannot perform a search or save a search on a compound constraint.</p>
Time Range	<p>The date range located in the upper right corner of the page sets a time range for events to include in the workflow. See Setting Event Time Constraints on page 1820 for more information.</p> <p>Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.</p>
Workflow Page Links	<p>Workflow page links appear in the upper left corner of predefined workflow table view and drill-down pages, above events and below the workflow name. Click a workflow page link to display that page using any active constraints.</p>
Workflow Name	<p>The name of the workflow appears at the top of the page. Beside it, when applicable, is the (switch workflows) link, which you can use to select other workflows of the same type.</p>

Using Geolocation

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

While monitoring your network, the *geolocation* feature provides you with additional data about the geographical sources of routable IP addresses (country, continent, and so on). You can use this data to determine if, for example, connections are originating from or terminating in countries unconnected with your organization. Geolocation information is available for intrusion events, connection events, file events, malware events, host profiles, and user profiles. Geolocation information is also available in the Context Explorer and the dashboard. Note that you cannot base any system rules (such as intrusion rules) on geolocation information.

By installing geolocation database (GeoDB) updates, you can view a Geolocation Details page with granular information available for an IP address, such as postal code, coordinates, time zone, Autonomous System Number (ASN), internet service provider (ISP), use type (home or business), organization, domain name, connection type, and proxy information. You can also pinpoint the detected location with any of four third-party map tools. Without a GeoDB update, only the flag icon and country name appear; you cannot view the Geolocation Details page. For information on installing and updating the GeoDB, see [Updating the Geolocation Database](#) on page 2092. You can view the current version of your GeoDB update by clicking **Help > About**.

Depending on availability, a number of fields may appear on the Geolocation Details page; fields with no information are not displayed. The [Geolocation Detail Fields](#) table contains information on these fields.

Geolocation Detail Fields

FIELD	CONTENTS
Country	Country associated with the host's IP address, accompanied by the country's flag. The continent is listed in parentheses. Examples: United States (North America) , Equatorial Guinea (Africa)
Region	State, province, or other subregion of the country where the host is located. Examples: VA , 35
City	City where the host is located. Examples: Seattle , Fukuoka
Postal Code	Postal code of the region where the host is located. Examples: 361000 , 90210
Latitude/Longitude	Exact coordinates of the host's location. Examples: 40.0375 , -76.1053 ; 53.4050 , -0.5484
Maps	Links to external mapping sites (Google Maps, Yahoo Maps, Bing Maps, and OpenStreetMap). Click any link to view a contextual map of the host's approximate location.

Geolocation Detail Fields (Continued)

FIELD	CONTENTS
Timezone	Time zone of the host's location, with Daylight Savings Time noted where applicable. Examples: GMT+8:00, GMT-4:00 (In DST)
ASN	Autonomous System Number (ASN) associated with the host's IP address, and any additional information about that ASN. Examples: 14618 (Amazon.com Inc.); 4837 (Cncgroup China169 Backbone)
ISP	Internet service provider (ISP) associated with the host's IP address. Examples: Atlantic Broadband; China Unicom Ip Network
Home/Business	Whether the host's connection is used for Home or Business purposes.
Organization	Organization associated with the host's IP address. Examples: Amazon.com, Bank of America
Domain Name	Domain name associated with the host's IP address. Examples: amazonaws.com, xmcnc.net
Connection Type	Connection type associated with the host's IP address. Examples: Broadband, DSL
Proxy Type	The type of proxy used. Examples: Anonymous, Corporate

To view geolocation details:

ACCESS: Any

- ▶ In an event view, host profile, or other geolocation-supporting page, click the small country flag icon or ISO country code that appears by an individual data point. (You cannot view geolocation details for aggregate geolocation information, such as on the Connection Summary dashboard, despite the presence of flag icons.)

TIP! In event views, hover your pointer over the flag icon to view a tooltip with the country's name.

The Geolocation Details page appears in a new window.



Geolocation for 72.00.00.000

Country	United States  (North America)
Region	VA
City	Ashburn
Postal Code	20147
Latitude/Longitude	39.0329, -77.4866
Maps	
Timezone	GMT-4:00 (In DST)
▼ Additional Information	
ASN	14618 (Amazon.com Inc.)
ISP	Amazon.com Inc.
Home/Business	Business
Organization	Amazon.com
Domain Name	amazonaws.com
Connection Type	Broadband
Proxy Type	Corporate

Using Table View Pages

LICENSE: Any

Table views include a column for each of the fields in the database if the column is enabled by default. Note that when you disable a column on a table view, the Sourcefire 3D System adds the Count column to the event view if disabling the column would create two or more identical rows. When you click on a value in a table view page, you constrain by that value. When you create a custom workflow, you add a table view to it by clicking **Add Table View**.

Table view pages provide some additional features not available on drill-down, host view, packet view, or vulnerability detail pages. The [Additional Table View Page Features](#) table provides more information on those features.

Additional Table View Page Features

FEATURE	DESCRIPTION
	<p>Click this icon in the column heading that you want to hide. In the pop-up window that appears, click Apply.</p> <p>TIP! To hide or show other columns, select or clear the appropriate check boxes before you click Apply.</p>
Disabled Columns list	<p>When you remove columns from a page, or columns are disabled by default, the column names appear in the Disabled Columns list, which is located above the table and hidden by default.</p> <p>To add a disabled column back to the event view, click the Search Constraints expand arrow () to expand the search constraints, then click the column name under Disabled Columns.</p> <p>See Sorting Drill-Down Workflow Pages on page 1835 for more information.</p>

Using Drill-Down Pages

LICENSE: Any

Drill-down pages contain a subset of columns that are available in the database. Note that drill-down pages for predefined workflows always have a Count column. Drill-down pages allow you to narrow the scope of events you are viewing and to move forward in the workflow. If you click on a value in a drill-down page, for example, you constrain by that value and move to the next page in the workflow, focusing more closely on events that match your selected values. Clicking a value in a drill-down page does not disable the column where the value is, even if the page you advance to is a table view. When you create a custom workflow, you add a drill-down page to it by clicking **Add Page**.

For more information on using features on drill-down pages to constrain the set of events as you go through a workflow, see [Using Common Table View or Drill-Down Page Functionality](#) on page 1813.

Using the Host View, Packet View, or Vulnerability Detail Pages

LICENSE: Any

The final page in a discovery event, host, host attributes, servers, client applications, or connection data workflow is the host view. The final page in a vulnerability workflow is the vulnerability detail page. An intrusion event workflow always ends with the packet view. On the final page of a workflow, you can

expand detail sections to view specific information about each object in the set you focused on over the course of the workflow. Although the web interface does not list the constraints on the final page of a workflow, previously set constraints are retained and applied to the set of data.

Setting Event Time Constraints

LICENSE: Any

Each event has a time stamp that indicates when the event occurred. You can constrain the information that appears in some workflows by setting the time window, sometimes called the time range.

Workflows based on events that can be constrained by time include a time range line at the top of the page, as shown in the following graphic.



2012-06-08 13:06:00 - 2012-06-08 15:38:10
Expanding

By default, workflows on Sourcefire appliances use an expanding time window set to the past hour. For example, if you log in at 11:30 AM, you will see events that occurred between 10:30 AM and 11:30 AM. As time moves forward, the time window expands. At 12:30 PM, you will see events that occurred between 10:30 AM and 12:30 PM.

You can change this behavior by setting your own default time window, which governs three properties:

- time window type (static, expanding, or sliding)
- time window length
- the number of time windows (either multiple time windows or a single global time window)

For general information on the default time window, see [Default Time Windows](#) on page 2219.

Regardless of the default time window setting, you can manually change the time window during your event analysis by clicking the time range at the top of the page, which displays the Date/Time pop-up window. Depending on the number of time windows you configured and the type of appliance you are using, you can also use the Date/Time window to change the default time window for the type of event you are viewing.

Finally, you can pause the time window, which allows you to examine the data provided by the workflow without the time window changing and removing or adding events that you are not interested in. Note that to avoid displaying the same events on different workflow pages, the time window automatically pauses when you click a link at the bottom of the page to display another page of events; you can unpause the time window when you are ready.

For more information, see the following sections:

- [Changing the Time Window](#) on page 1821
- [Changing the Default Time Window for Your Event Type](#) on page 1826
- [Pausing the Time Window](#) on page 1829

Changing the Time Window

LICENSE: Any

Regardless of the default time window, you can manually change the time window during your event analysis.

IMPORTANT! Manual time window settings are valid for only the current session. When you log out and then log back in, time windows are reset to the default.

Depending on the number of time windows you configured, changing the time window for one workflow may affect other workflows on the appliance. For example, if you have a single, global time window, changing the time window for one workflow changes it for all other workflows on the appliance. On the other hand, if you are using multiple time windows, changing the audit log or health event workflow time windows has no effect on any other time window, while changing the time window for other kinds of events affects all events that can be constrained by time (with the exception of audit events and health events).

Note that because not all workflows can be constrained by time, time window settings have no effect on workflows based on hosts, host attributes, applications, application details, vulnerabilities, users, or white list violations.

Use the Time Window tab on the Date/Time window to manually configure a time window. Depending on the number of time windows you configured in your default time window settings, the tab's title is one of the following:

- **Events Time Window**, if you configured multiple time windows and are setting the time window for a workflow other than the audit log or health events workflow
- **Health Monitoring Time Window**, if you configured multiple time windows and are setting the time window for the health events workflow
- **Audit Log Time Window**, if you configured multiple time windows and are setting the time window for the audit log
- **Global Time Window**, if you configured a single time window

The first decision you must make when configuring a time window is the type of time window you want to use:

- A *static* time window displays all the events generated from a specific start time to a specific end time.
- An *expanding* time window displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view.
- A *sliding* time window displays all the events generated from a specific start time (for example, one week ago) to the present; as time moves forward, the time window “slides” so that you see only the events for the range you configured (in this example, for the last week).

Depending on what type you select, the Date/Time window changes to give you different configuration options. The following graphic shows the Date/Time window, specifying that you want to use an expanding time window. With expanding time windows, the End Time calendar is grayed out and specifies that the end time is “Now.”

The screenshot shows the 'Events Time Window' configuration interface. At the top, there are two tabs: 'Events Time Window' (selected) and 'Preferences'. Below the tabs is a dropdown menu set to 'Expanding Time Window'. The interface is divided into two main sections: 'Start Time' and 'End Time'. The 'Start Time' section features a calendar for October 2011 with the 14th selected, and a time selector set to 14:25. The 'End Time' section has a grayed-out calendar with the word 'Now' overlaid, indicating that the end time is the current time. Below these sections, the configured time window is displayed as '2011-10-14 14:25' to '2011-10-14 16:19', with a duration of '1 hour, 54 minutes'. A 'Presets' section offers options like 'Last', 'Current', and 'Synchronize with'. At the bottom, there are 'Apply' and 'Reset' buttons, and a note: 'Any changes made will take effect on the next page load.'

If you use a static time window, you can set an end time.

The screenshot shows a configuration window titled "Static Time Window". It features two calendar views for October 2011. The left calendar is labeled "Start Time" and has the date 14 selected. Below it, the time is set to 14:25. The right calendar is labeled "End Time" with a checked checkbox, and the date 15 is selected. Below it, the time is set to 15:25.

If you choose to use a sliding time window, your options change further.

The screenshot shows a configuration window titled "Events Time Window" with a "Preferences" tab. It features a "Sliding Time Window" dropdown menu. Below it, there is a field "Show the Last" followed by an input box and a "month(s)" dropdown menu. A red error message reads "Please enter a valid Integer." Below this is a "Presets" section with buttons for "Last", "1 hour", "6 hours", "1 day", "1 week", "2 weeks", and "1 month". At the bottom, there are "Synchronize with" options for "Audit Log Time Window" and "Health Monitoring Time Window", along with "Apply" and "Reset" buttons. A note at the bottom states "Any changes made will take effect on the next page load."

IMPORTANT! The Sourcefire 3D System uses a 24-hour clock based on the time you specified in your time zone preferences. See [Setting Your Default Time Zone](#) on page 2222 for information about configuring a time zone.

The [Time Window Settings](#) table explains the various settings you can configure on the Time Window tab.

Time Window Settings

SETTING	TIME WINDOW TYPE	DESCRIPTION
time window type drop-down list	n/a	Select the type of time window you want to use: static, expanding, or sliding. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
Start Time calendar	static and expanding	Specify a start date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). TIP! Instead of using the calendar, you can use the Presets options, described below.
End Time calendar	static	Specify an end date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Note that If you are using an expanding time window, the End Time calendar is grayed out and specifies that the end time is "Now." TIP! Instead of using the calendar, you can use the Presets options, described below.
Show the Last field and drop-down list	sliding	Configure the length of the sliding time window.

Time Window Settings (Continued)

SETTING	TIME WINDOW TYPE	DESCRIPTION
Presets: Last	all	Click one of the time ranges in the list to change the time window, based on the local time of the appliance. For example, clicking 1 week changes the time window to reflect the last week. Clicking a preset changes the calendars to reflect the preset you choose.
Presets: Current	static and expanding	Click one of the time ranges in the list to change the time window, based on the local time and date of the appliance. Clicking a preset changes the calendars to reflect the preset you choose. Note that: <ul style="list-style-type: none"> • the current day begins at midnight • the current week begins at midnight Sunday • the current month begins at midnight on the first of the month
Presets: Synchronize with	all (not available if you are using a global time window)	Click one of: <ul style="list-style-type: none"> • Events Time Window to synchronize the current time window with the events time window • Health Monitoring Time Window to synchronize the current time window with the health monitoring time window • Audit Log Time Window to synchronize the current time window with the audit log time window

To change the time window during event analysis:

ACCESS: Admin/Maint/Any Security Analyst

1. On a workflow constrained by time, click the time range icon (🕒).
The Date/Time window appears.
2. On the **Time Window** tab, set the time window as described in the [Time Window Settings table](#) on page 1824.

TIP! Click **Reset** to change the time window back to the default settings.

3. Click **Apply**.
The window closes and the event view page displays events from the new time range.

Changing the Default Time Window for Your Event Type

LICENSE: Any

During your event analysis, you can use the Preferences tab on the Date/Time window to change the default time window for the type of event you are viewing without having to use the event view settings (see [Default Time Windows](#) on page 2219).

Keep in mind that changing the default time window in this way changes the default time window for only the type of event you are viewing. For example, if you configured multiple time windows, changing the default time window on the Preferences tab changes the settings for either the events, health monitoring, or audit log window, in other words, whichever time window is indicated by the first tab. If you configured a single time window, changing the default time window on the Preferences tab changes the default time window for all types of events.

The following graphic shows the Defense Center version of the Preferences tab, on an appliance that has multiple time windows configured.

The screenshot shows the 'Events Time Window' Preferences tab. The title is 'Time Window Preferences'. It contains the following settings:

- Refresh Interval (minutes): 0 (Set to 0 to disable)
- Number of Time Windows: Single Multiple
- Default Time Window: Show the Last - Static/Expanding (dropdown menu)
- 1 hour(s) (input field and dropdown menu)
- Use End Time

At the bottom right is a 'Save Preferences' button. At the bottom center is the text: 'Any changes made will take effect on the next page load.'

The [Time Window Preferences](#) table explains the various settings you can configure on the Preferences tab.

Time Window Preferences

PREFERENCE	DESCRIPTION
Refresh Interval	Sets the refresh interval for event views, in minutes. Entering zero disables the refresh option.
Number of Time Windows	Specify how many time windows you want to use: <ul style="list-style-type: none">• Select Multiple to configure separate default time windows for the audit log, for health events, and for workflows based on events that can be constrained by time.• Select Single to use a global time window that applies to all events,
Default Time Window: Show the Last - Sliding	<p>This setting allows you to configure a sliding default time window of the length you specify.</p> <p>The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window "slides" so that you always see events from the last hour.</p>

Time Window Preferences (Continued)

PREFERENCE	DESCRIPTION
Default Time Window: Show the Last - Static/ Expanding	<p>This setting allows you to configure either a static or expanding default time window of the length you specify.</p> <p>For static time windows (enable the Use End Time check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.</p> <p>For expanding time windows (disable the Use End Time check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the present. As you change event views, the time window expands to the present time.</p>
Default Time Window: Current Day - Static/ Expanding	<p>This setting allows you to configure either a static or expanding default time window for the current day. The current day begins at midnight, based on the time zone setting for your current session.</p> <p>For static time windows (enable the Use End Time check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.</p> <p>For expanding time windows (disable the Use End Time check box), the appliance displays all the events generated from midnight to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 24 hours before you log out, this time window can be more than 24 hours.</p>
Default Time Window: Current Week - Static/ Expanding	<p>This setting allows you to configure either a static or expanding default time window for the current week. The current week begins at midnight on the previous Sunday, based on the time zone setting for your current session.</p> <p>For static time windows (enable the Use End Time check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.</p> <p>For expanding time windows (disable the Use End Time check box), the appliance displays all the events generated from midnight Sunday to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 1 week before you log out, this time window can be more than 1 week.</p>

To change time window preferences during event analysis:

ACCESS: Admin/Maint/Any Security Analyst

1. On a workflow constrained by time, click the time range icon (🕒).
The Date/Time window appears.

2. Select the **Preferences** tab and change your preferences, as described in the [Time Window Preferences table](#) on page 1827.
3. Click **Save Preferences**.
Your preferences are saved.
4. You have two options:
 - To apply your new default time window settings to the event view you are using, click **Apply** to close the Date/Time window and refresh the event view.
 - To continue with your analysis without applying the default time window settings, close the Date/Time window without clicking **Apply**.

Pausing the Time Window

LICENSE: Any

You can pause the time window, which allows you to examine a snapshot of the data provided by the workflow. This is useful because when an unpaused workflow updates, it may remove events that you want to examine or add events that you are not interested in.

Note that you cannot pause a static time window. In addition, pausing an event time window has no effect on dashboards, nor does pausing a dashboard have any effect on pausing an event time window.

When you are finished with your analysis, you can unpause the time window. Unpausing the time window updates it according to your preferences, and also updates the event view to reflect the unpaused time window.

If the database contains more events than can be displayed on a single workflow page, you can click the links at the bottom of the page to display more events (see [Navigating to Other Pages in the Workflow](#) on page 1836). When you do this, the time window automatically pauses so that you do not see the same events twice. You can unpause the time window when you are ready.

To pause the time window:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ On the time range control, click the pause icon (⏸).
- The time window is paused until you unpause it.

To unpause the time window:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ On the time range control, click the play icon (▶).
- The time window is unpaused and updates according to your preferences. The event view updates to reflect the current time window.

Constraining Events

LICENSE: Any

The information that you see on a workflow page is determined by the constraints that you impose. For example, when you initially open an event workflow, the information is constrained to events that were generated in the previous hour.

To advance to the next page in the workflow and constrain the data you are viewing by specific values, select the rows with those values on the page and click **View**. To advance to the next page in the workflow retaining the current constraints and carrying forward all events, select **View All**.

IMPORTANT! If you select a row with multiple non-count values and click **View**, you create a compound constraint. For more information on compound constraints, see [Using Compound Constraints](#) on page 1832.

There is a third method for constraining data in a workflow. To constrain the page to the rows with values that you selected and also add the selected value to the list of constraints at the top of the page, click a value within a row on the page.

For example, if you click **10.10.60.119** in the Initiator IP column on a page with the following events:

<input type="checkbox"/>	▼ First Packet ×	Action ×	Initiator IP ×	Responder IP ×	Source Port / ICMP Type ×
↓	2013-03-10 23:27:34	Block	10.10.60.119	10.1.1.57	820 / tcp
↓	2013-03-10 23:27:34	Block	10.10.60.119	10.1.1.57	820 / tcp
↓	2013-03-10 22:19:28	Block	10.10.60.119	10.1.1.57	753 (rrh) / tcp
↓	2013-03-10 16:13:39	Block	10.10.32.124	10.10.60.165	856 / tcp

...then the constrained page includes only the events with that IP address:

▼ Search Constraints (Edit Search Save Search)

Initiator IP

Connections						Intrusion	Malware	Files	Hosts	Applications	Application Details	Servers	Hc
<input type="checkbox"/>	▼ First Packet ×	Action ×	Initiator IP ×	Responder IP ×	Source Port / ICMP Type ×								
↓	2013-03-10 23:27:34	Block	10.10.60.119	10.1.1.57	820 / tcp								
↓	2013-03-10 23:27:34	Block	10.10.60.119	10.1.1.57	820 / tcp								
↓	2013-03-10 22:19:28	Block	10.10.60.119	10.1.1.57	753 (rrh) / tcp								
↓	2013-03-09 23:21:59	Block	10.10.60.119	10.1.1.57	822 / tcp								

TIP! The procedure for constraining connection events based on Monitor rule criteria is slightly different and you may need to take some extra steps. Additionally, you cannot constrain connection events by associated file or intrusion information. For more information, see [Working with Connection Data Tables](#) on page 570.

You can also use searches to constrain the information in a workflow. The search criteria you enter on the search page are listed as the constraints at the top of the page, with the resulting events constrained accordingly. On the Defense Center, the current constraints are also applied when navigating to other workflows, unless they are compound constraints (see [Navigating Between Workflows](#) on page 1836).

When searching, you must pay careful attention to whether your search constraints apply to the table you are searching. For example, client data is not available in connection summaries. If you search for connection events based on the detected client in the connection and then view the results in a connection summary event view, the Defense Center displays connection data as if you had not constrained it at all. Invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough, as shown in the following graphic.



The [Search Constraint Functions](#) table describes each of the actions you can perform when applying a constraint.

Search Constraint Functions

To...	CLICK...
constrain the view to events that match a single value	the value in the table. For example, if you are viewing a list of logged connections and want to constrain the list to only those you allowed using access control, click Allow in the Action column. As another example, if you are viewing intrusion events and want to constrain the list to only events where the destination port is 80, click 80 (http)/tcp in the DST Port/ICMP Code column.
constrain the view to events that match multiple values	the check box for events with those values and click View . Note that a compound constraint is added if the row contains multiple non-count values. For more information on compound constraints, see Using Compound Constraints on page 1832.
remove a constraint	the name of the constraint in the Search Constraints box.

Search Constraint Functions (Continued)

To...	Click...
edit constraints using the search page	Edit Search in the Search Constraints box. Use this feature when you want to constrain against multiple values in a single column. For example, if you want to view the events related to two IP addresses, click Edit Search , then modify the appropriate IP address field on the Search page to include both addresses, and then click Search .
save constraints as a saved search	Save Search in the Search Constraints box and give the query a name. Note that you cannot save queries containing compound constraints. For more information on compound constraints, see Using Compound Constraints on page 1832.
use the same constraints with another event view	the link at the top of the table that corresponds to the event view. See Navigating Between Workflows on page 1836 for more information. Note that you do not retain compound constraints when you switch to another workflow. For more information on compound constraints, see Using Compound Constraints on page 1832.
toggle the display of constraints	the expand arrow (►). This is useful when the list of constraints is large and takes up most of the screen.

Using Compound Constraints

LICENSE: Any

Compound constraints are based on all non-count values for a specific event. When you select a row with multiple non-count values, you set a compound constraint that only retrieves events matching all the non-count values in that row on that page. For example, if you select a row that has a source IP address of 10.10.31.17 and a destination IP address of 10.10.31.15 and a row that has a source IP address of 172.10.10.17 and a destination IP address of 172.10.10.15, you retrieve all of the following:

- Events that have a source IP address of 10.10.31.17 AND a destination IP address of 10.10.31.15
- OR**
- Events that have a source IP address of 172.10.31.17 AND a destination IP address of 172.10.31.15

When you combine compound constraints with simple constraints, the simple constraints are distributed across each set of compound constraints. If, for example, you added a simple constraint for a protocol value of `tcp` to the compound constraints listed above, you retrieve all of the following:

- Events that have a source IP address of 10.10.31.17 AND a destination IP address of 10.10.31.15 AND a protocol of `tcp`

OR

- Events that have a source IP address of 172.10.31.17 AND a destination IP address of 172.10.31.15 AND a protocol of `tcp`

You cannot perform a search or save a search on a compound constraint. You also cannot retain compound constraints when you use the event view links or click **(switch workflow)** to switch to another workflow. If you bookmark an event view with compound constraints applied, the constraints are not saved with the bookmark.

To clear all compound constraints, click **Compound Constraints**.

Sorting Table View Pages and Changing Their Layout

LICENSE: Any

When viewing data in a workflow, you can sort the data based on any available column and remove and restore columns to view. You can sort data in ascending or descending order by column.

TIP! If you create a custom workflow, you can fully customize the arrangement of columns on the pages and predefine the page sort order. See [Creating Custom Workflows](#) on page 1841 for more information.

Sorting and Layout Functions

To...	Click...
sort a column	<p>the column title. Click the column title again to reverse the sort order.</p> <p>TIP! The direction icon (▼) indicates which column the data is sorted by, and whether the sort is ascending (upward-pointing icon) or descending (downward-pointing icon).</p>
remove a column from a table view	<p>the close icon (✕) in the column heading that you want to hide. In the pop-up window that appears, click Apply.</p> <p>When you disable a column, it is disabled for the duration of your session (unless you add it back later). Note that when you disable the first column, the Count column is added. You cannot disable the Count column.</p> <p>TIP! To hide or show other columns, select or clear the appropriate check boxes before you click Apply. To add a disabled column back to the view, click the expand arrow (►) to expand the search constraints, then click the column name under Disabled Columns.</p>
add a disabled column back to the view	<p>the column name under Disabled Columns.</p> <p>When you enable a column that is disabled by default, it is enabled for the duration of your session (unless you disable it later). Note that the Count column is removed if enabling results in no identical rows.</p>

Sorting Drill-Down Workflow Pages

LICENSE: Any

When viewing data in a workflow or event view, you can sort the data based on any available column and remove and restore columns to view. You can sort data in ascending or descending order by column. The direction icon (▼) indicates which column the data is sorted by, and whether the sort is ascending (upward-pointing icon) or descending (downward-pointing icon).

TIP! If you create a custom workflow, you can fully customize the arrangement of columns on the pages and predefine the page sort order. See [Creating Custom Workflows](#) on page 1841 for more information.

To sort a column:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ Click the column title.

To reverse the sort order:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ Click the column title again.

Selecting Rows on a Workflow Page

LICENSE: Any

There are several different ways to select and then act on the rows on workflow pages:



- To select all rows on the page, select the check box at the top of the page. You can then click any of the buttons at the bottom of the page (**View**, **Delete**, and so on) to perform that action on all of the events on that page.

- To select a single row, select the check box next to the individual row. You can then click any of the buttons at the bottom of the page to perform that action on only the events associated with that row.
- To select a single row and view its associated events on the next page of the workflow, click the arrow icon (↘).

IMPORTANT! You cannot select rows from multiple pages at once.

Navigating to Other Pages in the Workflow

LICENSE: Any

If the database contains more events than can be displayed on a single workflow page, you can click the links at the bottom of the page to display more events.

Displaying rows 26–50 of 417 rows << Page 2 of 17 >>

When you click one of these links, the time window automatically pauses so that you do not see the same events twice; you can unpause the time window when you are ready. For more information, see [Setting Event Time Constraints](#) on page 1820.

The [Navigating Pages](#) table describes how to use the navigation links.

Navigating Pages

To...	CLICK...
view a different page	the page number, enter the page you wish to view, then press Enter
view the next page	>
view the previous page	<
jump to the last page	>
jump to the first page	<

Navigating Between Workflows

LICENSE: Any

You can navigate to other workflows using the links arrayed above the column headings on a workflow page. When too many links are applicable for the screen to display, a **More** drop-down tab appears on the far right of the link row. You can move your cursor over the **More** tab to view and select additional workflows.

When you click a navigation link, properties shared by the rows you select and the constraints you set are used in the new workflow, if they are applicable. If configured constraints or event properties do not map to fields in the new workflow, they are dropped. In addition, compound constraints are not retained when you switch from one workflow to another.

IMPORTANT! When you view event counts over a time range, the total number of events may not reflect the number of events for which more detailed data is available. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment. For more information, see [Logging Connection, File, and Malware Information](#) on page 515.

Note that unless you have either paused the time window or have configured a static time window, the time window changes when you change workflows. For more information, see [Setting Event Time Constraints](#) on page 1820.



The links provide quick access to workflow pages for the following:

- Connection Events
- Intrusion Events
- Malware
- Files
- Hosts
- Applications
- Application Details
- Servers
- Host Attributes
- Discovery Events
- Users
- Vulnerabilities
- Third-Party Vulnerabilities
- Correlation Events
- White List Events

This feature enhances your ability to investigate suspicious activity. For example, if you are viewing connection data and notice that an internal host is transmitting

an abnormally large amount of data to an external site, you can select the responder IP address and the port as constraints and then click **Applications**. The applications workflow will use the responder IP and port as IP Address and Port constraints and display additional information about the application, such as what kind of application it is. You can also click **Hosts** at the top of the page to view the host profile for the remote host.

After finding more information about the application, you can click **Correlation Events** to return to the connection data workflow, remove the Responder IP from the constraints, add the Initiator IP to constraints, and click **Application Details** to see what client the user on the initiating host used when transferring data to the remote host. Note that the Port constraint is not transferred to the Application Details page. While keeping the local host as a constraint, you can also use other navigation buttons to find additional information:

- To discover if any policies have been violated by the local host, keep the IP address as a constraint and click **Correlation Events**.
- To find out if an intrusion rule triggered against the host, indicating a compromise, click **Intrusion Events**.
- To view the host profile for the local host and determine if the host is susceptible to any vulnerabilities that may have been exploited, click **Hosts**.

Using Bookmarks

LICENSE: Any

Create a bookmark if you want to return quickly to a specific location and time in an event analysis. Bookmarks retain information about:

- the workflow you are using
- the part of the workflow you are viewing
- the page number within the workflow
- any search constraints
- any disabled columns
- the time range you are using

The bookmarks you create are available to all user accounts with bookmark access. This means that if you uncover a set of events that require more in-depth analysis, you can easily create a bookmark and turn over the investigation to another user with the appropriate privileges.

IMPORTANT! If the events that appear in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.

See these sections for more information about using bookmarks:

- [Creating Bookmarks](#) on page 1839 describes how to create a new bookmark.
- [Viewing Bookmarks](#) on page 1839 describes how to view and use existing bookmarks.
- [Deleting Bookmarks](#) on page 1840 describes how to delete bookmarks.

Creating Bookmarks

LICENSE: Any

Use the following procedure to create a new bookmark.

To create a bookmark:

ACCESS: Admin/Maint/Any Security Analyst

1. During an event analysis, with the events of interest displayed, click **Bookmark This Page**.

The Create a Bookmark page appears.



The screenshot shows a form titled "Create a Bookmark". It has a blue header bar with the text "Bookmark Name" and an adjacent text input field. Below the input field is a button labeled "Save Bookmark".

2. In the **Bookmark Name** field, type a name (up to 80 alphanumeric characters and spaces) for the bookmark, then click **Save Bookmark**.

The bookmark is saved and the event page you bookmarked appears again.

Viewing Bookmarks

LICENSE: Any

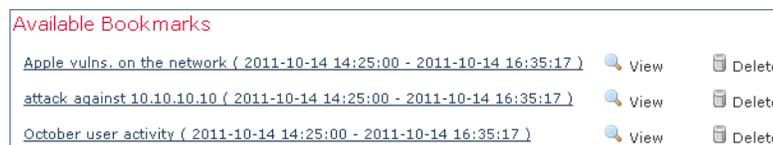
Use the following procedure to view and use existing bookmarks.

To view a bookmark:

ACCESS: Admin/Maint/Any Security Analyst

1. From any event view, click **View Bookmarks**.

The Bookmarks page appears. The Defense Center version of the page is displayed below.



The screenshot shows a table titled "Available Bookmarks". It contains three rows of bookmark information. Each row includes the bookmark name in a blue link, a time range in parentheses, a "View" icon, and a "Delete" icon.

Bookmark Name	Time Range	View	Delete
Apple vulns. on the network	(2011-10-14 14:25:00 - 2011-10-14 16:35:17)	View	Delete
attack against 10.10.10.10	(2011-10-14 14:25:00 - 2011-10-14 16:35:17)	View	Delete
October user activity	(2011-10-14 14:25:00 - 2011-10-14 16:35:17)	View	Delete

2. Next to the bookmark you want to use, click **View**.
The page you bookmarked appears.

IMPORTANT! If the events that originally appeared in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.

Deleting Bookmarks

LICENSE: Any

Use the following procedure to delete bookmarks. Note that deleting a bookmark does not affect the events retrieved by that bookmark.

To delete a bookmark:

ACCESS: Admin/Maint/Any Security Analyst

1. From any event view, click **View Bookmarks**.
The Bookmarks page appears.
2. Click **Delete** next to the bookmark you want to remove.
The bookmark is deleted.

Using Custom Workflows

LICENSE: Any

If the predefined and Sourcefire-provided custom workflows do not meet your needs, you can create custom workflows.

For more information, see:

- [Creating Custom Workflows](#) on page 1841 for a procedure to create custom workflows
- [Creating Custom Connection Data Workflows](#) on page 1843 for a procedure to create a custom workflow based on connection data
- [Viewing Custom Workflows](#) on page 1846 for procedures for viewing custom workflows based on event and custom tables
- [Editing Custom Workflows](#) on page 1847 for a procedure for editing custom workflows
- [Deleting Custom Workflows](#) on page 1847 for a procedure for deleting custom workflows

Creating Custom Workflows

LICENSE: Any

If the predefined and Sourcefire-provided custom workflows do not meet your needs, you can create custom workflows.

TIP! Instead of creating a new custom workflow, you can export a custom workflow from another appliance and then import it onto your appliance. You can then edit the imported workflow to suit your needs. For more information, see [Importing and Exporting Objects](#) on page 2225.

When you create a custom workflow, you:

- select a table to be the source of the workflow
- provide a workflow name
- add drill-down pages and table view pages to the workflow

For each drill-down page in the workflow, you can:

- provide a name that appears at the top of the page in the web interface
- include up to five columns per page
- specify a default sort order, ascending or descending

You can add table view pages in any position in the sequence of workflow pages. They do not have any editable properties, such as a page name, sort order, or user-definable column positions.

The final page of a custom workflow depends on the table on which you base the workflow, as described in the [Custom Workflow Final Pages](#) table. These final pages are added by default when you create the workflow.

Custom Workflow Final Pages

WORKFLOWS BASED ON...	HAVE THIS FINAL PAGE...
discovery events	hosts
vulnerabilities	vulnerability detail
third-party vulnerabilities	hosts
users	users
intrusion events	packets

The appliance does not add a final page to custom workflows based on other kinds of events (for example, audit log or malware events).

IMPORTANT! The procedure for creating a custom workflow based on connection data is slightly different. For more information, see the next section, [Creating Custom Connection Data Workflows](#).

To create a custom workflow:

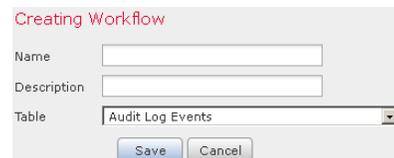
ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Custom > Custom Workflows**.

The Custom Workflows page appears.

2. Click **Create Custom Workflow**.

The Edit Custom Workflow page appears.



The screenshot shows a form titled "Creating Workflow". It contains three input fields: "Name", "Description", and "Table". The "Table" field is a dropdown menu with "Audit Log Events" selected. At the bottom of the form are two buttons: "Save" and "Cancel".

3. Type a name for the workflow in the **Name** field.
You can use up to 60 alphanumeric characters and spaces in the name.
4. Optionally, type a description for the workflow in the **Description** field.
You can use up to 80 alphanumeric characters and spaces.
5. Select the table you want to include from the **Table** drop-down list.

- Optionally, click **Add Page** to add one or more drill-down pages to the workflow.

A drill-down page section appears.

The screenshot shows a configuration window titled 'Page 1'. At the top, there is a 'Page Name' text box containing 'Custom Workflow Page' and a 'Sort Type' dropdown menu set to 'Descending'. Below this is a table configuration section with five columns, labeled 'Column 1' through 'Column 5'. Each column has a 'Sort Priority' dropdown menu and a 'Field' dropdown menu. The 'Sort Priority' dropdowns are currently empty, and the 'Field' dropdowns are also empty.

Begin by typing a name for the page in the **Page Name** field, using up to 80 alphanumeric characters, but no spaces.

Under Column 1, select a sort priority and a table column. This column will appear in the leftmost column of the page. For example, to create a page showing the destination ports that are targeted, and to sort the page by count, select **2** from the **Sort Priority** drop-down list and **DST Port/ICMP Code** from the **Field** drop-down list.

Continue selecting fields to include and setting their sort priority until all the fields to appear on the page have been specified. You can specify up to five fields per page.

IMPORTANT! If you selected **Vulnerabilities** as the Table Type in step 5, then add **IP Address** as a table column, the IP Address column does not appear when you are viewing vulnerabilities using your custom workflow, unless you use the search feature to constrain the workflow to view a specific IP address or range of addresses. For more information on searching for vulnerabilities, see [Searching for Sourcefire Vulnerabilities](#) on page 1439.

- Optionally, click **Add Table View** to add a table view page to the workflow.

IMPORTANT! You must add at least one drill-down page or a table view of events to a custom workflow.

- Click **Save**.

The new workflow is saved and added to the list of custom workflows.

Creating Custom Connection Data Workflows

LICENSE: FireSIGHT

Custom workflows based on connection data are like other custom workflows, except you can include connection data graph pages as well as drill-down pages and table view pages. You can include as many of each type of page in the workflow as you want, in any order. Each connection data graph page contains a single graph, which can be a line graph, bar graph, or pie chart. On line and bar

graphs, you may include more than one dataset. For more information on connection data, including connection summaries, connection graphs, and datasets, see [Understanding Connection Data](#) on page 540.

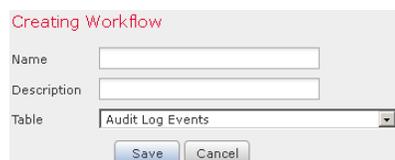
TIP! Instead of creating a new custom workflow, you can export a custom workflow from another appliance and then import it onto your appliance. You can then edit the imported workflow to suit your needs. For more information, see [Importing and Exporting Objects](#) on page 2225.

To create a custom workflow based on connection data:

ACCESS: Admin

1. Select **Analysis > Custom > Custom Workflow**.
2. Click **Create Custom Workflow**.

The Edit Custom Workflow page appears.



Creating Workflow

Name

Description

Table

3. Type a name for the workflow in the **Name** field.
You can use up to 60 alphanumeric characters and spaces.
4. Optionally, type a description for the workflow in the **Description** field.
You can use up to 80 alphanumeric characters and spaces.
5. From the **Table** drop-down list, select **Connection Events**.

6. Optionally, add one or more drill-down pages to the workflow:
 - To add a drill-down page that contains data on individual connections, click **Add Page**.
 - To add a drill-down page that contains connection summary data, click **Add Summary Page**.

In either case, a drill-down page section appears.



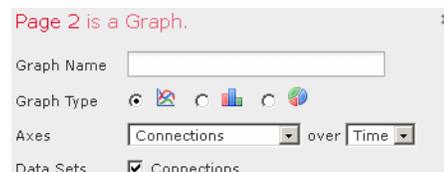
Begin by typing a name for the page in the **Page Name** field using up to 80 alphanumeric characters, but no spaces.

Under **Column 1**, select a sort priority and a table column. This column will appear in the leftmost column of the page.

Continue selecting fields to include and setting their sort priority until all the fields to appear on the page have been specified. You can specify up to five fields per page.

For example, to create a page showing the amount of traffic transmitted over your monitored network and to sort the page by the responders that transmitted the most traffic, select **1** from the **Sort Priority** drop-down list and **Responder Bytes** from the **Field** drop-down list.

7. Optionally, click **Add Graph** to add one or more graph pages to the workflow. A graph section appears.



Begin by typing a name for the page in the **Graph Name** field using up to 80 alphanumeric characters, but no spaces.

Then, select the type of graph you want to include on the page: line graph, bar graph, or pie chart.

Then, specify what kind of data you want to graph by selecting the x- and y-axes of the graph. On a pie chart, the x-axis represents the independent variable and the y-axis represents the dependent variable.

Finally, select the datasets you want to include on the graph. Note that pie charts can only include one data set.

8. Optionally, add a table view of connection data by clicking **Add Table View**.

9. Click **Save**.

The new workflow is saved and added to the list of custom workflows.

Viewing Custom Workflows

LICENSE: Any

The method you use to view a workflow depends on whether the workflow is based on one of the predefined event tables or on a custom table.

If your custom workflow is based on a predefined event table, access it in the same way that you would access a workflow that ships with the appliance. For example, to access a custom workflow based on the Hosts table, select **Analysis Hosts**. If, on the other hand, your custom workflow is based on a custom table, you must access it from the Custom Tables page.

TIP! You can set a custom workflow as the default workflow for any event type; see [Configuring Event View Settings](#) on page 2218.

For more information, see:

- [Viewing Custom Workflows for Predefined Tables](#) on page 1846
- [Viewing Custom Workflows for Custom Tables](#) on page 1846

Viewing Custom Workflows for Predefined Tables

LICENSE: Any

Use the following procedure to view a custom workflow that is **not** based on a custom table. Keep in mind that workflow access depends on your platform and user role, as described in [Selecting Workflows](#) on page 1809.

To view a custom workflow based on a predefined table:

ACCESS: Admin/Any Security Analyst

- ▶ Select the appropriate menu path and option for the table on which you based your custom workflow, as described in the [Features Using Workflows](#) table. The first page of the default workflow for that table appears. To use a different workflow, including a custom workflow, click **(switch workflow)** beside the current workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

Viewing Custom Workflows for Custom Tables

LICENSE: FireSIGHT

Use the following procedure to view a custom workflow that is based on a custom table.

To view a custom workflow based on a custom table:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Custom > Custom Tables**.

The Custom Tables page appears, listing the available custom tables.

2. Click the view icon next to the custom table you want to view, or click the name of the custom table.

The first page of the default workflow for that table appears. To use a different workflow, including a custom workflow, click **(switch workflow)** beside the current workflow title. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see [Setting Event Time Constraints](#) on page 1820.

Editing Custom Workflows

LICENSE: Any

If your event evaluation process changes, you can edit custom workflows to meet your new needs. Note that you cannot edit any of the predefined workflows.

To edit a custom workflow:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Custom > Custom Workflows**.

The Custom Workflows page appears, listing the existing custom workflows.

2. Click the edit icon () next to the name of the workflow that you want to edit.

The Edit Workflow page appears.

3. Make any changes that you want to the workflow and click **Save**.

The changes you made to the workflow are saved.

Deleting Custom Workflows

LICENSE: Any

The following procedure explains how to delete a custom workflow that you no longer need.

To delete a custom workflow:

ACCESS: Admin/Any Security Analyst

1. Select **Analysis > Custom > Custom Workflows**.

The Custom Workflows page appears, listing the available custom workflows.

2. Click the delete icon () next to the name of the workflow that you want to delete.

The workflow is deleted.

CHAPTER 46

MANAGING USERS

If your user account has Administrator access, you can manage the user accounts that can access the web interface on your Defense Center or managed device. On the Defense Center, you can also set up user authentication via an external authentication server, rather than through the internal database.

For more information, see the following sections:

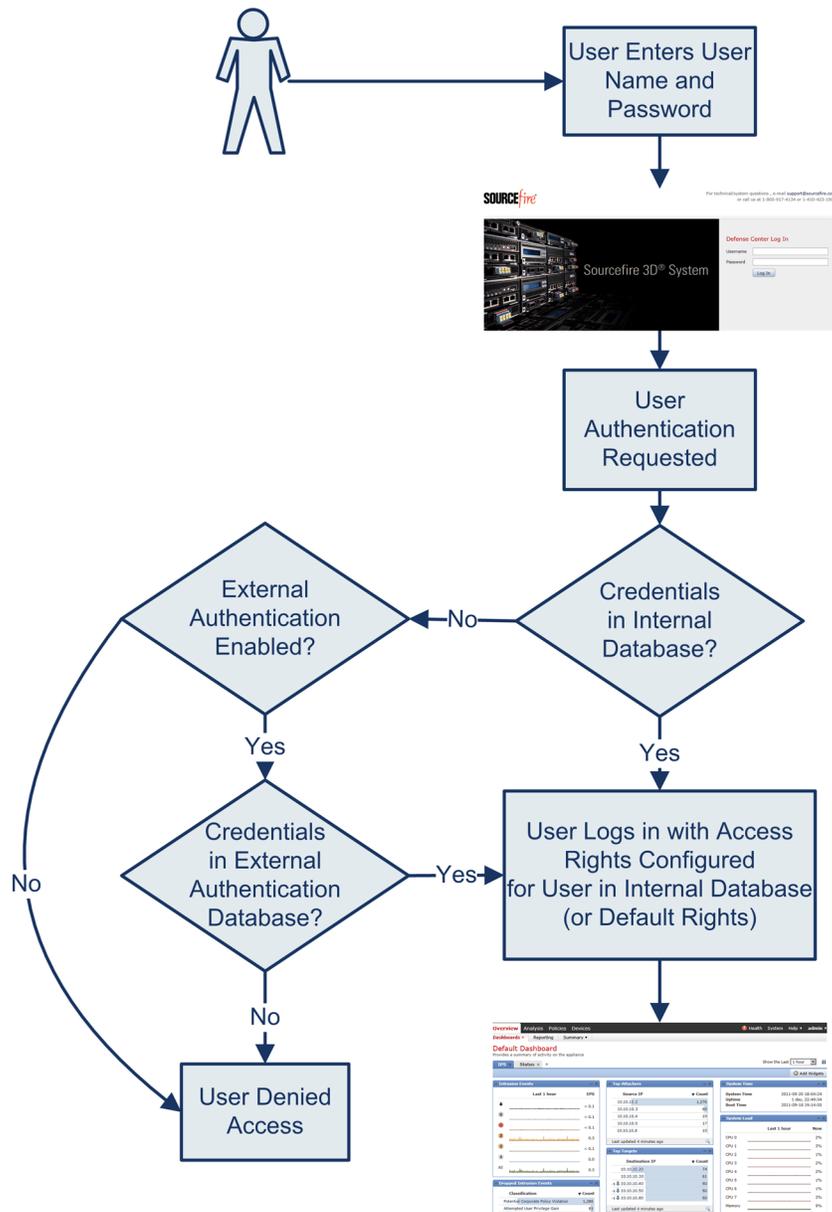
- [Understanding Sourcefire User Authentication](#) on page 1849
- [Managing Authentication Objects](#) on page 1854
- [Managing User Accounts](#) on page 1897
- [Managing User Role Escalation](#) on page 1925

Understanding Sourcefire User Authentication

LICENSE: Any

When a user logs into the web interface, the appliance looks for a match for the user name and password in the local list of users. This process is called *authentication*. There are two kinds of authentication: internal and external. If the user's account uses *internal authentication*, the authentication process checks the local database for this list. If the account uses *external authentication*, the process checks the local database to see if the user exists there and, if the user is not found locally, it queries an external server, such as a Lightweight Directory

Access Protocol (LDAP) directory server or a Remote Authentication Dial In User Service (RADIUS) authentication server, for a list of users.



For users with either internal or external authentication, you can control user permissions. Users with external authentication receive the permissions either for the group or access list they belong to, or based on the default user access role you set in the server authentication object or in a system policy on the managing Defense Center, unless you change the user permissions manually.

For more information, see the following sections:

- [Understanding Internal Authentication](#) on page 1851
- [Understanding External Authentication](#) on page 1851
- [Understanding User Privileges](#) on page 1852

Understanding Internal Authentication

LICENSE: Any

By default, the Sourcefire 3D System uses internal authentication to check user credentials when a user logs in. Internal authentication occurs when the user name and password are verified against records in the internal Sourcefire 3D System database. If you do not enable external authentication when you create a user, the user credentials are managed in the internal database.

Because you manually create each internally authenticated user, you set the access settings when you create the user and you do not need to set default settings.

IMPORTANT! Note that an internally authenticated user is converted to external authentication if you enable external authentication, the same user name exists for the user on the external server, and the user logs in using the password stored for that user on the external server. After an internally authenticated user converts to an externally authenticated user, you cannot revert to internal authentication for that user.

Understanding External Authentication

LICENSE: Any

External authentication occurs when the Defense Center or managed device retrieves user credentials from an external repository, such as an LDAP directory server or RADIUS authentication server. LDAP authentication and RADIUS authentication are types of external authentication. Note that you can only use one form of external authentication for an appliance.

If you want to use external authentication, you must configure an *authentication object* for each external authentication server where you want to request user information. The authentication object contains your settings for connecting to and retrieving user data from that server. You can then enable that object in a system policy on the managing Defense Center and apply the policy to an appliance to enable authentication. When any externally authenticated user logs in, the web interface checks each authentication server to see if that user is listed, in the order the servers are listed in the system policy.

When you create a user, you can specify whether that user is internally or externally authenticated.

TIP! You can use the Import/Export feature to export system policies. When you export a policy with external authentication enabled, the authentication objects are exported with the policy. You can then import the policy and object on another Defense Center. Do **not** import policies with authentication objects onto managed devices.

You can push a system policy to a managed device to enable external authentication on that device, but you cannot control the authentication object from the device's web interface. The only configuration of external authentication on the device occurs when you select the type of authentication for a new user. If you want to disable external authentication on a managed device, disable it in the system policy on the managing Defense Center and reapply the policy to the device. If you apply a local system policy (created on the managed device) to the device itself, external authentication is also disabled.

For more information on specific types of external authentication, see the following sections:

- [Understanding LDAP Authentication](#) on page 1854
- [Understanding RADIUS Authentication](#) on page 1885

Understanding User Privileges

LICENSE: Any

The Sourcefire 3D System lets you allocate user privileges based on the user's role. For example, an analyst typically needs access to event data to analyze the security of monitored networks, but might never require access to administrative functions for the Sourcefire 3D System itself. You can grant analysts predefined roles such as Security Analyst and Discovery Admin and reserve the Administrator role for the network administrator managing the Sourcefire 3D System. You can also create custom user roles with access privileges tailored to your organization's needs.

In the system policy on the Defense Center, you set a default access role for all users who are externally authenticated. After an externally authenticated user logs in for the first time, you can add or remove access rights for that user on the User Management page. If you do not modify the user's rights, the user has only the rights granted by default. Because you create internally authenticated users manually, you set the access rights when you create them.

If you configured management of access rights through LDAP groups, the access rights for users are based on their membership in LDAP groups. They receive the default access rights for the group that they belong to that has the highest level of access. If they do not belong to any groups and you have configured group access, they receive the default user access rights configured in the

authentication object for the LDAP server. If you configure group access, those settings override the default access setting in the system policy.

Similarly, if you assign a user to specific user role lists in a RADIUS authentication object, the user receives all assigned roles, unless one or more of those roles are mutually incompatible. If a user is on the lists for two mutually incompatible roles, the user receives the role that has the highest level of access. If the user does not belong to any lists and you have configured a default access role in the authentication object, the user receives that role. If you configure default access in the authentication object, those settings override the default access setting in the system policy.

The Sourcefire 3D System supports the following predefined user roles, listed in order of precedence, depending on the features you have licensed:

- *Access Admins* can view and modify access control and file policies, but cannot apply their policy changes.
- *Administrators* can set up the appliance's network configuration, manage user accounts and FireAMP cloud connections, and configure system policies and system settings. Users with the Administrator role have all rights and privileges of all other roles (with the exception of lesser, restricted versions of those privileges).
- *Discovery Admins* can review, modify and delete network discovery policies, but cannot apply their policy changes.
- *External Database* users can query the Sourcefire 3D System database using an external application that supports JDBC SSL connections. On the web interface, they can access the online help and user preferences.
- *Intrusion Admins* can review, modify, and delete intrusion policies and intrusion rules.
- *Maintenance Users* can access monitoring functions (including health monitoring, host statistics, performance data, and system logs) and maintenance functions (including task scheduling and backing up the system).

Note that maintenance users do not have access to the functions in the Policies menu and can only access the dashboard from the Analysis menu.

- *Network Admins* can review, modify, and apply device configurations as well as review and modify access control policies (but not file policies).
- *Security Approvers* can view and apply, but not create, configuration and policy changes.
- *Security Analysts* can review, analyze, and delete intrusion, discovery, user activity, connection, correlation, and network change events. They can review, analyze, and (when applicable) delete hosts, host attributes, services, vulnerabilities, and client applications. Security Analysts can also generate reports and view (but not delete or modify) health events.
- *Security Analysts (Read Only)* have all the same rights as Security Analysts, except that they cannot delete events.

In addition to the above predefined roles, you can also configure custom user roles with specialized access privileges. Any role can be the default access role for externally authenticated users.

You can grant user role escalation privileges to externally authenticated user accounts; you can also use an externally authenticated user's password as the escalation password. For more information, see [Managing User Role Escalation](#) on page 1925.

Managing Authentication Objects

LICENSE: Any

Authentication objects are server profiles for external authentication servers, containing connection settings and authentication filter settings for those servers. You can create, manage, and delete authentication objects on the Defense Center. See the following sections for details on these tasks:

- [Understanding LDAP Authentication](#) on page 1854
- [Preparing to Create an LDAP Authentication Object](#) on page 1859
- [Quick Start to LDAP Authentication](#) on page 1860
- [Tuning Your LDAP Authentication Connection](#) on page 1864
- [Creating Advanced LDAP Authentication Objects](#) on page 1866
- [LDAP Authentication Object Examples](#) on page 1880
- [Editing LDAP Authentication Objects](#) on page 1884
- [Creating RADIUS Authentication Objects](#) on page 1885
- [RADIUS Authentication Object Examples](#) on page 1894
- [Editing RADIUS Authentication Objects](#) on page 1896
- [Deleting Authentication Objects](#) on page 1897

Understanding LDAP Authentication

LICENSE: Any

LDAP, or the Lightweight Directory Access Protocol, allows you to set up a directory on your network that organizes objects, such as user credentials, in a centralized location. Multiple applications can then access those credentials and the information used to describe them. If you ever need to change a user's credentials, you can change them in one place, rather than having to change them on each Sourcefire 3D System appliance.

You can create LDAP authentication objects on a Defense Center, but not on other Sourcefire 3D System appliances. However, you can use the external authentication object on any appliance by applying a system policy where the object is enabled to the appliance. When you apply the policy, the object is copied to the appliance.

Note that you can use LDAP naming standards for address specification and for filter and attribute syntax in your authentication object. For more information, see the RFCs listed in the Lightweight Directory Access Protocol (v3): Technical Specification, RFC 3377. Examples of syntax are provided throughout this procedure. Note that when you set up an authentication object to connect to a Microsoft Active Directory Server, you can use the address specification syntax documented in the Internet RFC 822 (Standard for the Format of ARPA Internet Text Messages) specification when referencing a user name that contains a domain. For example, to refer to a user object, you might type `JoeSmith@security.example.com` rather than the equivalent user distinguished name of `cn=JoeSmith,ou=security,dc=example,dc=com` when using Microsoft Active Directory Server.

IMPORTANT! Currently, Sourcefire supports LDAP external authentication on LDAP servers running Microsoft Active Directory on Windows Server 2003 and Windows Server 2008, Oracle Directory Server Enterprise Edition 7.0 on Windows Server 2003 and Windows Server 2008, or OpenLDAP on Linux.

For more information, see the following sections:

- [Setting Defaults](#) on page 1855
- [Setting a Base DN](#) on page 1856
- [Setting a Base Filter](#) on page 1856
- [Selecting an Impersonation Account](#) on page 1856
- [Encrypting Your LDAP Connection](#) on page 1856
- [Setting the User Name Template](#) on page 1857
- [Setting a Connection Timeout](#) on page 1857
- [Using Attributes to Manage Access](#) on page 1857
- [Using Group Membership to Manage Access](#) on page 1858
- [Setting up Shell Access](#) on page 1858
- [Testing the Connection](#) on page 1859

Setting Defaults

LICENSE: Any

You can populate several fields using default values based on the server type you plan to connect to. Default values propagate the User Name Template, UI Access Attribute, Shell Access Attribute, Group Member Attribute, and Group Member URL Attribute fields when you select a server type and set defaults.

Setting a Base DN

LICENSE: Any

When the local appliance searches the LDAP server to retrieve user information on the authentication server, it needs a starting point for that search. You can specify the tree that the local appliance should search by providing a base distinguished name, or *base DN*.

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`.

After you identify a primary server, you can automatically retrieve a list of available base DNs from it and select the appropriate base DN.

Setting a Base Filter

LICENSE: Any

You can add a *base filter* that sets a specific value for a specific attribute. The base filter focuses your search by only retrieving objects in the base DN that have the attribute value set in the filter. Enclose the base filter in parentheses. For example, to filter for only users with a common name starting with F, use the filter `(cn=F*)`.

To test your base filter more specifically by entering a test user name and password, see [Testing User Authentication](#) on page 1879.

Selecting an Impersonation Account

LICENSE: Any

To allow the local appliance to access the user objects, you must supply user credentials for an impersonation account. The *impersonation account* is a user account with appropriate rights to browse the directory named by the base DN and retrieve the user objects you want to retrieve. Remember that the distinguished name for the user you specify must be unique to the tree for the server.

Encrypting Your LDAP Connection

LICENSE: Any

You can manage the encryption method for your LDAP connection. You can choose no encryption, Transport Layer Security (TLS), or Secure Sockets Layer (SSL) encryption.

Note that if you are using a certificate to authenticate when connecting via TLS or SSL, the name of the LDAP server in the certificate **must** match the name that you use in the Host Name/IP Address field. For example, if you enter `10.10.10.250` in the authentication profile and `computer1.example.com` in the certificate, the connection fails. Changing the name of the server in the authentication profile to `computer1.example.com` causes the connection to succeed.

Note that if you change the encryption method after specifying the port, the port resets to the default value for the selected server type.

Setting the User Name Template

LICENSE: Any

Selecting a user name template lets you indicate how user names entered on login should be formatted, by mapping the string conversion character (`%s`) to the value of the shell access attribute for the user. The user name template is the format for the distinguished name used for authentication. When a user enters a user name into the login page, the name is substituted for the string conversion character and the resulting distinguished name is used to search for the user credentials.

For example, to set a user name template for the Security organization of the Example company, you might enter `%s@security.example.com`.

Setting a Connection Timeout

LICENSE: Any

If you specify a backup authentication server, you can set a timeout for the connection attempt to the primary server. If the timeout period elapses without a response from the primary authentication server, the appliance then queries the backup server. For example, if the primary server has LDAP disabled, the appliance queries the backup server.

If LDAP is running on the port of the primary LDAP server and for some reason refuses to service the request (due to misconfiguration or other issues), however, the failover to the backup server does not occur.

Using Attributes to Manage Access

LICENSE: Any

Different types of LDAP servers use different attributes to store user data. If your LDAP server uses a UI access attribute of `uid`, the local appliance checks the `uid` attribute value for each object in the tree indicated by the base DN you set. If you do not set a specific UI access attribute, the local appliance checks the distinguished name for each user record on the LDAP server to see if it matches the user name. If one of the objects has a matching user name and password, the user login request is authenticated.

You can substitute a different LDAP attribute to make the local appliance match a user name with that attribute rather than the value of the distinguished name. Selecting a server type and setting defaults fills in a UI access attribute appropriate for that type of server. If one of the objects has a matching user name and password as a value for the attribute you specify, the user login request is authenticated. You can use any attribute, if the value of the attribute is a valid user name for the Sourcefire 3D System web interface. Valid user names are unique, and can include underscores (`_`), periods (`.`), hyphens (`-`), and alphanumeric characters.

The shell access attribute of your LDAP server acts as a shell access attribute. If your LDAP server uses `uid`, the local appliance checks the user name entered on login against the attribute value of `uid`. You can also set a custom shell access attribute other than `uid`.

Note that selecting a server type and setting defaults prepopulates a shell access attribute typically appropriate for that type of server. You can use any attribute, if the value of the attribute is a valid user name for shell access. Valid user names are unique, and can include underscores (`_`), periods (`.`), hyphens (`-`), and alphanumeric characters.

Using Group Membership to Manage Access

LICENSE: Any

If you prefer to base default access settings on a user's membership in an LDAP group, you can specify distinguished names for existing groups on your LDAP server for each of the access roles used by your Sourcefire 3D System. When you do so, you can configure a default access setting for those users detected by LDAP that do not belong to any specified groups. When a user logs in, the Sourcefire 3D System dynamically checks the LDAP server and assigns default access rights according to the user's current group membership.

When a user authenticated by an LDAP server logs into a local Sourcefire 3D System appliance for the first time, the user receives the default access settings for groups the user belongs to, or if groups are not configured, the default access setting you selected in the system policy.

You can then modify those settings, unless the settings are granted through group membership.

Setting up Shell Access

LICENSE: Any

You can use the LDAP server to authenticate accounts for shell access on a managed device or Defense Center. Specify a search filter that retrieves entries for users to whom you want to grant shell access. Note that you can only configure shell access for the first authentication object in your system policy. For more information on managing authentication object order, see [Configuring Authentication Profiles](#) on page 1975.

With the exception of the admin account, shell access is controlled entirely through the shell access attribute you set. Shell users are configured as local users on the appliance. The filter you set here determines which set of users on the LDAP server can log into the shell.

Note that a home directory for each shell user is created on login, and when an LDAP shell access user account is disabled (by disabling the LDAP connection), the directory remains, but the user shell is set to `/bin/false` in `/etc/passwd` to disable the shell. If the user then is re-enabled, the shell is reset, using the same home directory.

If all users qualified in the base DN are also qualified for shell access privileges, you can configure the shell access filter to search more efficiently by making the shell access filter the same as the base filter. Normally, the LDAP query to retrieve users combines the base filter with the shell access filter. If the shell access filter was the same as the base filter, the same query runs twice, which is unnecessarily time-consuming.

Shell users should log in using user names with all lowercase letters.

WARNING! On Series 3 Defense Centers, all shell users have **sudoers** privileges. Make sure that you restrict the list of users with shell access appropriately. On Series 3 and virtual devices, shell access granted to externally authenticated users defaults to the **Configuration** level of command line access, which also grants **sudoers** privileges.

Testing the Connection

LICENSE: Any

After you configure LDAP server and authentication settings, you can specify user credentials for a user who should be able to authenticate to test those settings.

For the user name, you can enter the value for the **uid** attribute for the user you want to test with. If you are connecting to a Microsoft Active Directory Server and supply a UI access attribute in place of **uid**, use the value for that attribute as the user name.

Preparing to Create an LDAP Authentication Object

LICENSE: Any

Before you configure a connection to your LDAP server, you should collect the information that you need to create the LDAP authentication object. For more information on specific aspects of configuration, see [Understanding LDAP Authentication](#) on page 1854.

You need the following for any authentication object:

- the server name or IP address for the server where you plan to connect
- the server type of the server where you plan to connect
- the user name and password for a user account with sufficient privileges to browse the LDAP tree; Sourcefire recommends that you use a domain admin user account for this purpose
- if there is a firewall between the appliance and the LDAP server, an entry in the firewall to allow outgoing connections
- if possible, the base distinguished name for the server directory where the user names reside

Note that you can use a third-party LDAP client to browse the LDAP tree and see base DN and attribute descriptions. You can also use that client to confirm that

your selected user can browse the base DN you select. Ask your LDAP administrator to recommend an approved LDAP client for your LDAP server.

Depending on how you plan to customize your LDAP authentication object configuration, you might also need the information in the following table.

Additional LDAP Configuration Information

To...	YOU NEED...
connect over a port other than 389	the port number
connect via an encrypted connection	the certificate for the connection
filter the users who can access your appliance based on an attribute value	the attribute-value pair to filter by
use an attribute as a UI access attribute rather than checking the user distinguished name	the name of the attribute
use an attribute as a shell login attribute rather than checking the user distinguished name	the name of the attribute
filter the users who can access your appliance via the shell based on an attribute value	the attribute-value pair to filter by
associate groups with specific user roles	the distinguished name of each group, as well as the group member attribute if the groups are static groups or the group member URL attribute if the groups are dynamic groups

Quick Start to LDAP Authentication

LICENSE: Any

You can set up an LDAP authentication object where you customize many of the values. However, if you just want to authenticate all the users in a particular directory, you can create an authentication object with the base DN for that directory. If you set defaults to those for your server type and supply authentication credentials for the account used to retrieve user data from the

server, you can quickly create an authentication object. Follow the procedure below to do so.

IMPORTANT! If you prefer to consider and possibly customize each authentication setting when creating the authentication object, use the procedure in [Creating Advanced LDAP Authentication Objects](#) on page 1866 to create the object. If you plan to encrypt your connection to the server, set user timeouts, customize the user name template, or assign Sourcefire 3D System user roles based on LDAP group membership, use the advanced procedure.

Before you configure a connection to your LDAP server, you should collect the information that you need to create the LDAP authentication object. For more information on specific aspects of configuration, see [Understanding LDAP Authentication](#) on page 1854.

You need the following:

- the server name or IP address for the server where you plan to connect
- the server type of the server where you plan to connect
- the user name and password for a user account with sufficient privileges to browse the LDAP tree; Sourcefire recommends that you use a domain admin user account for this purpose

Optionally, if you want to constrain your user search further, you can add a base filter to set a specific value for a specific attribute. The base filter focuses your search by only retrieving objects in the base DN that have the attribute value set in the filter. Enclose the base filter in parentheses. For example, to filter for only users with a common name starting with F, use the filter `(cn=F*)`. When you save the authentication object, the local appliance queries using the base filter to test it and indicates whether or not the filter appears to be correct.

To create an LDAP authentication object:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears
2. Click the **Login Authentication** tab.
The Login Authentication page appears.
3. Click **Create Authentication Object**.

4. Select **LDAP** from the **Authentication Method** drop-down list.
LDAP configuration options appear.

The screenshot shows a configuration form titled "Authentication Object". It contains the following fields and controls:

- Authentication Method:** A dropdown menu with "LDAP" selected.
- Name *:** An empty text input field.
- Description:** An empty text input field.
- Server Type:** A dropdown menu with "MS Active Directory" selected, and a "Set Defaults" button to its right.
- Primary Server:** A section header.
- Host Name/IP Address *:** An empty text input field.
- Port *:** A text input field containing "389".
- Backup Server (Optional):** A section header.
- Host Name/IP Address:** An empty text input field.
- Port:** A text input field containing "389".

5. Type a name and description for the authentication server in the **Name** and **Description** fields.
6. Select your server type from the **Server Type** drop-down list, then click the **Set Defaults** button to configure default settings for that type. You have the following options:
 - If you are connecting to a Microsoft Active Directory Server, select **MS Active Directory**, then click **Set Defaults**.
 - If you are connecting to a Sun Java Systems Directory Server or Oracle Directory Server, select **Oracle Directory**, then click **Set Defaults**.
 - If you are connecting to an OpenLDAP server, select **OpenLDAP**, then click **Set Defaults**.
 - If you are connecting to a server other than those listed above and want to clear default settings, select **Other**, then click **Set Defaults**.
7. Type the IP address or host name for the primary server where you want to obtain authentication data in the **Primary Server Host Name/IP Address** field.

IMPORTANT! If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

8. To fetch a list of all base DN's, click **Fetch DN's** and select the appropriate base DN from the drop-down list.

For example, to authenticate names in the Security organization at the Example company, select `ou=security,dc=example,dc=com`.

9. Optionally, to set a filter that retrieves only specific objects within the directory you specified as the Base DN, type the attribute type, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses, in the **Base Filter** field.

For example, if the user objects in a tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, type `(physicalDeliveryOfficeName=NewYork)`.

10. In the **User Name** and **Password** fields, type the distinguished name and password for a user who has sufficient credentials to browse the LDAP server.

For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute and the object for the administrator in the Security division at our example company has a `uid` value of `NetworkAdmin`, you might type `uid=NetworkAdmin,ou=security,dc=example,dc=com`.

WARNING! If you are connecting to a Microsoft Active Directory Server, you cannot provide a server user name that ends with the `$` character.

11. Retype the password in the **Confirm Password** field.
12. Optionally, to retrieve users for shell access, type the attribute type you want to filter on in the **Shell Access Attribute** field.

For example, on a Microsoft Active Directory Server, use the `SAMAccountName` shell access attribute to retrieve shell access users by typing `SAMAccountName` in the **Shell Access Attribute** field.

IMPORTANT! IPv6 addresses are not supported for shell authentication.

13. In the **User Name** and **Password** fields, type the `uid` value or shell access attribute value and password for the user whose credentials should be used to validate access to the LDAP server. Note, again, that server user names associated with a Microsoft Active Directory Server cannot end with the character `$`.

For example, to test to see if you can retrieve the `JSmith` user credentials at the Example company, type `JSmith`.

14. Click **Test** to test the connection.

A message appears, either indicating success of the test or detailing what settings are missing or need to be corrected. If the test succeeds, the test output appears at the bottom of the page, including a list of the users retrieved by the connection. If the number of users that appear in the test output is limited by the number of user records your LDAP server returns, the test output indicates this limitation.

15. You have two options:

- If the test succeeds, click **Save**.

The Login Authentication page appears, with the new object listed.

To enable LDAP authentication using the object on an appliance, you must apply a system policy with that object enabled to the appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

- If the test fails, or if you want to refine the list of users retrieved, continue with the next section, [Tuning Your LDAP Authentication Connection](#).

Tuning Your LDAP Authentication Connection

LICENSE: Any

If you create an LDAP authentication object and it either does not succeed in connecting to the server you select, or does not retrieve the list of users you want, you can tune the settings in the object.

If the connection fails when you test it, try the following suggestions to troubleshoot your configuration:

- Use the messages displayed at the top of the screen and in the test output to determine which areas of the object are causing the issue.
- Check that the user name and password you used for the object are valid:
 - Check that the user has the rights to browse to the directory indicated in your base distinguished name by connecting to the LDAP server using a third-party LDAP browser.
 - Check that the user name is unique to the directory information tree for the LDAP server.
 - Check that the user name contains only underscores, periods, hyphens, and alphanumeric characters.
 - If you see an LDAP bind error 49 in the test output, the user binding for the user failed. Try authenticating to the server through a third-party application to see if the binding fails through that connection as well.
- Check that you have correctly identified the server:
 - Check that the server IP address or host name is correct.
 - Check that you have TCP/IP access from your local appliance to the authentication server where you want to connect.
 - Check that access to the server is not blocked by a firewall and that the port you have configured in the object is open.
 - If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used for the server.

- Check that you have not used an IPv6 address for the server connection if you are authenticating shell access.
- If you used server type defaults, check that you have the correct server type and click **Set Defaults** again to reset the default values.

For more information, see [Identifying the LDAP Authentication Server](#) on page 1867.

- If you typed in your base distinguished name, click **Fetch DNs** to retrieve all the available base distinguished names on the server, and select the name from the list.
- If you are using any filters, access attributes, or advanced settings, check that each is valid and typed correctly.
- If you are using any filters, access attributes, or advanced settings, try removing each setting and testing the object without it.
- If you are using a base filter or a shell access filter, make sure that the filter is enclosed in parentheses and that you are using a valid comparison operator. For more information, see [Setting a Base Filter](#) on page 1856 and [Setting up Shell Access](#) on page 1858.
- To test a more restricted base filter, try setting it to the base distinguished name for the user to retrieve just that user.
- If you are using an encrypted connection:
 - Check that the name of the LDAP server in the certificate matches the host name that you use to connect.
 - Check that you have not used an IPv6 address with an encrypted server connection.
- If you are using a test user, make sure that the user name and password are typed correctly.
- If you are using a test user, remove the user credentials and test the object.
- Test the query you are using by connecting to the LDAP server via the command line on the appliance you want to connect from using this syntax:

```
ldapsearch -x -b 'base_distinguished_name'  
-h LDAPserver_ip_address -p port -v -D  
'user_distinguished_name' -w 'base_filter'
```

For example, if you are trying to connect to the security domain on `myrtle.example.com` using the `domainadmin@myrtle.example.com` user and a base filter of `(cn=*)`, you could test the connection using this statement:

```
ldapsearch -x -b 'CN=security,DC=myrtle,DC=example,DC=com'  
-h myrtle.example.com -p 389 -v -D  
'domainadmin@myrtle.example.com' -w '(cn=*)'
```

If you can test your connection successfully but authentication does not work after you apply a system policy, check that authentication and the object you want to use are both enabled in the system policy that is applied to the appliance.

If you connect successfully but want to adjust the list of users retrieved by your connection, you can add or change a base filter or shell access filter or use a more restrictive or less restrictive base DN. For more information, see the following topics:

- [Setting a Base DN](#) on page 1856
- [Setting a Base Filter](#) on page 1856
- [Configuring LDAP-Specific Parameters](#) on page 1869

Creating Advanced LDAP Authentication Objects

LICENSE: Any

You can create LDAP authentication objects to provide user authentication services for an appliance.

When you create an authentication object, you define settings that let you connect to an authentication server. You also select the directory context and search criteria you want to use to retrieve user data from the server. Optionally, you can configure shell access authentication.

Make sure you have TCP/IP access from your local appliance to the authentication server where you want to connect.

Although you can use the default settings for your server type to quickly set up a basic LDAP configuration, you can also customize advanced settings to control whether the appliance makes an encrypted connection to the LDAP server, the timeout for the connection, and which attributes the server checks for user information.

For the LDAP-specific parameters, you can use LDAP naming standards and filter and attribute syntax. For more information, see the RFCs listed in the Lightweight Directory Access Protocol (v3): Technical Specification, RFC 3377. Examples of syntax are provided throughout this procedure. Note that when you set up an authentication object to connect to a Microsoft Active Directory Server, you can use the address specification syntax documented in the Internet RFC 822 (Standard for the Format of ARPA Internet Text Messages) specification when referencing a user name that contains a domain. For example, to refer to a user object, you might type `JoeSmith@security.example.com` rather than the equivalent user distinguished name of `cn=JoeSmith,ou=security,dc=example,dc=com` when using Microsoft Active Directory Server.

To create an advanced authentication object:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears
2. Click the **Login Authentication** tab.
The Login Authentication page appears.

3. Click **Create Authentication Object**.

The Create Authentication Object page appears.

The screenshot shows a web-based configuration form for an authentication object. The form is titled "Authentication Object" and is divided into several sections. The "Authentication Method" is set to "LDAP". The "Name" field is required and is empty. The "Description" field is also empty. The "Server Type" is set to "MS Active Directory", and there is a "Set Defaults" button next to it. The "Primary Server" section includes "Host Name/IP Address *" and "Port *" fields. The "Host Name/IP Address" field is empty, and the "Port" field is set to "389". The "Additional Test Parameters" section includes "User Name" and "Password" fields, both of which are empty. At the bottom of the form, there are "Save", "Test", and "Cancel" buttons. A legend indicates that an asterisk (*) denotes a required field.

4. Identify the authentication server where you want to retrieve user data for external authentication. For more information, see [Identifying the LDAP Authentication Server](#) on page 1867.
5. Configure authentication settings to build a search request that retrieves the users you want to authenticate. Specify a user name template to format the user names that users enter on login. For more information, see [Configuring LDAP-Specific Parameters](#) on page 1869.
6. Optionally, configure LDAP groups to use as the basis for default access role assignments. For more information, see [Configuring Access Settings by Group](#) on page 1875.
7. Optionally, configure authentication settings for shell access. For more information, see [Configuring Administrative Shell Access](#) on page 1878.
8. Test your configuration by entering the name and password for a user who can successfully authenticate. For more information, see [Testing User Authentication](#) on page 1879.

Your changes are saved. Remember that you have to apply a system policy with the object enabled to an appliance before the authentication changes take place on that appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

Identifying the LDAP Authentication Server

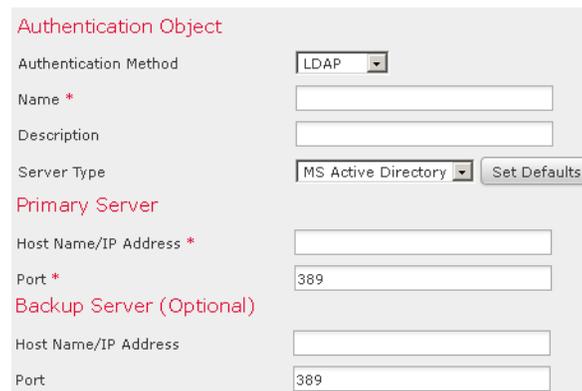
LICENSE: Any

When you create an authentication object, you first specify the primary and backup server and server port where you want the managed device or Defense Center to connect for authentication.

To identify an LDAP authentication server:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears
2. Click the **Login Authentication** tab.
The Login Authentication page appears.
3. Click **Create Authentication Object**.
The Create Authentication Object page appears.
4. Select **LDAP** from the **Authentication Method** drop-down list.
LDAP configuration options appear.



5. Type a name and description for the authentication server in the **Name** and **Description** fields.
6. Optionally, in the **Server Type** field, select the type of LDAP server you plan to connect to and click **Set Defaults** to populate the User Name Template, UI Access Attribute, Shell Access Attribute, Group Member Attribute, and Group Member URL Attribute fields with default values. You have the following options:
 - If you are connecting to a Microsoft Active Directory server, select **MS Active Directory** and click **Set Defaults**.
 - If you are connecting to a Sun Java Systems Directory Server or Oracle Directory Server, select **Oracle Directory** and click **Set Defaults**.
 - If you are connecting to an OpenLDAP server, select **OpenLDAP** and click **Set Defaults**.
 - If you are connecting to a LDAP server other than those listed above and want to clear default settings, select **Other** and click **Set Defaults**.

7. Type the IP address or host name for the primary server where you want to obtain authentication data in the **Primary Server Host Name/IP Address** field.

IMPORTANT! If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

8. Optionally, modify the port used by the primary authentication server in the **Primary Server Port** field.
9. Optionally, type the IP address or host name for the backup server where you want to obtain authentication data in the **Backup Server Host Name/IP Address** field.
10. Optionally, modify the port used by the primary authentication server in the **Backup Server Port** field.

Continue with [Configuring LDAP-Specific Parameters](#).

Configuring LDAP-Specific Parameters

LICENSE: Any

The settings in the LDAP-specific parameters section determine the area of the LDAP directory where the appliance searches for user names, and control details of how the appliance connects to the LDAP server.

When configuring these settings, note that valid user names are unique, and can include underscores (_), periods (.), and hyphens (-), but otherwise only alphanumeric characters are supported.

In addition for most LDAP-specific settings, you can use LDAP naming standards and filter and attribute syntax. For more information, see the RFCs listed in the Lightweight Directory Access Protocol (v3): Technical Specification, RFC 3377. Examples of syntax are provided throughout this procedure. Note that when you set up an authentication object to connect to a Microsoft Active Directory Server, you can use the address specification syntax documented in the Internet RFC 822 (Standard for the Format of ARPA Internet Text Messages) specification when referencing a user name that contains a domain. For example, to refer to a user object, you might type `JoeSmith@security.example.com` rather than the equivalent user distinguished name of `cn=JoeSmith,ou=security,dc=example,dc=com` when using Microsoft Active Directory Server.

The following table describes each of the LDAP-specific parameters.

LDAP-Specific Parameters

SETTING	DESCRIPTION	EXAMPLE
Base DN	<p>Supplies the base distinguished name of the directory where the appliance searches for user information on the LDAP server.</p> <p>Typically, the base DN has a basic structure indicating the company domain and operational unit.</p> <p>Note that after you identify a primary server, you can automatically retrieve a list of available base DN's from the server and select the appropriate base DN.</p>	<p>The Security organization of the Example company might have a base DN of <code>ou=security,dc=example,dc=com</code></p>
Base Filter	<p>Focuses your search by only retrieving objects in the base DN that have the specific attribute-value pair set in the filter. Note that you must enclose the base filter in parentheses.</p> <p>To test your base filter more specifically by entering a test user name and password, see Testing User Authentication on page 1879.</p>	<p>To filter for only users with a common name starting with F, use the filter <code>(cn=F*)</code>.</p>
User Name/ Password	<p>Allow the local appliance to access the user objects. Supply user credentials for a user with appropriate rights to the authentication objects you want to retrieve. The distinguished name for the user you specify must be unique to the directory information tree for the LDAP server. Note that server user names associated with a Microsoft Active Directory Server cannot end with the \$ character.</p>	<p>The user name for the <code>admin</code> user in the Security organization of the Example company might have a user name of <code>cn=admin,ou=security,dc=example,dc=com</code></p>
Encryption	<p>Determines whether and how the communications are encrypted. You can choose no encryption, Transport Layer Security (TLS), or Secure Sockets Layer (SSL) encryption. Note that if you are using a certificate to authenticate when connecting via TLS or SSL, the name of the LDAP server in the certificate must match the name that you use to connect.</p> <p>If you change the encryption method after specifying the port, the port resets to the default value for the selected server type.</p>	<p>If you enter <code>10.10.10.250</code> in the authentication profile and <code>computer1.example.com</code> in the certificate, the connection fails, even if <code>computer1.example.com</code> has an IP address of <code>10.10.10.250</code>. Changing the name of the server in the authentication profile to <code>computer1.example.com</code> causes the connection to succeed.</p>

LDAP-Specific Parameters (Continued)

SETTING	DESCRIPTION	EXAMPLE
SSL Certificate Upload Path	Indicates the path on your local computer to the certificate to be used for encryption.	c:/server.crt
User Name Template	Indicates how user names entered on login should be formatted, by mapping the string conversion character (%s) to the value of the shell access attribute for the user. The user name template is the format for the distinguished name used for authentication. When a user enters a user name into the login page, the appliance substitutes the name for the string conversion character and uses the resulting distinguished name to search for the user credentials.	To set a user name template for the Security organization of the Example company, enter %s@security.example.com.

LDAP-Specific Parameters (Continued)

SETTING	DESCRIPTION	EXAMPLE
Timeout	<p>Sets a timeout for the connection attempt to the primary server, so the connection rolls over to the backup server. If the number of seconds indicated in this field (or the timeout on the LDAP server) elapses without a response from the primary authentication server, the appliance then queries the backup server.</p> <p>However, if LDAP is running on the port of the primary LDAP server and for some reason refuses to service the request, the failover to the backup server does not occur.</p>	<p>If the primary server has LDAP disabled, the appliance queries the backup server.</p>
UI Access Attribute	<p>Tells the local appliance to match the value of a specific attribute rather than the value of the user distinguished name. You can use any attribute, if the value of the attribute is a valid user name for the Sourcefire 3D System web interface. If one of the objects has a matching user name and password, the user login request is authenticated.</p> <p>Selecting a server type and setting defaults prepopulates the UI Access Attribute with a value typically appropriate for that type of server.</p> <p>If you leave this field blank, the local appliance checks the user distinguished name value for each user record on the LDAP server to see if it matches the user name.</p>	<p>SAMAccountName</p>
Shell Access Attribute	<p>If you want to check a specific attribute for shell access credentials, you must explicitly set this field to match the attribute. You can use any attribute if the value of the attribute is a valid user name for shell access.</p> <p>If you leave this field blank, the user distinguished name is used for shell access authentication.</p> <p>Note that selecting a server type and setting defaults prepopulates this field with an attribute typically appropriate for that type of server.</p>	<p>SAMAccountName</p>

To configure the LDAP-specific parameters for a server:

ACCESS: Admin

1. In the LDAP-Specific Parameters section of the Create Authentication Object page, you have two options for setting the base DN:
 - To fetch a list of all available domains, click **Fetch DNs** and select the appropriate base domain name from the drop-down list.
 - Type the base distinguished name for the LDAP directory you want to access in the **Base DN** field.

For example, to authenticate names in the Security organization at the Example company, type or select `ou=security,dc=example,dc=com`.

The screenshot shows the 'LDAP-Specific Parameters' configuration form. It includes the following fields and controls:

- Base DN ***: A text input field with a 'Fetch DNs' button to its right.
- Base Filter**: A text input field.
- User Name ***: A text input field.
- Password ***: A text input field.
- Confirm Password ***: A text input field.
- Show Advanced Options**: A toggle switch currently set to 'off' (indicated by a right-pointing triangle).
- Attribute Mapping**: A section header.
- UI Access Attribute ***: A text input field with a 'Fetch Attrs' button to its right.
- Shell Access Attribute ***: A text input field.

2. Optionally, to set a filter that retrieves only specific objects within the directory you specified as the Base DN, type the attribute type, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses, in the **Base Filter** field.

For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, type `(physicalDeliveryOfficeName=NewYork)`.

3. Type the distinguished name and password for the user whose credentials should be used to validate access to the LDAP directory in the **User Name** and **Password** fields.

For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute and the object for the administrator in the Security division at our example company has a `uid` value of `NetworkAdmin`, you might type `uid=NetworkAdmin,ou=security,dc=example,dc=com`.

WARNING! If you are connecting to a Microsoft Active Directory Server, you cannot provide a server user name that ends with the \$ character.

4. Retype the password in the **Confirm Password** field.
5. After you configure the basic LDAP-specific parameters, you have several options:
 - To access advanced options, click the arrow next to **Show Advanced Options** and continue with the next step.
 - If you want to configure user default roles based on LDAP group membership, continue with [Configuring Access Settings by Group](#) on page 1875.
 - If you are not using LDAP groups for authentication, continue with [Configuring Administrative Shell Access](#) on page 1878.
6. Optionally, select one of the following encryption modes:
 - To connect using Secure Sockets Layer (SSL), select **SSL**.
 - To connect using Transport Layer Security (TLS), select **TLS**.
 - To connect without encryption, select **None**.

IMPORTANT! Note that if you change the encryption method after specifying a port, you reset the port to the default value for that method. For none or TLS, the port uses the default value of 389. If you select SSL encryption, the port uses the default of 636.

7. If you selected TLS or SSL encryption and you want to use a certificate to authenticate, click **Browse** to browse to the location of a valid TLS or SSL certificate or, in the **SSL Certificate Upload Path** field, type the path to the certificate.

A message appears, indicating a successful certificate upload.

IMPORTANT! If you previously uploaded a certificate and want to replace it, upload the new certificate and reapply the system policy to your appliances to copy over the new certificate.

8. Optionally, in the **User Name Template** field, type the string conversion character (%s) used to determine the user name from the value found in the **UI Access Attribute**.

For example, to authenticate all users who work in the Security organization of our example company by connecting to an OpenLDAP server where the shell access attribute is `uid`, you might type `uid=%s,ou=security,dc=example,dc=com` in the **User Name Template** field. For a Microsoft Active Directory server, you could type `%s@security.example.com`.

9. Optionally, in the **Timeout** field, type the number of seconds that should elapse before rolling over to the backup connection.

10. Optionally, to retrieve users based on an attribute instead of the Base DN and Base Filter, you have two options:

- Click **Fetch Attrs** to retrieve a list of available attributes and select the appropriate attribute.
- Type the attribute in the **UI Access Attribute** field.

For example, on a Microsoft Active Directory Server, you may want to use the UI Access Attribute to retrieve users, because there may not be a `uid` attribute on Active Directory Server user objects. Instead, you can search the `userPrincipalName` attribute by typing `userPrincipalName` in the **UI Access Attribute** field.

11. Optionally, to retrieve users for shell access, type the attribute to filter by in the **Shell Access Attribute** field.

For example, on a Microsoft Active Directory Server, use the `sAMAccountName` shell access attribute to retrieve shell access users by typing `sAMAccountName` in the **Shell Access Attribute** field.

12. For the next step, you have two choices:

- If you want to configure user default roles based on LDAP group membership, continue with [Configuring Access Settings by Group](#).
- If you are not using LDAP groups for authentication, continue with [Configuring Administrative Shell Access](#) on page 1878.

Configuring Access Settings by Group

LICENSE: Any

If you prefer to base default access settings on a user's membership in an LDAP group, you can specify distinguished names for existing groups on your LDAP server for each of the access roles used by your Sourcefire 3D System. When you do so, you can configure a default access setting for those users detected by LDAP that do not belong to any specified groups. When a user logs in, the Sourcefire 3D System dynamically checks the LDAP server and assigns default access rights according to the user's current group membership.

Any group you reference must exist on the LDAP server. You can reference static LDAP groups or dynamic LDAP groups. Static LDAP groups are groups where membership is determined by group object attributes that point to specific users, and dynamic LDAP groups are groups where membership is determined by creating an LDAP search that retrieves group users based on user object attributes. Group access settings for a role only affect users who are members of the group.

The access rights granted when a user logs into the Sourcefire 3D System depend on the LDAP configuration:

- If no group access settings are configured for your LDAP server, when a new user logs in, the Sourcefire 3D System authenticates the user against the LDAP server and then grants user rights based on the default minimum access role set in the system policy.
- If you configure any group settings, new users belonging to specified groups inherit the minimum access setting for the groups where they are members.
- If a new user does not belong to any specified groups, the user is assigned the default minimum access role specified in the Group Controlled Access Roles section of the authentication object.
- If a user belongs to more than one configured group, the user receives the access role for the group with the highest access as a minimum access role.

You cannot use the Sourcefire 3D System user management page to remove the minimum access rights for users assigned an access role because of LDAP group membership. You can, however, assign additional rights. When you modify the access rights for an externally authenticated user, the Authentication Method column on the User Management page provides a status of **External - Locally Modified**.

IMPORTANT! If you use a dynamic group, the LDAP query is used exactly as it is configured on the LDAP server. For this reason, the Sourcefire 3D System limits the number of recursions of a search to four to prevent search syntax errors from causing infinite loops. If a user's group membership is not established in those recursions, the default access role defined in the Group Controlled Access Roles section is granted to the user.

To configure default roles based on group membership:

ACCESS: Admin

1. On the Create Authentication Object page, click the down arrow next to **Group Controlled Access Roles**.

The section expands.

Group Controlled Access Roles (Optional) ▾

Access Admin	<input type="text"/>
Administrator	<input type="text"/>
External Database User	<input type="text"/>
Intrusion Admin	<input type="text"/>
Maintenance User	<input type="text"/>
Network Admin	<input type="text"/>
Discovery Admin	<input type="text"/>
Security Approver	<input type="text"/>
Security Analyst	<input type="text"/>
Security Analyst (Read Only)	<input type="text"/>
Default User Role	<input type="text" value="Access Admin"/>
Group Member Attribute	<input type="text"/>
Group Member URL Attribute	<input type="text"/>

2. Optionally, configure access defaults by group membership.

In the **DN** fields that correspond to Sourcefire 3D System user roles, type the distinguished name for the LDAP groups that contain users who should be assigned to those roles.

For example, you might type the following in the **Administrator** field to authenticate names in the information technology organization at the Example company:

```
cn=itgroup,ou=groups, dc=example,dc=com
```

For more information on user access roles, see [Adding New User Accounts](#) on page 1899.

3. From the **Default User Role** list, select the default minimum access role for users that do not belong to any of the specified groups.

TIP! Press the Ctrl key while clicking role names to select multiple roles.

4. If you used static groups, in the **Group Member Attribute** field, type the LDAP attribute that designates membership in a static group.

For example, if the **member** attribute is used to indicate membership in the static group you reference for default Security Analyst access, type **member**.

5. If you used dynamic groups, in the **Group Member URL Attribute** field, type the LDAP attribute that contains the LDAP search string used to determine membership in a dynamic group.
For example, if the `memberURL` attribute contains the LDAP search that retrieves members for the dynamic group you specified for default Admin access, type `memberURL`.
6. Continue with [Configuring Administrative Shell Access](#).

Configuring Administrative Shell Access

LICENSE: Any

You can also use the LDAP server to authenticate accounts for shell access on your managed device or Defense Center. Specify a search filter that retrieves entries for users you want to grant shell access. Note that you can only configure shell access for the first authentication object in your system policy. For more information on managing authentication object order, see [Configuring Authentication Profiles](#) on page 1975.

IMPORTANT! IPv6 is not supported for shell access authentication.

With the exception of the admin account, shell access is controlled entirely through the shell access attribute you set. The shell access filter you set determines which set of users on the LDAP server can log into the shell.

Note that a home directory for each shell user is created on login, and when an LDAP shell access user account is disabled (by disabling the LDAP connection), the directory remains, but the user shell is set to `/bin/false` in `/etc/passwd` to disable the shell. If the user then is re-enabled, the shell is reset, using the same home directory.

The **Same as Base Filter** check box allows you to search more efficiently if all users qualified in the base DN are also qualified for shell access privileges. Normally, the LDAP query to retrieve users combines the base filter with the shell access filter. If the shell access filter was the same as the base filter, the same query runs twice, which is unnecessarily time-consuming. You can use the **Same as Base Filter** option to run the query only once for both purposes.

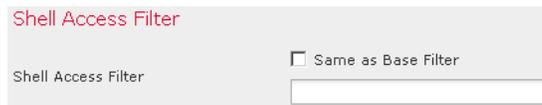
Shell users should log in using user names with all lowercase letters.

WARNING! On Series 3 Defense Centers, all shell users have `sudoers` privileges. Make sure that you restrict the list of users with shell access appropriately. On Series 3 and virtual devices, shell access granted to externally authenticated users defaults to the **Configuration** level of command line access, which also grants `sudoers` privileges.

To configure shell account authentication:

ACCESS: Admin

1. Optionally, on the Create Authentication Object page, set a shell access account filter. You have multiple options:
 - To retrieve administrative user entries based on attribute value, type the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses, in the **Shell Access Filter** field.
 - To use the same filter you specified when configuring authentication settings, select **Same as Base Filter**.
 - To prevent LDAP authentication of shell access, leave the field blank. If you choose not to specify a shell access filter, a warning displays when you save the authentication object to confirm that you meant to leave the filter blank.



The screenshot shows a configuration panel for the 'Shell Access Filter'. At the top, the title 'Shell Access Filter' is displayed in red. Below the title, there is a text input field labeled 'Shell Access Filter' and a checkbox labeled 'Same as Base Filter'.

For example, if all network administrators have a **manager** attribute which has an attribute value of **shell**, you can set a base filter of **(manager=shell)**.

2. Continue with [Testing User Authentication](#).

Testing User Authentication

LICENSE: Any

After you configure LDAP server and authentication settings, you can specify user credentials for a user who should be able to authenticate to test those settings.

For the user name, you can enter the value for the **uid** attribute for the user you want to test with. If you are connecting to a Microsoft Active Directory Server and supplied a shell access attribute in place of **uid**, use the value for that attribute as the user name. You can also specify a fully qualified distinguished name for the user.

The test output lists valid and invalid user names. Valid user names are unique, and can include underscores (**_**), periods (**.**), and hyphens (**-**), but otherwise only alphanumeric characters are supported. Invalid user names are user names containing other non-alphanumeric characters, such as spaces.

Note that testing the connection to servers with more than 1000 users only returns 1000 users because of web interface page size limitations.

TIP! If you mistype the name or password of the test user, the test fails even if the server configuration is correct. Test the server configuration without the additional test parameters first. If that succeeds supply a user name and password to test with the specific user.

To test user authentication:

ACCESS: Admin

1. In the **User Name** and **Password** fields, type the `uid` value or shell access attribute value and password for the user whose credentials should be used to validate access to the LDAP server.

For example, to test to see if you can retrieve the `JSmith` user credentials at the Example company, type `JSmith`.



2. Click **Test**.

A message appears, either indicating success of the test or detailing what settings are missing or need to be corrected. You have two options:

- If the test succeeds, the test output appears at the bottom of the page. Click **Save**. The Login Authentication page appears, with the new object listed.

To enable LDAP authentication using the object on an appliance, you must apply a system policy with that object enabled to the appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

- If the test fails, see [Tuning Your LDAP Authentication Connection](#) on page 1864 for suggestions for troubleshooting the connection. Note that the error message that appears indicates what caused the connection to fail.

LDAP Authentication Object Examples

LICENSE: Any

The following sections provide an example of LDAP configuration using basic settings and an example using more advanced configuration options:

- [Example: Basic LDAP Configuration](#) on page 1880
- [Example: Advanced LDAP Configuration](#) on page 1882

Example: Basic LDAP Configuration

LICENSE: Any

The following figure illustrates a basic configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in

this example has an IP address of 10.11.3.4. The connection uses port 389 for access.

Authentication Object

Authentication Method: LDAP

Name *: Basic Configuration Example

Description:

Server Type: MS Active Directory [Set Defaults]

Primary Server

Host Name/IP Address *: ex. IP or hostname

Port *: 389

Backup Server (Optional)

Host Name/IP Address: ex. IP or hostname

Port: 389

LDAP-Specific Parameters

Base DN *: OU=security,DC=it,DC=example,DC=com ex. dc=sourcefire,dc=com [Fetch DNs]

Base Filter: ex. (cn=jsmith), (lcn=jsmith), (&(cn=jsmith)(|(cn=bsmith)(cn=csmith*)))

User Name *: CN=admin,DC=example,DC=com ex. cn=jsmith,dc=sourcefire,dc=com

Password *:

Confirm Password *:

Show Advanced Options: [Show]

Attribute Mapping

UI Access Attribute *: sAMAccountName [Fetch Attrs]

Shell Access Attribute *: sAMAccountName

User Name:

Password:

*Required Field

[Save] [Test] [Cancel]

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company.

However, because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Selecting the MS Active Directory server type and clicking **Set Defaults** sets the **UI Access Attribute** to `sAMAccountName`. As a result, the Sourcefire 3D System checks the `sAMAccountName` attribute for each object for matching user names when a user attempts to log into the Sourcefire 3D System.

In addition, a **Shell Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a shell account on the appliance.

Note that because no base filter is applied to this server, the Sourcefire 3D System checks attributes for all objects in the directory indicated by the base distinguished name. Connections to the server time out after the default time period (or the timeout period set on the LDAP server).

Example: Advanced LDAP Configuration

LICENSE: Any

This example illustrates an advanced configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 636 for access.

The screenshot shows the 'Authentication Object' configuration page. The 'Authentication Method' is set to 'LDAP'. The 'Name' is 'Advanced Configuration Example'. The 'Server Type' is 'MS Active Directory'. Under the 'Primary Server' section, the 'Host Name/IP Address' is '10.11.3.4' and the 'Port' is '636'. There is a 'Set Defaults' button next to the 'Server Type' dropdown.

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company. However, note that this server has a base filter of `(cn=*smith)`. The filter restricts the users retrieved from the server to those with a common name ending in `smith`.

The screenshot shows the 'LDAP-Specific Parameters' configuration page. The 'Base DN' is 'OU=security,DC=it,DC=example,DC=com'. The 'Base Filter' is '(CN=*smith)'. The 'User Name' is 'CN=admin,DC=example,DC=com'. The 'Password' and 'Confirm Password' fields are masked with dots. Under 'Show Advanced Options', 'Encryption' is set to 'SSL'. The 'SSL Certificate Upload Path' is 'C:\certificate.pem'. The 'User Name Template' is '%s' and the 'Timeout (Seconds)' is '60'. Under the 'Attribute Mapping' section, both 'UI Access Attribute' and 'Shell Access Attribute' are set to 'sAMAccountName'. There are 'Fetch DNs' and 'Fetch Attrs' buttons.

The connection to the server is encrypted using SSL and a certificate named `certificate.pem` is used for the connection. In addition, connections to the server time out after 60 seconds because of the **Timeout** setting.

Because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Note that the configuration includes a **UI Access Attribute** of `sAMAccountName`. As a result, the Sourcefire 3D System checks the `sAMAccountName` attribute for each object for matching user names when a user attempts to log into the Sourcefire 3D System.

In addition, a **Shell Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a shell account on the appliance.

This example also has group settings in place. The Maintenance User role is automatically assigned to all members of the group with a `member` group attribute and the base domain name of `CN=SFmaintenance,DC=it,DC=example,DC=com`.

Group Controlled Access Roles (Optional) ▾

Access Admin	<input type="text"/>
Administrator	<input type="text"/>
External Database User	<input type="text"/>
Intrusion Admin	<input type="text"/>
Maintenance User	<input type="text" value="CN=SFmaintenance,DC=it,DC=example,DC=com"/>
Network Admin	<input type="text"/>
Discovery Admin	<input type="text"/>
Security Approver	<input type="text"/>
Security Analyst	<input type="text"/>
Security Analyst (Read Only)	<input type="text"/>
Default User Role	<input type="text" value="Access Admin"/>
Group Member Attribute	<input type="text" value="member"/>
Group Member URL Attribute	<input type="text"/>

The shell access filter is set to be the same as the base filter, so the same users can access the appliance through the shell as through the web interface.

Shell Access Filter

Same as Base Filter

Shell Access Filter

Additional Test Parameters

User Name

Password

*Required Field

Editing LDAP Authentication Objects

LICENSE: Any

You can edit an existing authentication object. Your changes do not take effect until you reapply the policy.

To edit an authentication object:

ACCESS: Admin

1. Select **System > Local > User Management**.

The User Management page appears

2. Click the **Login Authentication** tab.

The Login Authentication page appears.

3. Click the edit icon (✎) next to the object you want to edit.

The Create Authentication Object page appears.

4. Modify the object settings as needed.

For more information, see the following topics:

- [Quick Start to LDAP Authentication](#) on page 1860
- [Creating Advanced LDAP Authentication Objects](#) on page 1866
- [Identifying the LDAP Authentication Server](#) on page 1867
- [Configuring LDAP-Specific Parameters](#) on page 1869
- [Configuring Access Settings by Group](#) on page 1875
- [Configuring Administrative Shell Access](#) on page 1878
- [Testing User Authentication](#) on page 1879

5. Click **Test**.

A message appears, either indicating success of the test or detailing what settings are missing or need to be corrected. If the test succeeds, the test output appears at the bottom of the page.

If the test fails, see [Tuning Your LDAP Authentication Connection](#) on page 1864 for suggestions for troubleshooting the connection. Note that the error message that appears indicates what caused the connection to fail.

6. Click **Save**.

Your changes are saved and the Login Authentication page appears. Remember that you have to apply a system policy with the object enabled to an appliance before the authentication changes take place on that appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

Understanding RADIUS Authentication

LICENSE: Any

The Remote Authentication Dial In User Service (RADIUS) is an authentication protocol used to authenticate, authorize, and account for user access to network resources. You can create an authentication object for any RADIUS server that conforms to RFC 2865.

When a user authenticated on a RADIUS server logs in for the first time, the user receives the roles specified for that user in the authentication object, or if the user is not listed for any of the user roles, the default access role you selected in the authentication object, or failing that, the system policy. You can modify a user's roles, if needed, unless the settings are granted through the user lists in the authentication object. Note that when a user authenticated on a RADIUS server using attribute matching attempts to log in for the first time, the login is rejected as the user account is created. The user must log in a second time.

The Sourcefire 3D System implementation of RADIUS supports the use of SecurID® tokens. When you configure authentication by a server using SecurID, users authenticated against that server append the SecurID token to the end of their SecurID pin and use that as their password when they log into a Sourcefire appliance. As long as SecurID is configured correctly to authenticate users outside the Sourcefire 3D System, those users can log into a Sourcefire 3D System appliance using their PIN plus the SecurID token without any additional configuration on the appliance.

Creating RADIUS Authentication Objects

LICENSE: Any

When you create a RADIUS authentication object, you define settings that let you connect to an authentication server. You also grant user roles to specific and default users. If your RADIUS server returns custom attributes for any users you plan to authenticate, you must define those custom attributes. Optionally, you can also configure shell access authentication.

Note that to create an authentication object, you need TCP/IP access from your local appliance to the authentication server where you want to connect.

To create an authentication object:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears
2. Click the **Login Authentication** tab.
The Login Authentication page appears.
3. Click **Create Authentication Object**.
The Create Authentication Object page appears.

4. Identify the primary and backup authentication servers where you want to retrieve user data for external authentication and set timeout and retry values. For more information, see [Configuring RADIUS Connection Settings](#) on page 1886.
5. Set the default user role. Optionally, specify the users or user attribute values for users that you want to receive specific Sourcefire 3D System access roles. For more information, see [Configuring RADIUS User Roles](#) on page 1888.
6. Optionally, configure administrative shell access. For more information, see [Configuring Administrative Shell Access](#) on page 1891.
7. If the profiles for any of the users to authenticate return custom RADIUS attributes, define those attributes. For more information, see [Defining Custom RADIUS Attributes](#) on page 1891.
8. Test your configuration by entering the name and password for a user who should successfully authenticate. For more information, see [Testing User Authentication](#) on page 1893.

Your changes are saved. Remember that you have to apply a system policy with the object enabled to an appliance before the authentication changes take place on that appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

Configuring RADIUS Connection Settings

LICENSE: Any

When you create a RADIUS authentication object, you first specify the primary and backup server and server port where you want the local appliance (managed device or Defense Center) to connect for authentication.

IMPORTANT! For RADIUS to function correctly, you must open its authentication and accounting ports (by default, 1812 and 1813) on your firewall.

If you specify a backup authentication server, you can set a timeout for the connection attempt to the primary server. If the number of seconds indicated in the **Timeout** field (or the timeout on the LDAP server) elapses without a response from the primary authentication server, the appliance then re-queries the primary server.

After the appliance re-queries the primary authentication server the number of times indicated by the **Retries** field and the number of seconds indicated in the **Timeout** field again elapses without a response from the primary authentication server, the appliance then rolls over to the backup server.

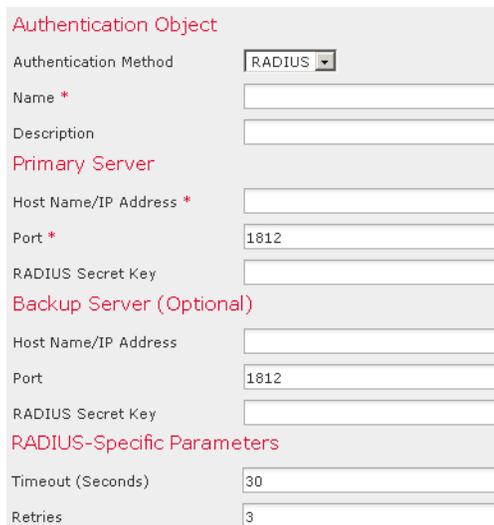
If, for example, the primary server has RADIUS disabled, the appliance queries the backup server. If RADIUS is running on the port of the primary RADIUS server

and for some reason refuses to service the request (due to misconfiguration or other issues), however, the failover to the backup server does not occur.

To identify a RADIUS authentication server:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears
2. Click the **Login Authentication** tab.
The Login Authentication page appears.
3. Click **Create Authentication Object**.
The Create Authentication Object page appears.
4. Select **RADIUS** from the **Authentication Method** drop-down list.
RADIUS configuration options appear.



Authentication Object

Authentication Method: RADIUS

Name *

Description

Primary Server

Host Name/IP Address *

Port * 1812

RADIUS Secret Key

Backup Server (Optional)

Host Name/IP Address

Port 1812

RADIUS Secret Key

RADIUS-Specific Parameters

Timeout (Seconds) 30

Retries 3

5. Type a name and description for the authentication server in the **Name** and **Description** fields.
6. Type the IP address or host name for the primary RADIUS server where you want to obtain authentication data in the **Primary Server Host Name/IP Address** field.

IMPORTANT! IPv6 addresses are not supported.

7. Optionally, modify the port used by the primary RADIUS authentication server in the **Primary Server Port** field.

IMPORTANT! If your authentication port and accounting port numbers are not sequential, leave this field blank. The system then determines RADIUS port numbers from the `radius` and `radacct` data in your appliance's `/etc/services` file.

8. Type the secret key for the primary RADIUS authentication server in the **RADIUS Secret Key** field.
9. Type the IP address or host name for the backup RADIUS authentication server where you want to obtain authentication data in the **Backup Server Host Name/IP Address** field.
10. Optionally, modify the port used by the backup RADIUS authentication server in the **Backup Server Port** field.

IMPORTANT! If your authentication port and accounting port numbers are not sequential, leave this field blank. The system then determines RADIUS port numbers from the `radius` and `radacct` data in your appliance's `/etc/services` file.

11. Type the secret key for the backup RADIUS authentication server in the **RADIUS Secret Key** field.
12. Type the number of seconds that should elapse before retrying the connection in the **Timeout** field.
13. Type the number of times the primary server connection should be tried before rolling over to the backup connection in the **Retries** field.
14. Continue with [Configuring RADIUS User Roles](#).

Configuring RADIUS User Roles

LICENSE: Any

You can specify the access roles for existing users on your RADIUS server by listing the user names for each of the access roles used by your Sourcefire 3D System. When you do so, you can also configure a default access setting for those users detected by RADIUS that are not specified for a particular role.

When a user logs in, the Sourcefire 3D System checks the RADIUS server and grants access rights depending on the RADIUS configuration:

- If specific access settings are not configured for a user and a default access role is not selected, when a new user logs in, the Sourcefire 3D System authenticates the user against the RADIUS server and then grants user rights based on the default access role (or roles) set in the system policy.
- If a new user is not specified on any lists and default access roles are selected in the **Default User Role** list of the authentication object, the user is assigned those access roles.
- If you add a user to the list for one or more specific role, that user receives all assigned access roles.

You can also use attribute-value pairs, rather than user names, to identify users who should receive a particular user role. For example, if you know all users who should be Security Analysts have the value `Analyst` for their `user-category` attribute, you can type `user-category=Analyst` in the Security Analyst List field to grant that role to those users. Note that you need to define any custom attributes before you use them to set user role membership. For more information, see [Defining Custom RADIUS Attributes](#) on page 1891.

You can assign a default user role (or roles) to be assigned to any users that are authenticated externally but not listed for a specific role. You can select multiple roles on the **Default User Role** list.

For more information on the user roles supported by the Sourcefire 3D System, see [Configuring User Roles](#) on page 1905.

You cannot remove the minimum access rights for users assigned an access role because of RADIUS user list membership through the Sourcefire 3D System user management page. You can, however, assign additional rights.

WARNING! If you want to change the minimum access setting for a user, you must not only move the user from one list to another in the RADIUS Specific Parameters section or change the user's attribute on the RADIUS server, you must reapply the system policy, and you must remove the assigned user right on the user management page.

To base access on user lists:

ACCESS: Admin

1. In the fields that correspond to Sourcefire 3D System user roles, type the name of each user or identifying attribute-value pair that should be assigned to those roles. Separate usernames and attribute-value pairs with commas. For example, to grant the Administrator role to the users `jsmith` and `jdoe`, type `jsmith, jdoe` in the **Administrator** field.

RADIUS-Specific Parameters

Timeout (Seconds)	<input type="text" value="30"/>
Retries	<input type="text" value="3"/>
Access Admin	<input type="text"/>
Administrator	<input type="text" value="jsmith, jdoe"/>
External Database User	<input type="text"/>
Intrusion Admin	<input type="text"/>
Maintenance User	<input type="text"/>
Network Admin	<input type="text"/>
Discovery Admin	<input type="text"/>
Security Approver	<input type="text"/>
Security Analyst	<input type="text"/>
Security Analyst (Read Only)	<input type="text"/>
Default User Role	<input type="list" value="Intrusion Admin"/>

As another example, to grant the Maintenance User role to all users with a **User-Category** value of `Maintenance`, type `User-Category=Maintenance` in the **Maintenance User** field.

For more information on user access roles, see [Configuring User Roles](#) on page 1905.

2. Select the default minimum access role for users that do not belong to any of the specified groups from the **Default User Role** list.

TIP! Press the Ctrl key while clicking role names to select multiple roles.

3. Continue with [Configuring Administrative Shell Access](#).

Configuring Administrative Shell Access

LICENSE: Any

You can also use the RADIUS server to authenticate accounts for shell access on your local appliance (managed device or Defense Center). Specify user names for users you want to grant shell access. Note that you can only configure shell access for the first authentication object in your system policy. For more information on managing authentication object order, see [Configuring Authentication Profiles](#) on page 1975.

With the exception of the admin account, the shell access list you set on the RADIUS authentication object entirely controls shell access on the appliance. Shell users are configured as local users on the appliance when the system policy is applied. Note that when a user authenticated on a RADIUS server using attribute matching attempts to log in for the first time, the login is rejected as the user account is created. The user must log in a second time.

Note that a home directory for each shell user is created on login, and when an RADIUS shell access user account is disabled (by disabling the RADIUS connection), the directory remains, but the user shell is set to `/bin/false` in `/etc/passwd` to disable the shell. If the user then is re-enabled, the shell is reset, using the same home directory.

Shell users should log in using user names with all lowercase letters.

WARNING! On Series 3 Defense Centers, all shell users have `sudoers` privileges. Make sure that you restrict the list of users with shell access appropriately. On Series 3 and virtual devices, shell access granted to externally authenticated users defaults to the **Configuration** level of command line access, which also grants `sudoers` privileges.

To configure shell account authentication:

ACCESS: Admin

1. Type the user names, separated by commas, in the **Administrator Shell Access User List** field.

IMPORTANT! If you choose not to specify a shell access filter, a warning displays when you save the authentication object to confirm that you meant to leave the filter blank.

2. Continue with [Defining Custom RADIUS Attributes](#) on page 1891.

Defining Custom RADIUS Attributes

LICENSE: Any

If your RADIUS server returns values for attributes not included in the dictionary file in `/etc/radiusclient/` and you plan to use those attributes to

set user roles for users with those attributes, you need to define those attributes in the login authentication object.

You can locate the attributes returned for a user by looking at the user's profile on your RADIUS server.

When you define an attribute, you provide the name of the attribute, which consists of alphanumeric characters. Note that words in an attribute name should be separated by dashes rather than spaces. You also provide the attribute ID, which should be an integer and should not conflict with any existing attribute IDs in the `etc/radiusclient/dictionary` file. You also specify the type of attribute: string, IP address, integer, or date.

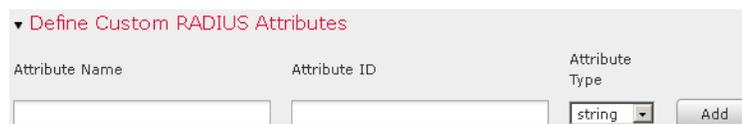
As an example, if a RADIUS server is used on a network with a Cisco router, you might want to use the `Ascend-Assign-IP-Pool` attribute to grant a specific role to all users logging in from a specific IP address pool. `Ascend-Assign-IP-Pool` is an integer attribute that defines the address pool where the user is allowed to log in, with the integer indicating the number of the assigned IP address pool. To declare that custom attribute, you create a custom attribute with an attribute name of `Ascend-IP-Pool-Definition`, an attribute ID of `218`, and an attribute type of `integer`. You could then type `Ascend-Assign-IP-Pool=2` in the **Security Analyst (Read Only)** field to grant read-only security analyst rights to all users with an `Ascend-IP-Pool-Definition` attribute value of `2`.

When you create a RADIUS authentication object, a new dictionary file for that object is created on the Sourcefire 3D System appliance in the `/var/sf/userauth` directory. Any custom attributes you add to the authentication object are added to the dictionary file.

To define a custom attribute:

ACCESS: Admin

1. Click the arrow to expand the Define Custom RADIUS Attributes section. The attribute fields appear.



Attribute Name	Attribute ID	Attribute Type
<input type="text"/>	<input type="text"/>	string

2. Type an attribute name consisting of alphanumeric characters and dashes, with no spaces, in the **Attribute Name** field.
3. Type the attribute ID, in integer form, in the **Attribute ID** field.
4. Select the type of attribute from the **Attribute Type** drop-down list.
5. Click **Add** to add the custom attribute to the authentication object.

TIP! You can remove a custom attribute from an authentication object by clicking **Delete** next to the attribute.

6. Continue with [Testing User Authentication](#) on page 1893.

Testing User Authentication

LICENSE: Any

After you configure RADIUS connection, user role, and custom attribute settings, you can specify user credentials for a user who should be able to authenticate to test those settings.

For the user name, you can enter the user name for the user you want to test with.

Note that testing the connection to servers with more than 1000 users only returns 1000 users because of UI page size limitations.

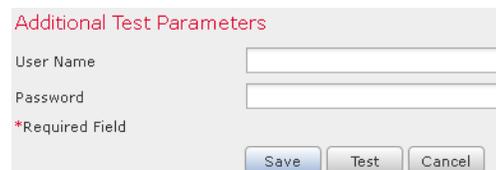
TIP! If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

To test user authentication:

ACCESS: Admin

1. In the **User Name** and **Password** fields, type the user name and password for the user whose credentials should be used to validate access to the RADIUS server.

For example, to test to see if you can retrieve the `jsmith` user credentials at our example company, type `jsmith`.



2. Select **Show Details** and click **Test**.

A message appears, either indicating success of the test or detailing what settings are missing or need to be corrected.

3. If the test succeeds, click **Save**.

The Login Authentication page appears, with the new object listed.

To enable RADIUS authentication using the object on an appliance, you must apply a system policy with that object enabled to the appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

RADIUS Authentication Object Examples

LICENSE: Any

This section provides examples of RADIUS server authentication objects to show how Sourcefire 3D System RADIUS authentication features can be used. See the following sections for more information:

- [Authenticating a User Using RADIUS](#) on page 1894
- [Authenticating a User with Custom Attributes](#) on page 1895

Authenticating a User Using RADIUS

LICENSE: Any

The following figure illustrates a sample RADIUS login authentication object for a server running FreeRADIUS with an IP address of 10.10.10.98. Note that the connection uses port 1812 for access, and note that connections to the server time out after 30 seconds of disuse, then retry three times before attempting to connect to a backup authentication server.

This example illustrates important aspects of RADIUS user role configuration:

- Users **ewharton** and **gsand** are granted administrative access to Sourcefire 3D System appliances where this authentication object is enabled.
- The user **cbronte** is granted Maintenance User access to Sourcefire 3D System appliances where this authentication object is enabled.
- The user **jausten** is granted Security Analyst access to Sourcefire 3D System appliances where this authentication object is enabled.
- The user **ewharton** can log into the appliance using a shell account.

The following graphic depicts the role configuration for the example:

The screenshot shows a configuration window titled "RADIUS-Specific Parameters". It contains several fields for user roles and a "Shell Access Filter" section. The fields are:

- Timeout (Seconds): 30
- Retries: 3
- Access Admin: (empty)
- Administrator: ewharton, gsand
- External Database User: (empty)
- Intrusion Admin: (empty)
- Maintenance User: cbronte
- Network Admin: (empty)
- Discovery Admin: (empty)
- Security Approver: (empty)
- Security Analyst: jausten
- Security Analyst (Read Only): (empty)
- Default User Role: A dropdown menu with options: Access Admin, Administrator, External Database User, and Intrusion Admin. "Intrusion Admin" is selected.

Below these fields is the "Shell Access Filter" section, which includes a field for "Administrator Shell Access User List" containing the value "ewharton".

Authenticating a User with Custom Attributes

LICENSE: Any

You can use an attribute-value pair to identify users who should receive a particular user role. If the attribute you use is a custom attribute, you must define the custom attribute.

The following figure illustrates the role configuration and custom attribute definition in a sample RADIUS login authentication object for the same FreeRADIUS server as in the previous example.

In this example, however, the `MS-RAS-version` custom attribute is returned for one or more of the users because a Microsoft remote access server is in use. Note the `MS-RAS-version` custom attribute is a string. In this example, all users logging in to RADIUS through a Microsoft v. 5.00 remote access server should

receive the Security Analyst (Read Only) role, so you type the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Security Analyst (Read Only)** field.

RADIUS-Specific Parameters

Timeout (Seconds)	30
Retries	3
Access Admin	
Administrator	ewharton, gsand
External Database User	
Intrusion Admin	
Maintenance User	
Network Admin	
Discovery Admin	
Security Approver	
Security Analyst	
Security Analyst (Read Only)	MS-RAS-Version=MSRASV5.00
Default User Role	Access Admin Administrator External Database User Intrusion Admin

Shell Access Filter

Administrator Shell Access	ewharton
User List	

Define Custom RADIUS Attributes

Attribute Name	Attribute ID	Attribute Type	
		string	Add
MS-Ras-Version	18	string	Delete

Editing RADIUS Authentication Objects

LICENSE: Any

You can edit an existing authentication object. If the object is in use in a system policy, the settings in place at the time the policy was applied stay in effect until you reapply the policy.

To edit an authentication object:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears

2. Click the **Login Authentication** tab.
The Login Authentication page appears.
3. Click the edit icon (✎) next to the object you want to edit.
The Create Authentication Object page appears.
4. Modify the object settings as needed.
For more information, see the following topics:
 - [Creating RADIUS Authentication Objects](#) on page 1885
 - [Configuring RADIUS Connection Settings](#) on page 1886
 - [Configuring RADIUS User Roles](#) on page 1888
 - [Configuring Administrative Shell Access](#) on page 1891
 - [Testing User Authentication](#) on page 1893
5. Click **Save**.
Your changes are saved and the Login Authentication page reappears. Remember that you have to apply a system policy with the object enabled to an appliance before the authentication changes take place on that appliance. For more information, see [Configuring Authentication Profiles](#) on page 1975 and [Applying a System Policy](#) on page 1965.

Deleting Authentication Objects

LICENSE: Any

You can delete an authentication object if it is not currently enabled in a system policy.

To delete an authentication object:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears
2. Click the **Login Authentication** tab.
The Login Authentication page appears.
3. Click the delete icon (🗑) next to the object you want to delete.
The object is deleted and the Login Authentication page appears.

Managing User Accounts

LICENSE: Any

If you have Administrator access, you can use the web interface to view and manage user accounts on a Defense Center or a managed device, including adding, modifying, and deleting accounts. You can also create and modify custom

user roles and configure user role escalation. User accounts without Administrator access are restricted from accessing management features. The navigation menu differs in appearance for each type of user.

See the following sections for more information about managing user accounts:

- [Viewing User Accounts](#) on page 1898 explains how to access the User Management page, where you can add, activate, deactivate, edit, and delete user accounts.
- [Adding New User Accounts](#) on page 1899 describes the different options you can use when you add a new user account.
- [Managing Externally Authenticated User Accounts](#) on page 1902 explains how externally authenticated users are added and what aspects of the user configuration you can manage within the Sourcefire 3D System.
- [Modifying User Privileges and Options](#) on page 1912 explains how to access and modify an existing user account.
- [Understanding Restricted User Access Properties](#) on page 1912 explains how to restrict the data available to a user account with restricted data access.
- [Deleting User Accounts](#) on page 1914 explains how to delete user accounts.
- [User Account Privileges](#) on page 1914 contains tables that list the menus and options each type of user account can access.

Viewing User Accounts

LICENSE: Any

From the User Management page, you can view, edit, and delete existing accounts. You can view the type of authentication for a user from the **Authentication Method** column. The **Password Lifetime** column indicates the days remaining on each user's password. The icons in the **Action** column allow you to edit users in more detail and set users active or inactive. Note that for externally authenticated users, if the authentication object for the server is disabled, the **Authentication Method** column displays **External (Disabled)**.

To access the User Management page:

ACCESS: Admin

- ▶ Select **System > Local > User Management**.

The User Management page appears, showing each user, with options to activate, deactivate, edit, or delete the user account.

See the following sections for information about the actions you can perform on the User Management page:

- [Adding New User Accounts](#) on page 1899
- [Configuring User Roles](#) on page 1905

- [Modifying User Privileges and Options](#) on page 1912
- [Understanding Restricted User Access Properties](#) on page 1912
- [Modifying User Passwords](#) on page 1913
- [Deleting User Accounts](#) on page 1914

Adding New User Accounts

LICENSE: Any

SUPPORTED DEVICES: feature dependent

When you set up a new user account, you can control which parts of the system the account can access. You can set password expiration and strength settings for the user account during creation. For a local account on a Series 3 device, you can also configure the level of command line access the user will have.

To add a new user:

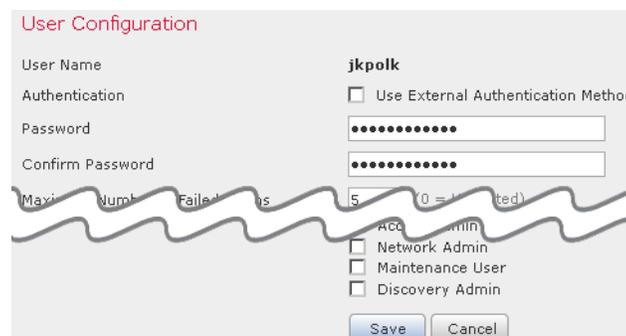
ACCESS: Admin

1. Select **System > Local > User Management**.

The User Management page appears.

2. Click **Create User**.

The Create User page appears.



3. In the **User Name** field, type a name for the new user.

New user names must contain alphanumeric or hyphen characters with no spaces, and must be no more than 32 characters. User names are case sensitive.

4. If you want this user to authenticate to an external directory server on login, select **Use External Authentication Method**.

If you enable this option, the password management options disappear. Skip to step 8 to continue with configuring an access role for the user.

Note that for users to authenticate to an external directory server, you must use the Defense Center to create an authentication object for the server you want to use, then apply a system policy with authentication enabled. In addition, the external authentication server must be available in order for

those users to log into Sourcefire 3D System appliances. For more information, see [Managing Authentication Objects](#) on page 1854 and [Configuring Authentication Profiles](#) on page 1975.

5. In the **Password** and **Confirm Password** fields, type a password (up to 32 alphanumeric characters).

If you enable password strength checking, the password must be at least eight alphanumeric characters of mixed case and must include at least one numeric character and one special character. It cannot be a word that appears in a dictionary or include consecutive repeating characters.

6. Configure the remaining user account login options.

For more information, see the [User Account Login Options table](#) on page 1904.

7. If you are creating a local user through the web interface of a Series 3 device, you can assign the level of **Command-Line Interface Access** for the user:

- Select **None** to disable access to the command line for the user.
- Select **Basic** to allow the user to log into the shell and to access a specific subset of commands.
- Select **Configuration** to allow the user to log into the shell and use any command line option, including expert mode if that is allowed on the appliance.

For more information on command line access, see [Managing Command Line Access](#) on page 1901.

8. Select access roles to grant to the user.

IMPORTANT! For all physical managed devices, the Sourcefire-provided predefined user roles are limited to Administrator, Maintenance User, and Security Analyst.

For more information, see [Configuring User Roles](#) on page 1905.

9. Click **Save**.

The user is created and the User Management page appears again.

TIP! Click the slider next to the name of an internally authenticated user on the User Management page to reactivate a deactivated user, or to disable an active user account without deleting it.

Managing Command Line Access

LICENSE: Any

SUPPORTED DEVICES: Series 3, virtual

On a Series 3 or virtual device, you can assign command line interface access to local device users.

Note that you can also assign command line access for users on a virtual device, but you use commands from the command line interface. For more information, see [Command Line Reference](#) on page 2242.

The commands a user can run depend on the level of access you assign to the user. When you set **Command-Line Interface Access** to **None**, the user cannot log into the appliance on the command line. Any session the user starts will close when the user provides credentials. The access level defaults to **None** on user creation. When you set **Command-Line Interface Access** to **Basic**, a specific set of commands can be run by the user:

Basic Command Line Commands

configure password	interfaces
end	lcd
exit	link-state
expert	managers
help	memory
history	model
logout	mpls-depth
?	NAT
??	network
access-control-config	network-modules
alarms	ntp
audit-log	perfstats
clustering	portstats
cpu	power-supply-status
database	process-tree

Basic Command Line Commands (Continued)

device-settings	processes
disk	routing-table
dns	serial-number
bypass	stacking
fan-status	summary
fastpath-rules	time
gui	traffic-statistics
hostname	version
hyperthreading	version
inline-sets	virtual-routers

When you set **Command-Line Interface Access** to **Configuration**, the user can access any of the command line options. Exercise caution in assigning this level of access to users.

WARNING! Shell access granted to externally authenticated users defaults to the **Configuration** level of command line access, granting rights to all command line utilities. For more information on shell access for externally authenticated users, see [Setting up Shell Access](#) on page 1858, [Configuring Administrative Shell Access](#) on page 1878, and [Configuring Administrative Shell Access](#) on page 1891.

Managing Externally Authenticated User Accounts

LICENSE: Any

When an externally authenticated user logs into an appliance that has external authentication enabled, the appliance grants the user the default access role you set by specifying group membership in the authentication object. If you did not configure access group settings, the appliance grants the default user role you set in the system policy. However, if you add users locally before they log into the appliance, the user privileges you configure on the User Management page override the default settings.

An internally authenticated user is converted to external authentication when all of the following conditions exist:

- You enable LDAP or RADIUS authentication.
- The same user name exists for the user on the LDAP or RADIUS server.
- The user logs in using the password stored for that user on the LDAP or RADIUS server.

For more information on selecting a default user role, see [Configuring Authentication Profiles](#) on page 1975 and [Understanding User Privileges](#) on page 1852. Note that you can set both predefined and custom user roles as the default user role for externally authenticated users. For more information, see [Configuring User Roles](#) on page 1905.

Note that you can only enable external authentication in a system policy on a Defense Center. You must use the Defense Center to apply the policy to managed devices if you want to use external authentication on them.

For more information on modifying user access, see [Modifying User Privileges and Options](#) on page 1912. Note that you cannot manage passwords for externally authenticated users or deactivate externally authenticated users through the Sourcefire 3D System interface. For externally authenticated users, you cannot remove the minimum access rights through the Sourcefire 3D System user management page for users assigned an access role because of LDAP group or RADIUS list membership or attribute values. On the Edit User page for an externally authenticated user, rights granted because of settings on an external authentication server are marked with a status of **Externally Modified**.

You can, however, assign additional rights. When you modify the access rights for an externally authenticated user, the Authentication Method column on the User Management page provides a status of **External - Locally Modified**.

should log in using user names with all lowercase letters.

WARNING! On Series 3 Defense Centers, all shell users have **sudoers** privileges. Make sure that you restrict the list of users with shell access appropriately. On Series 3 and virtual devices, shell access granted to externally authenticated users defaults to the **Configuration** level of command line access, which also grants **sudoers** privileges. For more information on setting up shell access, see [Setting up Shell Access](#) on page 1858, [Configuring Administrative Shell Access](#) on page 1878, and [Configuring Administrative Shell Access](#) on page 1891.

Managing User Login Settings

LICENSE: Any

You can control how and when the password for each user account is changed, as well as when user accounts are disabled. If you configured a timeout for web interface login sessions, you can exempt users from this timeout. The [User Account Login Options](#) table describes some of the options you can use to

regulate passwords and account access.

Note that for locally authenticated users on Series 3 managed devices, changing a user's password for the web interface also changes that password for the command line interface.

If you enable the **Check Password Strength** option, the minimum password length is automatically set to 8 characters. If you also set a value for **Minimum Password Length** that exceeds 8 characters, the higher value applies.

IMPORTANT! After you enable **Use External Authentication Method**, login options no longer appear. Use the external authentication server to manage login settings.

User Account Login Options

OPTION	DESCRIPTION
Use External Authentication Method	<p>Select this check box if you want this user's credentials to be externally authenticated.</p> <p>IMPORTANT! If you select this option for the user and the external authentication server is unavailable, that user can log into the web interface but cannot access any functionality.</p>
Maximum Number of Failed Logins	<p>Enter an integer, without spaces, that determines the maximum number of times each user can try to log in after a failed login attempt before the account is locked. The default setting is five tries; use 0 to allow an unlimited number of failed logins.</p>
Minimum Password Length	<p>Enter an integer, without spaces, that determines the minimum required length, in characters, of a user's password. The default setting is 8. A value of 0 indicates that no minimum length is required.</p>
Days Until Password Expiration	<p>Enter the number of days after which the user's password expires. The default setting is 0, which indicates that the password never expires.</p>
Days Before Password Expiration Warning	<p>Enter the number of warning days users have to change their password before their password actually expires. The default setting is 0 days.</p> <p>WARNING! The number of warning days must be less than the number of days before the password expires.</p>
Force Password Reset on Login	<p>Select this option to force users to change their passwords the first time they log in.</p>

User Account Login Options (Continued)

OPTION	DESCRIPTION
Check Password Strength	Select this option to require strong passwords. A strong password must be at least eight alphanumeric characters of mixed case and must include at least one numeric character and one special character. It cannot be a word that appears in a dictionary or include consecutive repeating characters.
Exempt from Browser Session Timeout	Select this option if you do not want a user's login sessions to terminate due to inactivity. Users with the Administrator role cannot be made exempt. For more information on session timeouts, see Configuring User Interface Settings on page 1994.

Configuring User Roles

LICENSE: Any

Each Sourcefire 3D System user has an associated user access role or roles. For example, an analyst needs access to event data to analyze the security of your network, but might not require access to administrative functions for the Sourcefire 3D System itself. Using user roles, you can, for example, grant Security Analyst access to analysts while reserving the Administrator role for the user or users managing the Sourcefire 3D System. The Sourcefire 3D System includes ten predefined user roles designed for a variety of administrators and analysts. You can also create custom user roles with specialized access privileges.

The menus and other options in the web interface that users can access depend on their roles. Predefined user roles have a set of predetermined access privileges, while custom user roles have granular access privileges that their creator determines.

You configure user roles on the User Roles page.

To access the User Roles page:

ACCESS: Admin

1. Select **System > Local > User Management**.

The User Management page appears.

2. Click the **User Roles** tab.

The User Roles page appears, showing all predefined and custom user roles, with options to activate, deactivate, edit, copy, delete, and export roles.

For more information on configuring the two types of user roles, see the following sections:

- [Managing Predefined User Roles](#) on page 1906
- [Managing Custom User Roles](#) on page 1908
- [Creating a Custom Copy of a Predefined User Role](#) on page 1911
- [Deleting a Custom User Role](#) on page 1911

Managing Predefined User Roles

LICENSE: Any

The Sourcefire 3D System includes ten predefined user roles that provide a range of access privilege sets to meet the needs of your organization. On the User Roles page, predefined user roles are labeled “Sourcefire Provided.” Note that managed devices have access to only three of the ten predefined user roles: Administrator, Maintenance User, and Security Analyst.

Although you cannot edit predefined user roles, you can use their access privilege sets as the basis for custom user roles. For information on creating and editing custom user roles, see [Managing Custom User Roles](#) on page 1908. In addition, because you cannot edit predefined user roles, you cannot configure them to escalate to another user role. For more information, see [Managing User Role Escalation](#) on page 1925.

The [Predefined User Roles](#) table briefly describes the predefined roles available to you. For a list of the menus and options available to each role, see [User Account Privileges](#) on page 1914.

Predefined User Roles

USER ROLE	PRIVILEGES
Access Admin	Provides access to access control policy features. Note, however, that Access Admins cannot apply access control policies. Access Admins have access to access control-related options in the Policies menu.
Administrator	Provides access to analysis and reporting features, rule and policy configuration, system management, and all maintenance features. Administrators have access to all menu options; their sessions present a higher security risk if compromised, so you cannot make them exempt from login session timeouts. Note that you should limit use of the Administrator role for security reasons. This role is also available on managed devices.
Discovery Admin	Provides access to network discovery, correlation, and user activity features. Discovery Admins have access to relevant options in the Policies menu.

Predefined User Roles (Continued)

USER ROLE	PRIVILEGES
External Database User	Provides read-only access to the Sourcefire 3D System database using an application that supports JDBC SSL connections. Note that for the third-party application to authenticate to the Sourcefire 3D System appliance, you must enable database access in the system settings as described in Enabling Access to the Database on page 2007. On the web interface, External Database Users have access only to online help-related options in the Help menu. Because this role's function does not involve the web interface, access is provided only for ease of support and password changes.
Intrusion Admin	Provides access to all intrusion policy and intrusion rule features. Intrusion Admins have access to intrusion-related options in the Policies menu. Note that Intrusion Admins cannot apply intrusion policies as part of access control policies.
Maintenance User	Provides access to monitoring and maintenance features. Maintenance Users have access to maintenance-related options in the Health and System menus. This role is also available on managed devices.
Network Admin	Provides access to access control and device configuration features. Network Admins have access to access control and device-related options in the Policies and Devices menus.
Security Approver	Provides limited access to access control policies. Security Approvers can view and apply intrusion and access control policies, but cannot make policy changes. They have access to applicable policy-related options in the Policies menu.
Security Analyst	Provides access to security event analysis features, including event views, reports, hosts, host attributes, services, vulnerabilities, client applications, and read-only access to health events. Security Analysts have access to analysis-related options in the Overview , Analysis , Health , and System menus. This role is also available on managed devices.
Security Analyst (Read Only)	Provides read-only access to security event analysis features, including event views, reports, hosts, host attributes, services, vulnerabilities, client applications, and health events. Security Analysts have access to analysis-related options in the Overview , Analysis , Health , and System menus.

Along with assigning an event analyst role to a user, you can restrict that user's deletion rights to only allow deletion of report profiles, searches, bookmarks, custom tables, and custom workflows created by that user. For more information, see [Adding New User Accounts](#) on page 1899.

Note that externally authenticated users, if assigned no other roles, have minimum access rights based on the settings in LDAP or RADIUS authentication objects and in the system policy. You can assign additional rights to these users,

but to remove or change minimum access rights, you must perform the following tasks:

- Move the user from one list to another in the authentication object or change the user's attribute value or group membership on the external authentication server.
- Reapply the system policy.
- Use the User Management page to remove the access from that user account.

You cannot delete predefined user roles, but you can deactivate them. Deactivating a role removes that role and all associated permissions from any user who is assigned that role.

WARNING! If a deactivated role is the only role assigned to a given user, that user can log in and access the User Preferences menu, but is otherwise unable to access the Sourcefire 3D System.

To activate or deactivate a user role:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Click the **User Roles** tab.
The User Roles page appears.
3. Click the slider next to the user role you want to activate or deactivate.

IMPORTANT! If you deactivate, then reactivate, a role with Lights-Out Management while a user with that role is logged in, or restore a user or user role from a backup during that user's login session, that user must log back into the web interface to regain access to IPMItool commands. For more information, see [Using Lights-Out Management](#) on page 2030.

Managing Custom User Roles

LICENSE: Any

In addition to the predefined user roles, you can also create custom user roles with specialized access privileges. Custom user roles can have any set of menu-based and system permissions, and may be completely original or based on a predefined user role. Like predefined user roles, custom roles can serve as the default role for externally authenticated users. Unlike predefined roles, you can modify and delete custom roles.

Selectable permissions are hierarchical, and are based on the Sourcefire 3D System menu layout. Permissions are expandable if they have sub-pages or if they have more fine-grained permissions available beyond simple page access. In that case, the parent permission grants page view access and the children granular access to related features of that page. For example, the Correlation Events permission grants access to the Correlation Events page, while the Modify Correlation Events check box allows the user to edit and delete the information available on that page. Permissions that contain the word “Manage” grant the ability to edit and delete information that other users create.

You can apply restricted searches to a custom user role. These constrain the data a user may see in the event viewer. You can configure a restricted search by first creating a private saved search and selecting it from the “Restricted Search” drop-down menu under the appropriate menu-based permission. For more information, see [Performing a Search](#) on page 1771.

When you configure a custom user role on a Defense Center, all menu-based permissions are available for you to grant. When you configure a custom user role on a managed device, only some permissions are available — those relevant to device functions. For more information on the menu-based permissions you can configure and their relationship with predefined user roles, see:

- [Analysis Menu](#) on page 1916
- [Policies Menu](#) on page 1920
- [Devices Menu](#) on page 1922
- [Object Manager](#) on page 1922
- [Health Menu](#) on page 1923
- [System Menu](#) on page 1923
- [Help Menu](#) on page 1925

The selectable options under System Permissions allow you to create a user role that can make queries to the external database or escalate to the permissions of a target user role. For more information, see [Enabling Access to the Database](#) on page 2007 and [Managing User Role Escalation](#) on page 1925.

Optionally, instead of creating a new custom user role, you can export a custom user role from another appliance, then import it onto your appliance. You can then edit the imported role to suit your needs before you apply it. For more information, see [Exporting Objects](#) on page 2226 and [Importing Objects](#) on page 2231.

To create a custom user role:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Click the **User Roles** tab.
The User Roles page appears.

3. Click **Create User Role**.

The User Role Editor page appears.

The screenshot shows the 'User Role Editor' interface. At the top, there are two input fields: 'Name' and 'Description'. Below these is a section titled 'Menu Based Permissions' which contains a tree view of permissions. The tree view has a plus sign next to each item, indicating it can be expanded. The items are: Overview, Analysis, Policies, Object Manager, Devices, Health, System, and Help. Below this is a section titled 'System Permissions' which contains a checkbox for 'External Database Access'. At the bottom right of the form are two buttons: 'Save' and 'Cancel'.

4. In the **Name** field, type a name for the new user role.

You can use alphanumeric or hyphen characters, without spaces. Role names must be no more than 75 characters. User role names are case sensitive.

5. Optionally, add a description for the new role in the **Description** field.

Role descriptions must be no more than 255 characters.

6. Select permissions for the new role.

When you select an unselected permission, all of its children are selected, and the multi-value permissions choose the first value. If you deselect a high-level permission, all of its children are deselected also. Selected permissions without all children selected appear in italic text.

Note that choosing to copy a predefined user role to use as the base for your custom role preselects the permissions associated with that predefined role. For more information on copying predefined user roles, see [Creating a Custom Copy of a Predefined User Role](#) on page 1911.

The current escalation target role is listed beside the role escalation check box. If you select this check box, you can then choose to authenticate escalations either with the assigned user's password or with the password of another specified user role. For more information, see [Managing User Role Escalation](#) on page 1925.

7. Click **Save**.

The custom user role is created and the User Roles page appears again.

Creating a Custom Copy of a Predefined User Role

LICENSE: Any

You can copy an existing role to use as the basis for a new custom role. This preselects the existing role's permissions in the User Role Editor so you can model one role on another.

To create a custom copy of a predefined user role:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Click the **User Roles** tab.
The User Roles page appears.
3. Click the copy icon () next to the user role you want to copy.
The User Role Editor page appears with the copied role's permissions preselected.
Note that you can copy both custom and predefined user roles in this way.

Deleting a Custom User Role

LICENSE: Any

Unlike predefined user roles, you can delete custom roles that are no longer necessary. If you want to disable a custom role without removing it entirely, you can deactivate it instead; for more information, refer to the procedure in [Managing Predefined User Roles](#) on page 1906. Note that you cannot delete your own user role or a role that is set as a default user role in the system policy. For more information, see [Configuring Authentication Profiles](#) on page 1975.

To delete a custom user role:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Click the **User Roles** tab.
The User Roles page appears.
3. Click the delete icon () next to the custom role you want to delete.
The custom role is deleted.
If a deleted role is the only role assigned to a given user, that user can log in and access the User Preferences menu, but is otherwise unable to access the Sourcefire 3D System.

Modifying User Privileges and Options

LICENSE: Any

After adding user accounts to the system, you can modify access privileges, account options, or passwords at any time. Note that password management options do not apply to users who authenticate to an external directory server. You manage those settings on the external server. However, you must configure access rights for all accounts, including those that are externally authenticated.

For externally authenticated users, you cannot remove the minimum access rights through the Sourcefire 3D System user management page for users assigned an access role because of LDAP group or RADIUS list membership or attribute values. You can, however, assign additional rights. When you modify the access rights for an externally authenticated user, the Authentication Method column on the User Management page provides a status of **External - Locally Modified**.

Note that if you change the authentication for a user from externally authenticated to internally authenticated, you must supply a new password for the user.

To modify user account privileges:

ACCESS: Admin

1. Select **System > Local > User Management.**

The User Management page appears.

2. Click the edit icon () next to the user you want to modify.

The Edit User page appears.

3. Modify the account or accounts as needed:

- See [Managing Externally Authenticated User Accounts](#) on page 1902 for a description of how users can be authenticated through external servers.
- See [Managing User Login Settings](#) on page 1903 for information on changing password settings for internally authenticated users.
- See [Configuring User Roles](#) on page 1905 for more information on configuring roles to grant access for Sourcefire 3D System functions.

Understanding Restricted User Access Properties

LICENSE: Any

You can restrict the data that a user role can view in the event viewer by applying a restricted search to that role. You can specify this information when creating or editing the role assigned to a user. To create a custom role with restricted access, you must choose the tables you want to restrict from the Menu Based Permissions list, then select private saved searches from the Restrictive Search drop-down lists. For more information, see [Managing Custom User Roles](#) on page 1908.

Modifying User Passwords

LICENSE: Any

You can modify user passwords from the User Management page for internally authenticated users. Note that you must manage externally authenticated user passwords on the LDAP or RADIUS server.

To change a user's password:

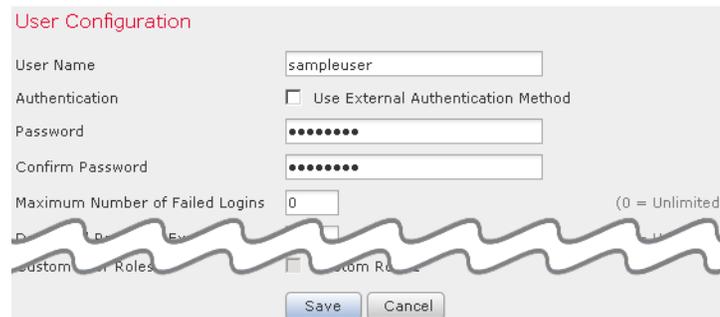
ACCESS: Admin

1. Select **System > Local > User Management**.

The User Management page appears.

2. Next to the user name, click the edit icon (✎).

The Edit User page appears.



3. In the **Password** field, type the new password (up to 32 alphanumeric characters).

4. In the **Confirm Password** field, retype the new password.

If password strength checking is enabled for the user account, the password must have at least eight alphanumeric characters of mixed case, with at least one number and one special character. It cannot be a word that appears in a dictionary or contain consecutive repeating characters.

5. Make any other changes you want to make to the user configuration:
 - For more information on password options, see [Managing User Login Settings](#) on page 1903.
 - For more information on user roles, see [Configuring User Roles](#) on page 1905.

6. Click **Save**.

The password is changed and any other changes saved.

Deleting User Accounts

LICENSE: Any

You can delete user accounts from the system at any time, with the exception of the admin account, which cannot be deleted.

To delete a user account:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Next to the user whose account you want to delete, click the delete icon (🗑️).
The account is deleted.

User Account Privileges

LICENSE: Any

The following sections provide a list of the configurable user permissions in the Sourcefire 3D System and the user roles that can access them. The permissions listed here follow the order of the Menu Based Permissions list that appears when you create a custom user role. Not all permissions are available on managed devices; permissions available only on the Defense Center are marked accordingly. For more information, see [Managing Custom User Roles](#) on page 1908.

Note that because the DC500 Defense Center and Series 2 devices support restricted features sets, not all permissions are applicable to these appliances. See the [Supported Capabilities by Appliance Model table](#) on page 42 for a summary of Series 2 appliance features.

For more information on the access notations used in the tables that follow and throughout this documentation, see [Access Conventions](#) on page 56. The following sections refer to the user role privileges associated with each main menu in the web-based interface:

- [Overview Menu](#) on page 1915
- [Analysis Menu](#) on page 1916
- [Policies Menu](#) on page 1920
- [Devices Menu](#) on page 1922
- [FireAMP](#) on page 1922
- [Devices Menu](#) on page 1922
- [Health Menu](#) on page 1923
- [System Menu](#) on page 1923
- [Help Menu](#) on page 1925

Overview Menu

LICENSE: Any

The [Overview Menu](#) table lists, in order, the user role privileges required to access each option in the Overview menu and the sub-permissions within. An X indicates that the user role has access. The Security Approver, Discovery Admin, Intrusion Admin, Access Admin, Network Admin, and External Database User roles have no permissions in the Overview menu.

Overview Menu

PERMISSION	ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
Dashboards	X	X	X	X
Manage Dashboards	X			
Appliance Information Widget	X	X	X	X
Appliance Status Widget (<i>Defense Center only</i>)	X	X	X	X
Correlation Events Widget	X		X	X
Current Interface Status Widget	X	X	X	X
Current Sessions Widget	X			
Custom Analysis Widget (<i>Defense Center only</i>)	X		X	X
Disk Usage Widget	X	X	X	X
Interface Traffic Widget	X	X	X	X
Intrusion Events Widget (<i>Defense Center only</i>)	X		X	X
Network Correlation Widget (<i>Defense Center only</i>)	X		X	X
Product Licensing Widget (<i>Defense Center only</i>)	X	X		
Product Updates Widget	X	X		
RSS Feed Widget	X	X	X	X
System Load Widget	X	X	X	X
System Time Widget	X	X	X	X

Overview Menu (Continued)

PERMISSION	ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
White List Events Widget <i>(Defense Center only)</i>	X		X	X
Reporting <i>(Defense Center only)</i>	X		X	X
Manage Report Templates <i>(Defense Center only)</i>	X		X	X
Summary	X		X	X
Intrusion Event Statistics <i>(Defense Center only)</i>	X		X	X
Intrusion Event Performance	X			
Intrusion Event Graphs <i>(Defense Center only)</i>	X		X	X
Discovery Statistics <i>(Defense Center only)</i>	X		X	X
Discovery Performance <i>(Defense Center only)</i>	X			
Connection Summary <i>(Defense Center only)</i>	X		X	X

Analysis Menu

LICENSE: Any

The [Analysis Menu](#) table lists, in order, the user role privileges required to access each option in the Analysis menu and the sub-permissions within. Permissions that appear multiple times under different headings will be listed on the table only where they first appear, except to indicate submenu headings. An X indicates that the user role has access. The Security Approver, Intrusion Admin, Access Admin, Network Admin, and External Database User roles have no permissions in the Analysis menu. The Analysis menu is only available on the Defense Center.

Analysis Menu

MENU	ADMIN	DISCOVERY ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
Application Statistics	X			X	X
Geolocation Statistics	X			X	X
User Statistics	X			X	X

Analysis Menu (Continued)

MENU	ADMIN	DISCOVERY ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
URL Category Statistics	X			X	X
URL Reputation Statistics	X			X	X
Intrusion Event Statistics by Application	X			X	X
Intrusion Event Statistics by User	X			X	X
Security Intelligence Category Statistics	X			X	X
Context Explorer	X			X	X
Connection Events	X			X	X
Modify Connection Events	X			X	
Connection Summary Events	X			X	X
Modify Connection Summary Events	X			X	
Intrusion	X			X	X
Intrusion Events	X			X	X
Modify Intrusion Events	X			X	
View Local Rules	X			X	X
Reviewed Events	X			X	X
Clipboard	X			X	X
Incidents	X			X	X
Malware Events	X			X	X
Modify Malware Events	X			X	
File Events	X			X	X
Modify File Events	X			X	

Analysis Menu (Continued)

MENU	ADMIN	DISCOVERY ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
Hosts	X			X	X
Network Map	X			X	X
Hosts	X			X	X
Modify Hosts	X			X	
Servers	X			X	X
Modify Servers	X			X	
Vulnerabilities	X			X	X
Host Attributes	X			X	X
Modify Host Attributes	X			X	
Applications	X			X	X
Application Details	X			X	X
Modify Application Details	X			X	
Host Attribute Management	X				
Discovery Events	X			X	X
Modify Discovery Events	X			X	
Users	X	X		X	X
User Activity	X	X		X	X
Modify User Activity Events	X	X		X	
Users	X	X		X	X
Modify Users	X	X		X	
Vulnerabilities	X			X	X

Analysis Menu (Continued)

MENU	ADMIN	DISCOVERY ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
Third-party Vulnerabilities	X			X	X
Modify Third-party Vulnerabilities	X			X	
Correlation	X	X		X	X
Correlation Events	X	X		X	X
Modify Correlation Events	X	X		X	
White List Events	X	X		X	X
Modify White List Events	X	X		X	
White List Violations	X	X		X	X
Remediation Status	X	X			
Modify Remediation Status	X	X			
Custom	X			X	X
Custom Workflows	X			X	X
Manage Custom Workflows	X			X	X
Custom Tables	X			X	X
Manage Custom Tables	X			X	X
Search	X		X	X	X
Manage Search	X				
Bookmarks	X			X	X
Manage Bookmarks	X			X	X

Policies Menu

LICENSE: Any

The [Policies Menu](#) table lists, in order, the user role privileges required to access each option in the Policies menu and the sub-permissions within. An X indicates that the user role has access. The External Database User, Maintenance User, Security Analyst, and Security Analyst (Read Only) roles have no permissions in the Policies menu. The Policies menu is only available on the Defense Center.

Policies Menu

MENU	ACCESS ADMIN	ADMIN	DISCOVERY ADMIN	INTRUSION ADMIN	NETWORK ADMIN	SECURITY APPROVER
Access Control	X	X			X	X
Access Control List	X	X			X	X
Modify Access Control Policy	X	X			X	
Modify Administrator Rules	X	X			X	
Modify Root Rules	X	X			X	
Apply Intrusion Policies		X				X
Apply Access Control Policies		X				X
Intrusion	X	X		X		X
Intrusion Policy		X		X		X
Modify Intrusion Policy		X		X		
Rule Editor		X		X		
Email		X		X		
File Policy	X	X				
Modify File Policy	X	X				
Network Discovery		X	X			X
Modify Network Discovery		X	X			
Apply Network Discovery		X				X

Policies Menu (Continued)

MENU	ACCESS ADMIN	ADMIN	DISCOVERY ADMIN	INTRUSION ADMIN	NETWORK ADMIN	SECURITY APPROVER
Custom Fingerprinting		X	X			
Custom Product Mappings		X	X			
User 3rd Party Mappings		X	X			
Custom Topology		X	X			
Application Detectors		X	X			
Users		X				
Correlation		X				
Policy Management		X				
Rule Management		X				
White List		X				
Traffic Profiles		X				
Actions		X	X			
Alerts		X	X			
Impact Flag Alerts		X	X			
Discovery Event Alerts		X	X			
Scanners		X	X			
Scan Results		X	X			
Modify Scan Results		X	X			
Groups		X				
Modules		X				
Instances		X				

Devices Menu

LICENSE: Any

The **Devices** menu table lists, in order, the user role privileges required to access each option in the Devices menu and the sub-permissions within. An X indicates that the user role has access. The Access Admin, Discovery Admin, External Database User, Intrusion Admin, Maintenance User, Security Approver, Security Analyst, and Security Analyst (Read Only) have no permissions in the Devices menu. The Devices menu is only available on the Defense Center.

Devices Menu

MENU	ADMIN	NETWORK ADMIN
Device Management	X	X
Modify Devices	X	X
Apply Device Changes	X	X
NAT	X	X
NAT List	X	X
Modify NAT Policy	X	X
Apply NAT Rules	X	
VPN	X	X
Modify VPN	X	X
Apply VPN Changes	X	X

Object Manager

LICENSE: Any

The Object Manager permission is available to the Access Admin, Administrator, and Network Admin user roles. The Object Manager permission is only available on the Defense Center.

FireAMP

LICENSE: Any

The FireAMP permission is available only to the Administrator user role. This permission is only available on the Defense Center.

Health Menu

LICENSE: Any

The [Health Menu](#) table lists, in order, the user role privileges required to access each option in the Health menu and the sub-permissions within. An X indicates that the user role has access. The Access Admin, Discovery Admin, Intrusion Admin, External Database User, Network Admin, and Security Approver roles have no permissions in the Health menu. The Health menu is only available on the Defense Center.

Health Menu

MENU	ADMIN	MAINT USER	SECURITY ANALYST	SECURITY ANALYST (RO)
Health Policy	X	X		
Modify Health Policy	X	X		
Apply Health Policy	X	X		
Health Events	X	X	X	X
Modify Health Events	X	X		

System Menu

LICENSE: Any

The [System Menu](#) table lists, in order, the user role privileges required to access each option in the System menu and the sub-permissions within. An X indicates that the user role has access. The Access Admin, Discovery Admin, Intrusion Admin, External Database User, and Security Analyst (Read Only) roles have no permissions in the System Menu.

System Menu

MENU	ADMIN	MAINT USER	NETWORK ADMIN	SECURITY APPROVER	SECURITY ANALYST
Local	X				
Configuration	X				
Registration	X				
High Availability (<i>DC1000, DC1500, DC3000, DC3500 only</i>)	X				

System Menu (Continued)

MENU	ADMIN	MAINT USER	NETWORK ADMIN	SECURITY APPROVER	SECURITY ANALYST
eStreamer	X				
Host Input Client <i>(Defense Center only)</i>	X				
User Management	X				
Login Authentication <i>(Defense Center only)</i>	X				
System Policy <i>(Defense Center only)</i>	X				
Modify System Policy <i>(Defense Center only)</i>	X				
Apply System Policy <i>(Defense Center only)</i>	X				
Updates	X				
Rule Updates <i>(Defense Center only)</i>	X				
Rule Update Import Log <i>(Defense Center only)</i>	X				
Licenses	X				
Monitoring	X	X	X	X	X
Audit	X				
Modify Audit Log	X				
Syslog	X	X			
Task Status	X	X	X	X	X
View Other Users' Tasks	X				
Statistics	X	X			
Tools	X	X			X
Backup Management	X	X			
Restore Backup	X	X			
Scheduling	X	X			

System Menu (Continued)

MENU	ADMIN	MAINT USER	NETWORK ADMIN	SECURITY APPROVER	SECURITY ANALYST
Delete Other Users' Scheduled Tasks	X				
Import/Export	X				
Discovery Data Purge <i>(Defense Center only)</i>	X				X
Whois	X	X			X

Help Menu

LICENSE: Any

The Help menu and its permissions are accessible to all user roles. You cannot restrict Help menu options.

Managing User Role Escalation

LICENSE: Any

You can give custom user roles the permission, with a password, to temporarily gain the privileges of another, targeted user role in addition to those of the base role. This allows you to easily substitute one user for another during an absence, or to more closely track the use of advanced user privileges.

For example, a user whose base role has very limited privileges may escalate to the Administrator role to perform administrative actions. You can configure this feature so that users can use their own passwords, or so they use the password of another user that you specify. The second option allows you to easily manage one escalation password for all applicable users. For more information, see [Configuring a Custom User Role for Escalation](#) on page 1926.

Note that only one user role at a time can be the escalation target role. You can use a custom or predefined user role. Each escalation lasts for the duration of a login session and is recorded in the audit log.

For more information on configuring and using this feature, please see the following sections:

- [Configuring the Escalation Target Role](#) on page 1926
- [Configuring a Custom User Role for Escalation](#) on page 1926
- [Escalating Your User Role](#) on page 1928

Configuring the Escalation Target Role

LICENSE: Any

You can assign any of your user roles, predefined or custom, to act as the system-wide escalation target role. This is the role to which any other role may escalate, if it has the ability.

To configure the escalation target role:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Click **User Roles**.
The User Roles page appears.
3. Click **Configure Permission Escalation**.
The Configure Permission Escalation dialog box appears.



4. Select a user role from the drop-down list.
5. Click **OK** to save your changes.
Your changes are saved and the User Roles page appears.

IMPORTANT! Changing the escalation target role is effective immediately. Users in escalated sessions now have the permissions of the new escalation target.

Configuring a Custom User Role for Escalation

LICENSE: Any

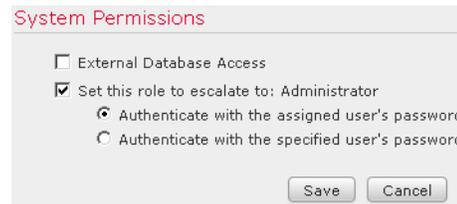
To use the user role escalation feature, you must first configure a custom user role with the escalation permission, select its escalation password, and assign that role to a user. For more information, see [Adding New User Accounts](#) on page 1899 and [Configuring User Roles](#) on page 1905.

Consider the needs of your organization when you configure the escalation password for a custom role. If you want to easily manage many escalating users, you may want to select another user whose password serves as the escalation password. If you change that user's password or deactivate that user, all escalating users who require that password are affected. This allows you to manage user role escalation more efficiently, especially if you select an externally authenticated user that you can manage centrally.

To configure a custom user role for escalation:

ACCESS: Admin

1. Select **System > Local > User Management**.
The User Management page appears.
2. Click **User Roles**.
The User Roles page appears.
3. Click **Create User Role** to create a new custom user role, or the edit icon (✎) next to an existing custom user role.
The User Role Editor page appears.
4. Choose a name, description and menu-based permissions for the custom user role.
For more information, see the procedure in [Managing Custom User Roles](#) on page 1908.
5. In System Permissions, select the **Set this role to escalate to:** check box.
The escalation password options appear.



System Permissions

External Database Access

Set this role to escalate to: Administrator

- Authenticate with the assigned user's password
- Authenticate with the specified user's password

Save Cancel

6. Select the password that this role uses to escalate. You have two options:
 - If you want users with this role to use their own passwords when they escalate, select **Authenticate with the assigned user's password**.
 - If you want users with this role to use the password of another user, select **Authenticate with the specified user's password** and type that username.

IMPORTANT! When authenticating with another user's password, you can enter any username, even that of a deactivated or nonexistent user. Deactivating the user whose password is used for escalation makes escalation impossible for users with the role that requires it. You can use this feature to quickly remove escalation powers if necessary.

7. Click **Save**.
Your changes are saved and the User Roles page appears again. Users with this role can now escalate to the target user role. For more information on assigning roles to a user, see [Adding New User Accounts](#) on page 1899.

Escalating Your User Role

LICENSE: Any

When a user has an assigned custom user role with permission to escalate, that user may escalate to the target role's permissions at any time. Note that escalation has no effect on user preferences. The **Escalate Permissions** option in the User menu does not appear if your assigned user role is not configured for user role escalation.

To escalate user permissions:

ACCESS: Any

1. Select **Local > User > Escalate Permissions**.

The Escalate User Permissions dialog box appears.



2. Enter the authentication password.
3. Click **Escalate**.

You now have all permissions of the escalation target role in addition to your current role.

Note that escalation lasts for the remainder of your login session. To return to the privileges of your base role only, you must log out, then begin a new session.

CHAPTER 47

SCHEDULING TASKS

You can schedule many different types of administrative tasks to run at designated times, either once or on a recurring basis.

IMPORTANT! Some tasks (such as those involving automated software updates or that require pushing updates to managed devices) 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](#) on page 1930 explains how to set up a scheduled task so that it runs at regular intervals.
- [Automating Backup Jobs](#) on page 1932 provides procedures for scheduling backup jobs.
- [Automating Certificate Revocation List Downloads](#) on page 1934 provides procedures for automatically refreshing the certificate revocation list (CRL) for an appliance.
- [Automating Nmap Scans](#) on page 1936 provides procedures for scheduling Nmap scans.
- [Automating Applying an Intrusion Policy](#) on page 1938 provides procedures for queuing an intrusion policy apply on managed devices.
- [Automating Reports](#) on page 1940 provides procedures for scheduling reports.

- [Automating Geolocation Database Updates](#) on page 1942 provides procedures for scheduling automatic updates of the geolocation database (GeoDB).
- [Automating FireSIGHT Recommendations](#) on page 1943 provides procedures for scheduling the automatic update of intrusion rule state recommendations.
- [Automating Software Updates](#) on page 1945 provides procedures for scheduling the download, push, and installation of software updates.
- [Automating Vulnerability Database Updates](#) on page 1951 provides procedures for scheduling the download and installation of VDB updates.
- [Automating URL Filtering Updates](#) on page 1955 provides procedures for automating updates of URL filtering data.
- [Viewing Tasks](#) on page 1957 describes how to view and manage tasks after they are scheduled.
- [Editing Scheduled Tasks](#) on page 1959 describes how to edit an existing task.
- [Deleting Scheduled Tasks](#) on page 1959 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.

IMPORTANT! You cannot configure a recurring task schedule on the inactive Defense Center in a high availability pair of Defense Centers. You must recreate the recurring task schedule on a newly activated Defense Center when it changes from inactive to active.

Note that the time displayed on most pages on the web interface is the local time, which is determined by using the time zone you specify in your local configuration. Further, the Defense Center 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:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
2. Click **Add Task**.
The New Task page appears.
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.
4. For the **Schedule task to run** option, select **Recurring**.
The page reloads with the recurring task options.

New Task

Job Type: Backup

Schedule task to run: Once Recurring

Start On: January 3, 2012 America/New York

Repeat Every: 1 Hours Days Weeks Months

Run At: 2:00 Pm

Repeat On: Sunday Monday Tuesday Wednesday Thursday Friday Saturday

Job Name:

Backup Profile:

Comment:

Email Status To:

Save Cancel

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.
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.

7. In the **Run At** field, specify the time when you want to start your recurring task.
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.
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](#) on page 1932
 - [Automating Certificate Revocation List Downloads](#) on page 1934
 - [Automating Nmap Scans](#) on page 1936
 - [Automating Reports](#) on page 1940
 - [Automating FireSIGHT Recommendations](#) on page 1943
 - [Automating Software Updates](#) on page 1945
 - [Automating Vulnerability Database Updates](#) on page 1951
 - [Automating URL Filtering Updates](#) on page 1955

Automating Backup Jobs

LICENSE: Any

You can use the scheduler to automate system backups of a Defense Center or a managed device.

TIP! You must design a backup profile before you can configure it as a scheduled task. For information on backup profiles, see [Creating Backup Profiles](#) on page 2208.

To automate backup tasks:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
2. Click **Add Task**.
The New Task page appears.

3. From the **Job Type** list, select **Backup**.

The page reloads to show the backup options. The Defense Center version of the page is shown below.

The screenshot shows a 'New Task' configuration form. At the top, the title 'New Task' is in red. Below it, the 'Job Type' dropdown is set to 'Backup'. The 'Schedule task to run' section has radio buttons for 'Once' (selected) and 'Recurring'. The 'Current time' is displayed as '2011-11-03 18:24'. The 'Start Time' is configured with dropdowns for 'November', '3', '2011', '1:00', 'Pm', and a time zone of 'America/New York'. There is an empty text field for 'Job Name'. The 'Backup Profile' dropdown is set to 'Test backup'. Below that is a large empty text area for 'Comment'. At the bottom, there is an 'Email Status To' text field and two buttons: 'Save' and 'Cancel'.

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](#) on page 1930 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. From the **Backup Profile** list, select the appropriate backup profile.
For more information on creating new backup profiles, see [Creating Backup Profiles](#) on page 2208.
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.

8. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want status messages sent.

You must have a valid email relay server configured on the Defense Center to send status messages. See [Configuring a Mail Relay Host and Notification Address](#) on page 1983 for more information about configuring a relay host.

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](#) on page 2238.

Automating Certificate Revocation List Downloads

LICENSE: Any

You can use the scheduler to automatically refresh the certificate revocation list (CRL) for the appliance web server on an appliance where you enable user certificates for the appliance. The Download CRL task is automatically created when you enable fetching of a CRL in the local appliance configuration, so this procedure explains how to open the scheduled task to set the frequency.

TIP! You must enable and configure user certificates and set a CRL download URL before scheduling this task. For information on configuring user certificates, see [Configuring User Certificates](#) on page 2006.

To automate download of certificate revocation lists:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

2. Locate the **download CRL** task in the Task Details and click the edit icon (✎). The Edit Task page appears, showing the download options. The Defense Center version of the page is shown below.

Edit Task

Job Type: Download CRL

Schedule task to run: Once Recurring

Current time: 2011-11-15 11:51

Start Time: November 15, 2011 4:30 Pm America/New York

Job Name: Test Download CRL

Comment:

Email Status To:

Save Cancel

3. Specify how you want to schedule the CRL download, **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](#) on page 1930 for details.
4. 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.

5. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want status messages sent.

You must have a valid email relay server configured on the Defense Center to send status messages. See [Configuring a Mail Relay Host and Notification Address](#) on page 1983 for more information about configuring a relay host.
6. 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](#) on page 2238.

Automating Nmap Scans

LICENSE: FireSIGHT

You can schedule regular Nmap scans of targets on your network. Automated scans allow you to refresh information previously supplied by an Nmap scan. Because the Sourcefire 3D System cannot update Nmap-supplied data, you need to rescan periodically to keep that data up to date. You can also schedule scans to automatically test for unidentified applications or servers on hosts in your network. See the following sections for more information:

- [Preparing Your System for an Nmap Scan](#)
- [Scheduling an Nmap Scan](#)

Note that a Discovery Administrator can also use an Nmap scan as a remediation. For example, when an operating system conflict occurs on a host, that conflict may trigger an Nmap scan. Running the scan obtains updated operating system information for the host, which resolves the conflict. For more information, see [Nmap Scan Remediations](#) on page 1625.

Preparing Your System for an Nmap Scan

LICENSE: FireSIGHT

If you have not used the Nmap scanning capability before, you must complete several Nmap configuration steps before defining a scheduled scan. See the following sections for more information:

- [Creating an Nmap Scan Instance](#) on page 1702 provides information on setting up an Nmap server connection profile.
- [Creating an Nmap Scan Target](#) on page 1704 provides information on setting up a scan target.
- [Creating an Nmap Remediation](#) on page 1705 provides information on setting up a remediation definition.

Scheduling an Nmap Scan

LICENSE: FireSIGHT

You can schedule a scan of a host or hosts on your network using the Nmap utility.

After Nmap replaces a host's operating system, applications, or servers detected by the system with the results from an Nmap scan, the system no longer updates the information replaced by Nmap for the host. Nmap-supplied service and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep Nmap-supplied operating systems, applications, or servers up to date. If the host is deleted from the network map and re-added, any Nmap scan results are discarded and the system resumes monitoring of all operating system and service data for the host.

To automate Nmap scanning:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

2. Click **Add Task**.

The New Task page appears.

3. From the **Job Type** list, select **Nmap Scan**.

The page reloads to show the options for automating Nmap scans.

The screenshot shows the 'New Task' configuration form. The 'Job Type' dropdown is set to 'Nmap Scan'. Under 'Schedule task to run', the 'Once' radio button is selected. The 'Current time' is '2011-11-03 18:29'. The 'Start Time' is configured with a month of 'November', a day of '3', a year of '2011', a time of '1:00', and a period of 'Pm', with the time zone set to 'America/New York'. The 'Job Name' field is empty. The 'Nmap Remediation' dropdown is set to 'SampleNmapRemediation' and the 'Nmap Target' dropdown is set to 'SampleNmapScanTarget'. There is a large empty text area for 'Comment' and an empty text field for 'Email Status To'. At the bottom, there are 'Save' and 'Cancel' buttons.

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](#) on page 1930 for details.

5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. In the **Nmap Remediation** field, select the Nmap remediation to use when running the scan.
7. In the **Nmap Target** field, select the scan target that defines the target hosts you want to scan.

8. 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.

9. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

10. 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](#) on page 2238.

Automating Applying an Intrusion Policy

LICENSE: Protection

You can queue an intrusion policy apply to a managed device. This task only applies the intrusion policy if an access control policy that references the intrusion policy is applied to the selected device 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 the following sections for more information:

- [Using Default Intrusion Policies](#) on page 692
- [Creating an Intrusion Policy](#) on page 672
- [Reapplying an Intrusion Policy](#) on page 680
- [Applying an Access Control Policy](#) on page 466

To queue a policy apply to a managed device:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The schedule calendar page for the current month appears.

2. Click **Add Task**.

The New Task page appears.

- From the **Job Type** list, select **Queue Intrusion Policy Apply**.
The page reloads to show the options for queuing a policy apply.

The screenshot shows a 'New Task' configuration form. At the top, the title is 'New Task'. Below it, the 'Job Type' is set to 'Queue Intrusion Policy Apply'. The 'Schedule task to run' section has 'Once' selected with a radio button, and 'Recurring' is unselected. The 'Current time' is '2012-01-05 13:31'. The 'Start Time' is configured with a date picker showing 'January 5, 2012' and a time picker showing '2:00 Pm' in the 'America/New York' time zone. There are empty text input fields for 'Job Name', 'Intrusion Policy' (set to 'All intrusion policies'), and 'Device' (set to 'All targeted devices'). A large text area is provided for 'Comment'. At the bottom, the 'Email Status To' field contains the message 'Not available. You must set up your mail relay host.' and there are 'Save' and 'Cancel' buttons.

- 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 Defense Center.
 - For recurring tasks, you have several options for setting the interval between instances of the task. See [Configuring a Recurring Task](#) on page 1930 for details.
- In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
- In the **Intrusion Policy** field, you have the following options:
 - Select an intrusion policy to apply to the selected target device.
 - Select **All intrusion policies** to apply all intrusion policies already applied to the device selected in the **Device** field.
- In the **Device** field, you have the following options:
 - Select a device to which you want to apply the intrusion policy selected in the **Intrusion Policy** field.
 - Select **All targeted devices** to apply the selected intrusion policy to all monitored devices which already have that intrusion policy applied.

TIP! This field only displays devices which have the intrusion policy selected in the **Intrusion Policy** field already applied.

8. 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.

9. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

10. 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](#) on page 2238.

11. 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 Reports

LICENSE: Any

You can automate reports so that they run at regular intervals. However, you must design a template for your report before you can configure it as a scheduled task. See [Understanding Report Templates](#) on page 1733 for more information about using the report designer to create a report template.

To automate report generation:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The schedule calendar page for the current month appears.

2. Click **Add Task**.

The New Task page appears.

- From the **Job Type** list, select **Report**.

The page reloads to show the options for setting up a report to run automatically.

New Task

Job Type: Report

Schedule task to run: Once Recurring

Current time: 2013-05-10 13:05

Start Time: May 10 2013 2:00 Pm America/New York

Job Name:

Report Template: Host Report: Edit

Comment:

Email Status To:

If report is empty, still attach to email:

Save Cancel

- 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 Defense Center.
 - For recurring tasks, you have several options for setting the interval between instances of the task. See [Configuring a Recurring Task](#) on page 1930 for details.
- In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
- In the **Report Template** field, select the report template that you want to use from the drop-down list.
- 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.

- Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

9. If you do not want to receive report email attachments when reports have no data (for example, when no events of a certain type occurred during the report period), select the **If report is empty, still attach to email** check box.
10. 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](#) on page 2238.
11. 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: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

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](#) on page 2092.

To automate geolocation database updates:

ACCESS: Admin

1. Select **System > Updates**.
The Product Updates page appears.
2. Click the **Geolocation Updates** tab.
The Geolocation Updates page appears.

Product Updates Rule Updates **Geolocation Updates**

Defense Center running geolocation update version: 2013-03-12-001

One-Time Geolocation Update

Note that updates may be large and can take up to 45 minutes.

Source

Upload and install geolocation update

Download and install geolocation update from the Support Site

Browse...

Import

Recurring Geolocation Updates

Enable Recurring Weekly Updates

Update Start Time

Monday 01:00 Pm America/New York

Save Cancel

3. Under **Recurring Geolocation Updates**, select the **Enable Recurring Weekly Updates** check box.
The Update Start Time field appears.
4. In the **Update Start Time** field, specify the time and day of the week when you want weekly GeoDB updates to occur.
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](#) on page 2238.

Automating FireSIGHT Recommendations

LICENSE: Protection

You can automatically generate rule state recommendations based on network discovery data for your network using the most recently saved configuration settings in your custom intrusion policy.

IMPORTANT! If the system automatically generates scheduled recommendations for an intrusion policy with unsaved changes, you must discard your changes in that policy and commit the policy if you want the policy to reflect the automatically generated recommendations. See [Committing Intrusion Policy Changes](#) on page 679 for more information.

When the task runs, the system automatically generates recommended rule states. Optionally, depending on the configuration of your policy, it also modifies the states of intrusion rules based on the criteria described in [Managing FireSIGHT Rule State Recommendations](#) on page 764. Modified rule states take effect the next time you apply your intrusion policy. See [Using FireSIGHT Recommendations](#) on page 768 for more information.

To automate rule state recommendation generation:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
2. Click **Add Task**.
The New Task page appears.

- From the **Job Type** list, select **FireSIGHT Recommended Rules**.
The page reloads to show the options for generating FireSIGHT recommendations.

New Task

Job Type: FireSIGHT Recommended Rules (FireSIGHT Recommended Rules must first be configured in the selected policies)

Schedule task to run: Once Recurring

Current time: 2011-11-03 18:36

Start Time: November 3, 2011, 1:00 Pm America/New York

Job Name: [Empty text box]

Policies: All Policies
Initial Passive Policy - katsura
Initial Inline Policy - katsura

Comment: [Empty text area]

Email Status To: [Empty text box]

Save Cancel

- Optionally, click the **policies** link next to the **Job Type** field to display the Detection & Prevention page, where you can configure FireSIGHT Recommended Rules in a policy. See [Managing FireSIGHT Rule State Recommendations](#) on page 764 for more information.
- 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](#) on page 1930 for details.
- In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
- Next to **Policies**, select one or more policies where you want to generate recommendations. You have the following options:
 - In the **Policies** field, select one or more policies. Use the Shift and Ctrl keys to select multiple policies.
 - Click the **All Policies** check box to select all policies.

8. 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.

9. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

10. 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](#) on page 2238.

Automating Software Updates

LICENSE: Any

You can automatically download and apply most patches and feature releases to the Sourcefire 3D System.

IMPORTANT! You must manually upload and install updates in two situations. First, you cannot schedule major updates to the Sourcefire 3D System. Second, you cannot schedule updates for or pushes from appliances that cannot access the [Sourcefire Support Site](#). If your appliance is not directly connected to the Internet, you should set up a proxy as described in [Configuring Network Settings](#) on page 2009 to allow it to download updates from the Support Site. For information on manually updating the Sourcefire 3D System, see [Updating System Software](#) on page 2054.

The tasks you must schedule to install software updates vary depending on whether you are updating the Defense Center or are using a Defense Center to update managed devices. Sourcefire **strongly** recommends that you use your Defense Centers to update the devices they manage.

To update the Defense Center, schedule the software installation using the Install Latest Update task. To use a Defense Center to automate software updates for its managed devices, you must schedule two tasks:

1. Push (copy) the update to managed devices using the Push Latest Update task.
2. Install the update on managed devices using the Install Latest Update task.

When scheduling updates, schedule the push and install tasks to happen in succession. That is, to automate software updates on your managed devices, you must first push the update to the device before you can install it. (Note that during the manual update process you do not have to push an update to managed devices before you install it. For more information, see [Updating Managed Devices](#) on page 2064.)

IMPORTANT! You cannot create individual update tasks for managed devices in a clustered or stacked configuration.

Always allow enough time between tasks for the process to complete. Tasks should be scheduled at least 30 minutes apart. For example, if you schedule a task to install an update and the update has not finished copying from the Defense Center to the device, the installation task will not succeed. However, if the scheduled installation task repeats daily, it will install the pushed update when it runs the next day.

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](#) on page 1946
- [Automating Software Pushes](#) on page 1948
- [Automating Software Installs](#) on page 1949

Automating Software Downloads

LICENSE: Any

You can create a scheduled task that automatically downloads the latest software updates from Sourcefire. You can use this task to schedule download of updates you plan to install manually.

To automate software update downloads:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
2. Click **Add Task**.
The New Task page appears.

- From the **Job Type** list, select **Download Latest Update**.
The New Task page reloads to show the update options.

New Task

Job Type:

Schedule task to run: Once Recurring

Current time: 2011-11-03 18:38

Start Time:

Job Name:

Update Items: Software Vulnerability Database

Comment:

Email Status To:

- 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](#) on page 1930 for details.
- In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
- In the **Update Items** section, select **Software**.
- 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.

- Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

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](#) on page 2238.

Automating Software Pushes

LICENSE: Any

If you want to automate the installation of software updates on managed devices, you must push the updates to the devices before installing.

When you push updates to managed devices, information about the push process status is reported on the Tasks page. See [Viewing the Status of Long-Running Tasks](#) on page 2238 for more information.

When you create the task to push software updates to managed devices, make sure you allow enough time between the push task and a scheduled install task for the updates to be copied to the device.

To push software updates to managed devices:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

2. Click **Add Task**.

The New Task page appears.

3. From the **Job Type** list, select **Push Latest Update**.

The page reloads to show the options for pushing updates.

The screenshot shows the 'New Task' configuration page. The 'Job Type' dropdown is set to 'Push Latest Update'. The 'Schedule task to run' section has 'Once' selected. The 'Current time' is 2011-12-13 16:38. The 'Start Time' is set to December 13, 2011, at 5:00 PM in the America/New York time zone. The 'Job Name' field is empty. The 'Device' dropdown is set to 'linden'. There is a large empty text area for 'Comment'. The 'Email Status To' field is empty. At the bottom are 'Save' and 'Cancel' buttons.

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](#) on page 1930 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. From the **Device** list, select the device that you want to receive updates.
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.

8. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.
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](#) on page 2238.

Automating Software Installs

LICENSE: Any

If you are using a Defense Center to create a task to install a software update on a managed device, make sure you allow enough time between the task that pushes the update to the device and the task that installs the update. See [Automating Software Pushes](#) on page 1948 for information about pushing updates to managed devices.

WARNING! Depending on the update being installed, the appliance may reboot after the software is installed.

To schedule a software installation task:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

2. Click **Add Task**.

The New Task page appears.

3. From the **Job Type** list, select **Install Latest Update**.

The page reloads to show the options for installing updates.

The screenshot shows the 'New Task' configuration interface. It includes a 'Job Type' dropdown menu with 'Install Latest Update' selected. Below it, there are radio buttons for 'Once' (selected) and 'Recurring'. The 'Current time' is displayed as '2011-12-13 16:41'. The 'Start Time' is configured with dropdowns for 'December', '13', '2011', '5:00', 'Pm', and 'America/New York'. There is an empty 'Job Name' text box, a 'Device' dropdown menu with 'linden' selected, and radio buttons for 'Software' (selected) and 'Vulnerability Database'. A large empty text area is provided for a 'Comment'. At the bottom, there is an 'Email Status To' text box and 'Save' and 'Cancel' buttons.

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](#) on page 1930 for details.

5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

6. From the **Device** list, you have the following options:

- Select the device where you want to install the update.
- Select the name of the Defense Center to install the update there.

7. In the **Update Items** section, select **Software**.

8. 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.

9. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

10. 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](#) on page 2238.

Automating Vulnerability Database Updates

LICENSE: FireSIGHT

Sourcefire uses vulnerability database (VDB) updates to expand the list of network assets, traffic, and vulnerabilities that the Sourcefire 3D System recognizes. You can use the scheduling feature to download and install the latest VDB update on your Defense Centers, thereby ensuring that you are using the most up-to-date information to evaluate the hosts on your network.

IMPORTANT! You cannot schedule updates for appliances that cannot access the [Sourcefire Support Site](#). If your appliance is not directly connected to the Internet, you should set up a proxy as described in [Configuring Network Settings](#) on page 2009 to allow it to download updates from the Support Site. For information on manually updating the Sourcefire 3D System, see [Updating System Software](#) on page 2054.

When automating VDB updates, you must automate two separate steps:

1. Downloading the VDB update.
2. Installing the VDB update.

Always allow enough time between tasks for the process to complete. For example, if you schedule a task to install an update and the update has not fully downloaded, the installation task will not succeed. However, if the scheduled installation task repeats daily, it will install the downloaded VDB update when the task runs the next day.

If you want to have more control over this process, you can use the **Once** option to download and install VDB updates during off-peak hours after you learn that an update has been released.

IMPORTANT! Installing a VDB update causes a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected.

See the following sections for more information:

- [Automating VDB Update Downloads](#) on page 1952
- [Automating VDB Update Installs](#) on page 1953

Automating VDB Update Downloads

LICENSE: FireSIGHT

You can create a scheduled task on the Defense Center that automatically downloads the latest VDB update from Sourcefire.

To automate VDB update downloads:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
2. Click **Add Task**.
The New Task page appears.
3. From the **Job Type** list, select **Download Latest Update**.
The New Task page reloads to show the update options.

New Task

Job Type:

Schedule task to run: Once Recurring

Current time: 2011-11-03 18:38

Start Time:

Job Name:

Update Items: Software Vulnerability Database

Comment:

Email Status To:

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](#) on page 1930 for details.

5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. In the **Update Items** section, select **Vulnerability Database**.
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.

8. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.
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](#) on page 2238.

Automating VDB Update Installs

LICENSE: FireSIGHT

You should allow enough time between the task that downloads the VDB update and the task that installs the update; see [Automating VDB Update Downloads](#) on page 1952 for information.

IMPORTANT! Installing a VDB update causes a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected.

To schedule a VDB update:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
2. Click **Add Task**.
The New Task page appears.

3. From the **Job Type** list, select **Install Latest Update**.

The page reloads to show the options for installing updates.

The screenshot shows a 'New Task' configuration form. At the top, the title 'New Task' is in red. Below it, the 'Job Type' dropdown is set to 'Install Latest Update'. The 'Schedule task to run' section has 'Once' selected with a radio button, and 'Recurring' is unselected. The 'Current time' is displayed as '2011-12-13 16:41'. The 'Start Time' section includes dropdowns for 'December', '13', '2011', '5:00', 'Pm', and 'America/New York'. The 'Job Name' field is empty. The 'Device' dropdown is set to 'linden'. The 'Update Items' section has 'Software' selected and 'Vulnerability Database' unselected. There is a large text area for 'Comment' and an 'Email Status To' field. At the bottom are 'Save' and 'Cancel' buttons.

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](#) on page 1930 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. From the **Device** drop-down list, select the name of the Defense Center.
7. In the **Update Items** section, select **Vulnerability Database**.
8. 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.

9. Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.

10. 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](#) on page 2238.

Automating URL Filtering Updates

LICENSE: URL Filtering

SUPPORTED DEFENSE CENTERS: Any except DC500

You can use the scheduler to automate updates of URL filtering data from the Sourcefire cloud. For a URL filtering update task to succeed:

- The Defense Center must have access to the internet or it cannot contact the cloud.
- You must enable URL filtering, as described in [Enabling Sourcefire Cloud Communications](#) on page 2032.

Note that when you enable URL filtering, you can also enable automatic updates. This forces the Defense Center 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:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

2. Click **Add Task**.

The New Task page appears.

- From the **Job Type** list, select **Update URL Filtering Database**.
The page reloads to show the URL filtering update options.

The screenshot shows a 'New Task' form with the following fields and options:

- Job Type:** A dropdown menu set to 'Update URL Filtering Database'.
- Schedule task to run:** Radio buttons for 'Once' (selected) and 'Recurring'.
- Current time:** A text field showing '2011-11-03 18:49'.
- Start Time:** A series of dropdown menus for month ('November'), day ('3'), year ('2011'), time ('1:00'), and period ('Pm'), followed by a text field for the time zone ('America/New York').
- Job Name:** A text input field.
- Comment:** A large text area.
- Email Status To:** A text input field.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom.

- 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](#) on page 1930 for details.
- In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
- 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.

- Optionally, in the **Email Status To** field, type the email address (or multiple email addresses separated by commas) where you want 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](#) on page 1983 for more information about configuring a relay host.
- 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](#) on page 2238.

Viewing Tasks

LICENSE: Any

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](#) on page 1957
- [Using the Task List](#) on page 1958

Using the Calendar

LICENSE: Any

The Calendar view option allows you to view which scheduled tasks occur on which day.

To view scheduled tasks using the calendar:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

2011 / 5						
Sun.	Mon.	Tues.	Wed.	Thurs.	Fri.	Sat.
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28

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.

IMPORTANT! For more information about using the task list, see [Using the Task List](#).

Using the Task List

LICENSE: Any

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](#) on page 1957 for more information.

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: <ul style="list-style-type: none">• 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

LICENSE: Any

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:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.

The Scheduling page appears.

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.

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.

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](#) on page 1932
- [Automating Certificate Revocation List Downloads](#) on page 1934
- [Automating Nmap Scans](#) on page 1936
- [Automating Reports](#) on page 1940
- [Automating FireSIGHT Recommendations](#) on page 1943
- [Automating Software Updates](#) on page 1945
- [Automating Vulnerability Database Updates](#) on page 1951
- [Automating URL Filtering Updates](#) on page 1955

5. Click **Save** to save your edits.

Your change are saved and the Scheduling page appears again.

Deleting Scheduled Tasks

LICENSE: Any

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](#) on page 1960.
- To delete a single instance of a task, see [Deleting a One-Time Task](#) on page 1960.

Deleting a Recurring Task

LICENSE: Any

When you delete one instance of a recurring task, you automatically delete all instances of that task.

To delete a recurring task:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
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.
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

LICENSE: Any

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:

ACCESS: Admin/Maint

1. Select **System > Tools > Scheduling**.
The Scheduling page appears.
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.
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.

CHAPTER 48

MANAGING SYSTEM POLICIES

A system policy allows you to manage the following on your Sourcefire 3D System appliances:

- access control preferences
- appliance access lists
- audit log settings
- authentication profiles
- dashboard settings
- database event limits
- DNS cache properties
- the mail relay host and notification address
- tracking intrusion policy changes
- specifying a different language
- custom login banners
- SNMP polling settings
- synchronizing time
- serving time from the Defense Center
- user interface and command line interface timeout settings
- mapping vulnerabilities for servers

You can use a system policy to control the aspects of your Defense Center that are likely to be similar for other appliances in your deployment. For example, your organization's security policies may require that your appliances have a "No Unauthorized Use" message when a user logs in. With system policies, you can

set the login banner once in a system policy on a Defense Center and then apply the policy to all the devices that it manages.

You can also benefit from having multiple system policies on a Defense Center. For example, if you have different mail relay hosts that you use under different circumstances or if you want to test different database limits, you can create several system policies and switch between them, rather than editing a single policy.

Contrast a system policy, which controls aspects of an appliance that are likely to be similar across a deployment, with system settings, which are likely to be specific to a single appliance. See [Configuring Appliance Settings](#) on page 1998 for more information.

See the following sections for more information:

- [Creating a System Policy](#) on page 1962
- [Editing a System Policy](#) on page 1963
- [Applying a System Policy](#) on page 1965
- [Comparing System Policies](#) on page 1965
- [Deleting System Policies](#) on page 1969

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 appliance and then import it onto your appliance. You can then edit the imported policy to suit your needs before you apply it. For more information, see [Importing and Exporting Objects](#) on page 2225.

To create a system policy:

ACCESS: Admin

1. Select **System > Local > System Policy**.

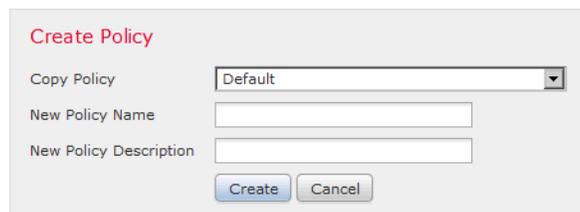
The System Policy page appears.

Policy Name	Applied To	Last Modified	
Default Default System Policy	None	2012-04-06 12:37:34	   
Initial_System_Policy 2012-04-06 12:43:18 Initial System Policy	None	2012-04-06 12:43:18	   

The **Policy Name** column includes the system policy's description. The **Applied To** column indicates the number of appliances where the policy is applied and a count of **out-of-date** appliances where the previously applied policy has changed and should be reapplied.

2. Click **Create Policy**.

The Create Policy page appears.



3. From the drop-down list, select an existing policy to use as a template for your new system policy.
4. Type a name for your new policy in the **New Policy Name** field.
5. Type a description for your new policy in the **New Policy Description** field.
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 the Access List for Your Appliance](#) on page 1971
- [Configuring Audit Log Settings](#) on page 1973
- [Configuring Authentication Profiles](#) on page 1975
- [Configuring Dashboard Settings](#) on page 1977
- [Configuring Database Event Limits](#) on page 1978
- [Configuring DNS Cache Properties](#) on page 1981
- [Configuring a Mail Relay Host and Notification Address](#) on page 1983
- [Configuring Access Control Policy Preferences](#) on page 1970
- [Configuring Intrusion Policy Preferences](#) on page 1985
- [Specifying a Different Language](#) on page 1986
- [Adding a Custom Login Banner](#) on page 1987
- [Configuring SNMP Polling](#) on page 1988
- [Synchronizing Time](#) on page 1991
- [Serving Time from the Defense Center](#) on page 1993
- [Configuring User Interface Settings](#) on page 1994
- [Mapping Vulnerabilities for Servers](#) on page 1996

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 appliance, reapply the policy after you have saved your changes. For

more information, see [Applying a System Policy](#) on page 1965.

To edit an existing system policy:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears, including a list of the existing system policies.

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 Access Control Policy Preferences](#) on page 1970
- [Configuring the Access List for Your Appliance](#) on page 1971
- [Configuring Audit Log Settings](#) on page 1973
- [Configuring Authentication Profiles](#) on page 1975
- [Configuring Dashboard Settings](#) on page 1977
- [Configuring Database Event Limits](#) on page 1978
- [Configuring DNS Cache Properties](#) on page 1981
- [Configuring a Mail Relay Host and Notification Address](#) on page 1983
- [Configuring Intrusion Policy Preferences](#) on page 1985
- [Specifying a Different Language](#) on page 1986
- [Adding a Custom Login Banner](#) on page 1987
- [Configuring SNMP Polling](#) on page 1988
- [Synchronizing Time](#) on page 1991
- [Serving Time from the Defense Center](#) on page 1993
- [Configuring User Interface Settings](#) on page 1994
- [Mapping Vulnerabilities for Servers](#) on page 1996

IMPORTANT! If you are editing a system policy applied to an appliance, make sure you reapply the updated policy when you are finished. See [Applying a System Policy](#) on page 1965.

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 appliance. If a system policy is already applied, any changes you make do not take effect until you reapply it.

To apply a system policy:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

Policy Name	Applied To	Last Modified	
Default Default System Policy	None	2012-04-06 12:37:34	   
Initial_System_Policy 2012-04-06 12:43:18 Initial System Policy	None	2012-04-06 12:43:18	   

2. Click the apply icon () next to the system policy that you want to apply.
The Apply page appears.
3. Select the appliances to which you want to apply the system policy.

TIP! You can sort the appliances by group, model, health policy, or applied system policy. You can select either an individual appliance or an entire group.

4. Click **Apply**.

The System Policy page appears. A message indicates the status of applying the system policy.

Comparing System Policies

LICENSE: Any

You can compare two system policies or two revisions of the same system policy, subject to the system policies you can access. This allows you to review policy changes for compliance with your organization's standards, or for optimization of system performance. To quickly compare your active system policy to another, you can select the **Running Configuration** option. Optionally, after you compare, you can generate a PDF report to record the differences between the system policies or system policy revisions.

There are two tools you can use to compare system policies or system policy revisions:

- The comparison view displays the differences between two system policies or system 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 web interface, with their differences highlighted.

- The comparison report creates a record of the differences between two system policies or system policy revisions in a format similar to the system policy report, but in PDF format.

You can use this to save, copy, print, and share your policy comparisons for further examination.

Using the System Policy Comparison View

LICENSE: Any

The comparison view displays both system 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. For all revisions, the system policy comparison view displays the time of last modification and the last user to the right of the policy name.

Default (2011-11-09 16:22:34 by admin)	Initial_System_Policy 2011-11-09 16:38:10 (2011-11-10 16:18:57 by admin)
Policy Information	
Name	Default
Description	Default System Policy
Modified	2011-11-09 16:22:34 by admin
Applied To	
Email Notification	
Mail Relay Host	
Time Synchronization	
Settings on Defense Center, Master Defense	
Set My Clock	Manually in System Settings
Settings on 3D Sensor	
Set My Clock	Manually in System Settings
Policy Information	
Name	Initial_System_Policy 2011-
Description	Initial System Policy
Modified	2011-11-10 16:18:57 by admin
Applied To	katsura
Email Notification	
Mail Relay Host	mail.example.com
Time Synchronization	
Settings on Defense Center, Master Defense	
Set My Clock	Via NTP from 0.sourcefire.pr
Settings on 3D Sensor	
Set My Clock	Via NTP from 0.sourcefire.pr

Differences between the two system 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 in the [System Policy Comparison View Actions](#) table.

System Policy Comparison View Actions

To...	YOU CAN...
navigate individually through changes	select 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 system policy comparison view	select New Comparison . The Select Comparison window appears. See Using the System Policy Comparison Report for more information.
generate a system policy comparison report	select Comparison Report . The system policy comparison report is a PDF that contains information identical to the system policy comparison view.

Using the System Policy Comparison Report

LICENSE: Any

A system policy comparison report is a record of all differences between two system policies or two revisions of the same system policy identified by the system policy comparison view, presented in PDF format. You can use this report to further examine the differences between two system policy configurations and to save and disseminate your findings.

You can generate a system policy comparison report from the comparison view for any system policies to which you have access. Changes you make to a system policy do not appear in the system policy comparison report until you save the changes.

Depending on your configuration, a system policy comparison report can contain one or more sections. The following sample graphic displays the Policy Information, User Detection Settings, and Time Synchronization sections of a system policy comparison report, and lists the configuration for each rule for both system policy configurations. Each section uses the same format and provides

the same level of detail. Note that the Value A and Value B columns represent the policies or policy revisions you configured in the comparison view.

Policy Information

Field	Value A	Value B
Name	katsura system policy	Test Policy 1
Description		Test Policy
Modified	2011-11-09 17:31:56 by admin	2011-11-11 15:09:08 by admin
Applied To	katsura	

Email Notification

Field	Value A	Value B
Mail Relay Host	mail.example1.com	mail.example.com
Encryption Method	None	TLS
From Address	admin1@example.com	admin1@example1.com

Time Synchronization

Field	Value A	Value B
Settings on Defense Center, Master Defense Center > Set My Clock	Via NTP from 0.sourcefire.pool.ntp.org	Manually in System Settings

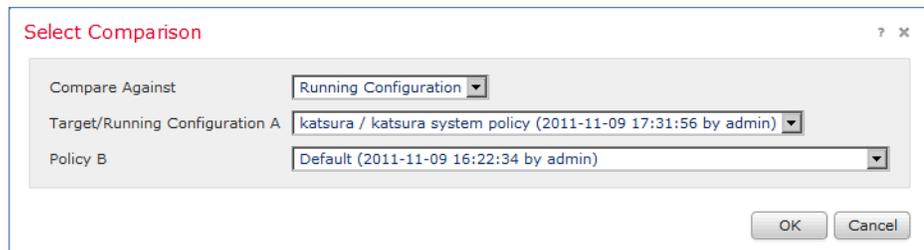
You use a similar procedure to compare other types of policies on the Sourcefire 3D System. For more information, see:

- [Comparing Two Intrusion Policies](#) on page 685
- [Comparing Health Policies](#) on page 2150

To compare two system policies or two revisions of the same policy:

ACCESS: Admin

1. Select **System > Local > System Policy**.
The System Policy page appears.
2. Click **Compare Policies**.
The Select Comparison pop-up window appears.



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**.
 - To compare another policy to a currently active policy, select **Running Configuration**.

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.
 - If you are comparing a running configuration to another policy, select the running configuration from the **Target/Running Configuration A** drop-down list, and the other policy from the **Policy B** drop-down list.
5. Click **OK** to display the system policy comparison view.
The comparison view appears.
6. Click **Comparison Report** to generate the system policy comparison report.
The system 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.

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:

ACCESS: Admin

1. Select **System > Local > System Policy**.
The System Policy page appears.
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 Access Control Policy Preferences](#) on page 1970
- [Configuring the Access List for Your Appliance](#) on page 1971

- [Configuring Audit Log Settings](#) on page 1973
- [Configuring Authentication Profiles](#) on page 1975
- [Configuring Dashboard Settings](#) on page 1977
- [Configuring Database Event Limits](#) on page 1978
- [Configuring DNS Cache Properties](#) on page 1981
- [Configuring a Mail Relay Host and Notification Address](#) on page 1983
- [Configuring Intrusion Policy Preferences](#) on page 1985
- [Specifying a Different Language](#) on page 1986
- [Adding a Custom Login Banner](#) on page 1987
- [Synchronizing Time](#) on page 1991
- [Serving Time from the Defense Center](#) on page 1993
- [Configuring User Interface Settings](#) on page 1994
- [Mapping Vulnerabilities for Servers](#) on page 1996

Configuring Access Control Policy Preferences

LICENSE: Protection

You can configure the system to prompt users for a comment when they add or modify a rule in an access control policy, prompting them to enter a rule comment. You can use this to track users' reasons for policy changes. If you enable comments on access control rule changes, you can make the rule comment optional or mandatory. The system prompts the user for a comment when each new change to a rule is saved.

The system adds the comment to the rule's comment history when the user saves the rule. For more information, see [Adding Rule Comments](#) on page 762.

To configure the access control policy rule comment settings:

ACCESS: Admin

1. Select **System > Local > System Policy.**

The System Policy page appears.

2. You have the following options:

- To modify the access control policy settings in an existing system policy, click the edit icon () next to the system policy.
- To configure the access control policy 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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Access Control Preferences**.

The Access Control Preferences page appears.



4. You have the following options:

- Select **Disabled** from the drop-down list to allow users to add or modify a rule in an access control policy without entering a comment.
- Select **Optional** from the drop-down list to display the Description of Changes (Optional) window to users when they save changes to access control policy rules. This allows users the option to describe changes in a comment.
- Select **Required** from the drop-down list to display the Description of Changes (Required) window to users when they save changes to access control policy rules. This requires users to describe changes in a comment before the changes are saved.

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](#) on page 1965 for more information.

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.

WARNING! 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 Access to the Database](#) on page 2007.

To configure the access list:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.



Host	Port	
any	443	✎
any	22	✎

3. Optionally, to delete one of the current settings, click the delete icon (✎). The setting is removed.

WARNING! 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.

4. Optionally, to add access for one or more IP addresses, click **Add Rules**. The Add IP Address page appears.



IP Address

Port SSH HTTPS SNMP

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 range using CIDR notation (for example, 192.168.1.1/24)
For information on using CIDR in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
 - any, to designate any IP address

6. Select **SSH, HTTPS, SNMP**, or a combination of these options to specify which ports you want to enable for these IP addresses.
7. Click **Add**.
The Access List page appears again, reflecting the changes you made.
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](#) on page 1965 for more information.

Configuring Audit Log Settings

LICENSE: Any

You can configure the system policy so that the appliance streams an audit log to an external host.

IMPORTANT! You must ensure that the external host is functional and accessible from the appliance 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 appliance 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:

ACCESS: Admin

1. Select **System > Local > System Policy**.
The System Policy page appears.

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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Audit Log Settings**.

The Audit Log Settings page appears.

The screenshot shows the 'Audit Log Settings' configuration page. It contains several fields and dropdown menus:

- Send Audit Log to Syslog:** A dropdown menu currently set to 'Disabled'.
- Host:** A text input field.
- Facility:** A dropdown menu currently set to 'USER'.
- Severity:** A dropdown menu currently set to 'INFO'.
- Tag (optional):** A text input field.
- Send Audit Log to HTTP Server:** A dropdown menu currently set to 'Disabled'.
- URL to Post Audit:** A text input field.

4. Select **Enabled** from the **Send Audit Log to Syslog** drop-down menu. (The default setting is **Disabled**.)
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.

WARNING! 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.

6. Select a syslog facility from the **Facility** field.
7. Select a severity from the **Severity** field.
8. Optionally, insert a reference tag in the **Tag (optional)** field.
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**.
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)

WARNING! To allow encrypted posts, you must use an HTTPS URL. Note that sending audit information to an external URL may affect system performance.

11. Click `Save Policy and Exit`.

The system policy is updated. Your changes do not take effect until you apply the system policy to the Defense Center and its managed devices. See [Applying a System Policy](#) on page 1965 for more information.

Configuring Authentication Profiles

LICENSE: Any

Normally, when a user logs into an appliance, the appliance verifies user credentials by comparing the credentials to a user account stored in the appliance's local database. However, if you create an authentication object referencing an external authentication server, you can apply the system policy to let users logging into the Defense Center or managed device authenticate to that server, rather than using the local database.

When you apply a system policy with authentication enabled to an appliance, the appliance verifies the user credentials against users on an LDAP or RADIUS server. In addition, if a user has internal authentication enabled and the user credentials are not found in the internal database, the appliance then checks the external server for a set of matching credentials. If a user has the same username on multiple systems, all passwords across all servers work. Note, however, that if authentication fails on the available external authentication servers, the appliance does not revert to checking the local database.

When you enable authentication, you can set the default user role for any user whose account is externally authenticated. You can select multiple roles, as long as those roles can be combined. For example, if you set up an authentication profile that retrieves only users in the Network Security group in your company, you may set the default user role to include the Security Analyst role so users can access collected event data without any additional user configuration on your part. However, if your authentication profile retrieves records for other personnel in addition to the security group, you would probably want to leave the default role unselected. For more information on available user roles, see [Understanding User Privileges](#) on page 1852.

Note that when you create an LDAP authentication object on your Defense Center, you can set a filter search attribute to specify the set of users who can successfully authenticate against the LDAP server. See [Configuring LDAP-Specific Parameters](#) on page 1869 for more information.

If no access role is selected, users can log in but cannot access any functionality. After a user attempts to log in, their account is listed on the User Management page, where you can edit the account settings to grant additional permissions. For more information on modifying a user account, see [Modifying User Privileges and Options](#) on page 1912. For a complete procedure for logging in initially as an externally authenticated user, see [Logging into the Appliance to Set Up an Account](#) on page 60.

If you configure the system policy to use one user role and apply the policy, then later modify the policy to use different default user roles and reapply, any user accounts created before the modification retain the first user role until you modify the accounts, or delete and recreate them.

You can enable authentication in a system policy on your Defense Center and then push that policy to managed devices. After you apply the policy to a device, eligible externally authenticated users can log into that device. To make changes to the authentication profile settings, you have to modify the system policy on the Defense Center, and then apply the policy to the device again. To disable authentication on a managed device, you can disable it in a system policy on the Defense Center and push that to the device.

If a user with internal authentication attempts to log in, the appliance first checks if that user is in the local user database. If the user exists, the appliance then checks the username and password against the local database. If a match is found, the user logs in successfully. If the login fails, however, and external authentication is enabled, the appliance checks the user against each external authentication server in the authentication order shown in the system policy. If the username and password match results from an external server, the appliance changes the user to an external user with the default privileges for that authentication object.

If an external user attempts to log in, the appliance checks the username and password against the external database. If a match is found, the user logs in successfully. If the login fails, the user login attempt is rejected. External users cannot authenticate against the user list in the local database. If the user is a new external user, an external user account is created in the local database with the default privileges for the external authentication object.

To enable authentication of users on external servers:

ACCESS: Admin

1. Select **System > Local > System Policy**.
The System Policy page appears.

2. You have the following options:
 - To modify the authentication profile settings in an existing system policy, click the edit icon () next to the system policy.
 - To configure the authentication profile 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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Authentication Profiles**.
The Authentication Profiles page appears.



4. From the **Status** drop-down list, select **Enabled**.
5. From the **Default User Role** drop-down list, select user roles to define the default permissions you want to grant to externally authenticated users.

TIP! Press Ctrl before selecting roles to select multiple default user roles. Note that although you can select both a Security Analyst role and the corresponding Security Analyst (Read Only) role, only the Security Analyst role is applied.

6. If you want to use the external server to authenticate shell access accounts as well, select **Enabled** from the **Shell Authentication** drop-down list.
7. Click **Save Policy and Exit**.
The system policy is updated. Your changes do not take effect until you apply the system policy to the Defense Center and its managed devices. See [Applying a System Policy](#) on page 1965 for more information.

Configuring Dashboard Settings

LICENSE: Any

You can configure the system policy so that Custom Analysis widgets are enabled on the dashboard. Dashboards provide you with at-a-glance views of current system status through the use of widgets: small, self-contained components that provide insight into different aspects of the Sourcefire 3D System.

The Custom Analysis widget allows you to create a visual representation of events based on a flexible, user-configurable query of the events in your

appliance's database. See [Understanding the Custom Analysis Widget](#) on page 79 for more information on how to use custom widgets.

To enable Custom Analysis widgets:

ACCESS: Admin

1. Select **System > Local > System Policy**.
The System Policy page appears.
2. You have the following options:
 - To modify the dashboard settings in an existing system policy, click the edit icon (✎) next to the system policy.
 - To configure the dashboard 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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Dashboard**.
The Dashboard Settings page appears.



Enable Custom Analysis Widgets

4. Select the **Enable Custom Analysis Widgets** check box to allow users to add Custom Analysis widgets to dashboards. Clear the check box to prohibit users from using those widgets.
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](#) on page 1965 for more information.

Configuring Database Event Limits

LICENSE: Any

Use the Database page to specify the maximum number of each type of event that the Defense Center can store. Note that the setting for audit records also applies to managed devices. To improve performance, you should tailor event limits to the number of events you regularly work with. For some event types, you

can disable storage. The following table lists the minimum and maximum number of records you can store for each event type.

Database Event Limits

EVENT TYPE	UPPER EVENT LIMIT	LOWER EVENT LIMIT
intrusion events	2.5 million (DC500) 10 million (DC1000, virtual Defense Center) 20 million (DC750) 30 million (DC1500) 100 million (DC3000) 150 million (DC3500)	10,000
discovery events	10 million	zero (disables storage)
connection events	10 million (DC500, DC1000, virtual Defense Center) 50 million (DC750) 100 million (DC1500, DC3000) 500 million (DC3500)	zero (disables storage)
connection summaries (aggregated connection events)	10 million (DC500, DC1000, virtual Defense Center) 50 million (DC750) 100 million (DC1500, DC3000) 500 million (DC3500)	zero (disables storage)
correlation and compliance white list events	1 million	one
malware events	10 million	10,000
file events	10 million	zero (disables storage)
health events	1 million	zero (disables storage)
audit records	100,000	one
remediation status events	10 million	one
the white list violation history of the hosts on your network	a 30-day history of violations	one day's history

Database Event Limits (Continued)

EVENT TYPE	UPPER EVENT LIMIT	LOWER EVENT LIMIT
user activity (user events)	10 million	one
user logins (user history)	10 million	one
rule update import log records	1 million	one

If the number of events in the intrusion event database exceeds the maximum, the oldest events and packet files are pruned until the database is back within limits. In addition, if the `/volume` disk partition reaches 80% of its capacity, unified files are deleted from the system until available disk space is again within acceptable limits, beginning with the oldest files. See [Configuring a Mail Relay Host and Notification Address](#) on page 1983 for information about generating automated email notifications when events are automatically pruned.

For information on manually pruning the discovery and user databases, see [Purging Discovery Data from the Database](#) on page 2236.

In addition, you can configure an email address that will receive notifications when intrusion events and audit records are pruned from the database.

To configure the maximum number of records in the database:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

2. You have the following options:

- To modify the database settings in an existing system policy, click the edit icon (✎) next to the system policy.
- To configure the database 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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Database**.

The Database page appears.

The screenshot shows a configuration page with four sections:

- Intrusion Event Database**: Supported Platforms (Defense Center), Maximum Intrusion Events (1000000)
- Discovery Event Database**: Supported Platforms (Defense Center), Maximum Discovery Events (0 - do not store) (1000000)
- Rule Update Import Log Database**: Supported Platforms (ALL), Maximum Rule Update Import Log (1000000)
- Data Pruning**: Supported Platforms (ALL), Data Pruning Notification Address (empty field)

4. For each of the databases, enter the number of records you want to store.

For information on how many records each database can maintain, see the [Database Event Limits table](#) on page 1979.

5. Optionally, in the **Data Pruning Notification Address** field, enter the email address you want to receive notifications when intrusion events and audit records are pruned from the appliance's database.

6. 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](#) on page 1965 for more information.

Configuring DNS Cache Properties

LICENSE: Any

If you have a DNS server configured on the Network page, you can configure the appliance to resolve IP addresses automatically on the event view pages. As a user assigned the Administrator role, you can also configure basic properties for DNS caching performed by the appliance. Configuring DNS caching allows you to identify IP addresses you previously resolved without performing additional lookups. This can reduce the amount of traffic on your network and speed the display of event pages when IP address resolution is enabled.

To configure the DNS cache properties:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

2. You have the following options:

- To modify the DNS cache settings in an existing system policy, click the edit icon (✎) next to the system policy.
- To configure the DNS cache 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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **DNS Cache**.

The DNS Cache page appears.



DNS Resolution Caching	Enabled ▾
DNS Cache Timeout (in minutes)	300

4. Select **Enabled** from the **DNS Resolution Caching** drop-down list to enable caching. Select **Disabled** to disable it.

IMPORTANT! DNS resolution caching is a system-wide setting that allows the caching of previously resolved DNS lookups. To configure IP address resolution on a per-user-account basis, users must also select **Event View Settings** from the **User Preferences** menu, enable **Resolve IP Addresses**, and then click **Save**. For information about configuring DNS servers, see [Configuring Network Settings](#) on page 2009. For information about configuring event view preferences, see [Configuring Event View Settings](#) on page 2218.

5. In the **DNS Cache Timeout (in minutes)** field, enter the number of minutes a DNS entry remains cached in memory before it is removed for inactivity. The default setting is 300 minutes (five hours).

6. 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](#) on page 1965 for more information.

WARNING! Although DNS caching is enabled for the appliance, IP address resolution is not enabled on a per-user basis unless it is configured on the Events page accessed from the User Preferences menu.

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
- use email for discovery event, impact flag, and correlation event alerting
- use email for intrusion event alerting
- use email for health 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:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

2. You have the following options:

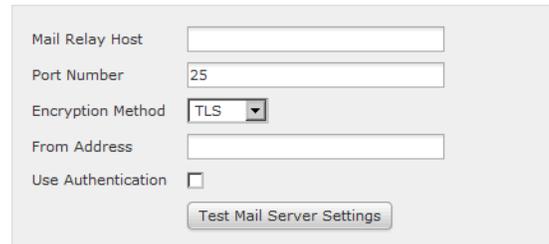
- To modify the email settings in an existing system policy, click the edit icon () next to the system policy.
- To configure the email settings as part of a new system policy, click **Create Policy**.

Provide a name and description for the system policy as described in [Creating a System Policy](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Email Notification**.

The Configure Email Notification page appears.



The screenshot shows a configuration form for email notifications. It contains the following fields and controls:

- Mail Relay Host:** A text input field.
- Port Number:** A text input field containing the value "25".
- Encryption Method:** A dropdown menu with "TLS" selected.
- From Address:** A text input field.
- Use Authentication:** An unchecked checkbox.
- Test Mail Server Settings:** A button located below the checkbox.

4. In the **Mail Relay Host** field, type the hostname or IP address of the mail server you want to use.

IMPORTANT! The mail host you enter must allow access from the appliance.

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.

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.

7. Enter a valid email address in the **From Address** field for use as the source email address for messages sent by the appliance.
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.
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.

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](#) on page 1965 for more information.

Configuring Intrusion Policy Preferences

LICENSE: Protection

You can configure the system to prompt users for a comment when they modify an intrusion policy. You can use this to track users' reasons for policy changes. If you enable comments on intrusion policy changes, you can make the comments optional or mandatory. The change description is written to the audit log.

You can also have all intrusion policy changes written to the audit log. For more information on the audit log, see [Managing Audit Records](#) on page 2187.

To configure the intrusion policy comment settings:

ACCESS: Admin

1. Select **System > Local > System Policy.**

The System Policy page appears.

2. You have the following options:

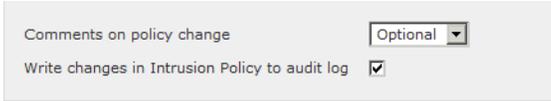
- To modify the intrusion policy preferences in an existing system policy, click the edit icon () next to the system policy.
- To configure the intrusion policy preferences as part of a new system policy, click **Create Policy**.

Provide a name and description for the system policy as described in [Creating a System Policy](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Intrusion Policy Preferences.**

The Intrusion Policy Preferences page appears.



Comments on policy change Optional ▾
Write changes in Intrusion Policy to audit log

4. From the **Comments on policy change** drop-down list, you have the following options:
 - Select **Disabled** to allow users to modify an intrusion policy without entering a change description.
 - Select **Optional** to display the Description of Changes window to users when they save changes to an intrusion policy. This allows users the option to describe changes in a comment.
 - Select **Required** to display the Description of Changes window to users when they save changes to an intrusion policy. This requires users to describe changes in a comment before the changes are saved.
5. Optionally, if you want to write all intrusion policy changes to the audit log, select **Write changes in Intrusion Policy to audit log**.
6. 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](#) on page 1965 for more information.

Specifying a Different Language

LICENSE: Any

You can use the Language page to specify a different language for the web interface.

WARNING! The language you select here is used for the web interface for every user who logs into the appliance.

To select a different language for the user interface:

ACCESS: Admin

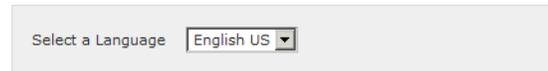
1. Select **System > Local > System Policy**.

The System Policy page appears.
 2. You have the following options:
 - To modify the language settings in an existing system policy, click the edit icon (✎) next to the system policy.
 - To configure the language 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](#) on page 1962, and click **Save**.
- In either case, the Access List page appears.

3. Click **Language**.

The Language page appears.



A screenshot of a web interface showing a language selection dropdown menu. The text 'Select a Language' is on the left, and a dropdown box on the right contains 'English US' with a downward arrow.

4. Select the language you want to use.
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](#) on page 1965 for more information.

Adding a Custom Login Banner

LICENSE: Any

You can create a custom login banner that appears when users log into the appliance using SSH and on the login page of the web interface. Banners can contain any printable characters except the less-than symbol (<) and the greater-than symbol (>).

To add a custom banner:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

2. You have the following options:

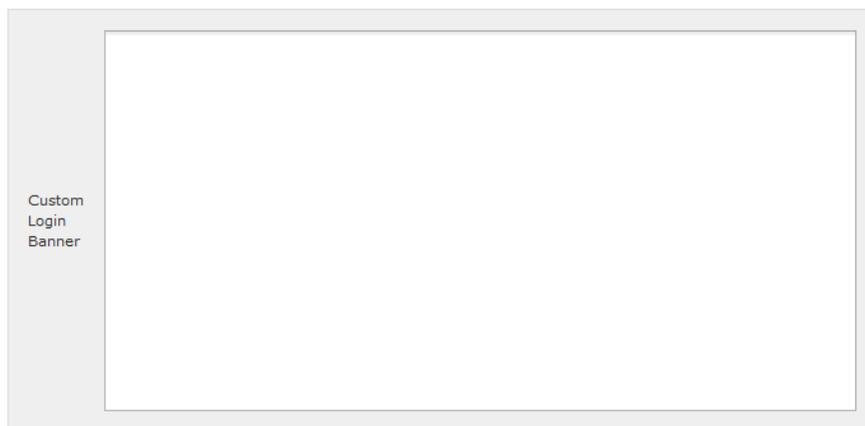
- To modify the login banner in an existing system policy, click the edit icon () next to the system policy.
- To configure the login banner as part of a new system policy, click **Create Policy**.

Provide a name and description for the system policy as described in [Creating a System Policy](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Login Banner**.

The Login Banner page appears.



4. In the **Custom Login Banner** field, enter the login banner you want to use with this system policy.
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](#) on page 1965 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.

This feature allows access to:

- the standard management information base (MIB) for the appliance, which includes system details such as contact, administrative, location, service information, IP addressing and routing information, and transmission protocol usage statistics
- additional MIBs for managed devices that include statistics on traffic passing through physical interfaces, logical interfaces, virtual interfaces, ARP, NDP, virtual bridges, and virtual routers

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.

IMPORTANT! 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](#) on page 1971. Note that the SNMP MIB contains information that could be used to attack your appliance. Sourcefire recommends that you restrict your access list for SNMP access to the specific hosts that will be used to poll for the MIB. Sourcefire also recommends you use SNMPv3 and use strong passwords for network management access.

To configure SNMP polling:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

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**.

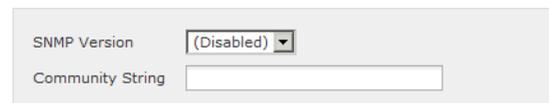
Provide a name and description for the system policy as described in [Creating a System Policy](#) on page 1962, and click **Create**.

In either case, the Access List page appears.

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](#) on page 1971.

4. Click **SNMP**.

The SNMP page appears.



The screenshot shows a configuration form with two fields. The first field is labeled "SNMP Version" and has a dropdown menu with "(Disabled)" selected. The second field is labeled "Community String" and has an empty text input box.

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.

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.

The screenshot shows a user definition form with the following fields and controls:

- Username**: A text input field.
- Authentication Protocol**: A dropdown menu currently set to "None".
- Authentication Password**: A text input field.
- Verify Password**: A text input field.
- Privacy Protocol**: A dropdown menu currently set to "None".
- Privacy Password**: A text input field.
- Verify Password**: A text input field.
- Add**: A button at the bottom of the form.

7. Enter a username in the **Username** field.
8. Select the protocol you want to use for authentication from the **Authentication Protocol** drop-down list.
9. Type the password required for authentication with the SNMP server in the **Authentication Password** field.
10. Retype the authentication password in the **Verify Password** field just below the **Authentication Password** field.
11. Select the privacy protocol you want to use from the **Privacy Protocol** list, or select **None** to not use a privacy protocol.
12. Type the SNMP privacy key required by the SNMP server in the **Privacy Password** field.
13. Retype the privacy password in the **Verify Password** field just below the **Privacy Password** field.
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.
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](#) on page 1965 for more information.

Synchronizing Time

LICENSE: Any

You can manage time synchronization on the appliance using the Time Synchronization page. You can choose to synchronize the time:

- manually
- using one or more NTP servers (one of which can be a Defense Center)

Time settings are part of the system policy. You can specify the time settings either by creating a new system policy or by editing an existing policy. In either case, the time setting is not used until you apply the system policy.

Note that time settings are displayed on most pages on the appliance in local time using the time zone you set on the Time Zone page (America/New York by default), but are stored on the appliance itself using UTC time. In addition, the current time appears in UTC at the top of the Time Synchronization page (local time is displayed in the Manual clock setting option, if enabled).

You can synchronize the appliance's time with an external time server. If you specify a remote NTP server, your appliance must have network access to it. Do not specify an untrusted NTP server. Connections to NTP servers do not use configured proxy settings. To use the Defense Center as an NTP server, see [Serving Time from the Defense Center](#) on page 1993.

Sourcefire recommends that you synchronize your virtual appliances to a physical NTP server. Do not synchronize your managed devices (virtual or physical) to a Virtual Defense Center.

The procedure for synchronizing time differs slightly depending on whether you are using the web interface on a Defense Center or a managed device. Each procedure is explained separately below.

To synchronize time:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Time Synchronization**.

The Time Synchronization page appears.

Supported Platforms	Defense Center
Serve Time via NTP	Enabled
Set My Clock	<input type="radio"/> Manually in Local Configuration <input checked="" type="radio"/> Via NTP from example.ntp.org
Supported Platforms	Managed Device
Set My Clock	<input type="radio"/> Manually in Local Configuration <input checked="" type="radio"/> Via NTP from Defense Center <input type="radio"/> Via NTP from example.ntp.org

4. If you want to serve time from the Defense Center to your managed devices, in the **Serve time via NTP** drop-down list, select **Enabled**.

5. You have the following options for specifying how the time is synchronized on the Defense Center:

- To set the time manually, select **Manually in Local Configuration**. See [Setting the Time Manually](#) on page 2015 for information about setting the time after you apply the system policy.
- To receive time through NTP from a different server, select **Via NTP from** and, in the text box, type a comma-separated list of IP addresses for the NTP servers you want to use or, if DNS is enabled, type the fully qualified host and domain names.

WARNING! If the appliance is rebooted and your DHCP server sets an NTP server record different than the one you specify here, the DHCP-provided NTP server will be used instead. To avoid this situation, configure your DHCP server to set the same NTP server.

6. You have the following options for specifying how time is synchronized on any managed devices:
 - Select **Manually in Local Configuration** to set the time manually. See [Setting the Time Manually](#) on page 2015 for information about setting the time after you apply the system policy.
 - Select **Via NTP from Defense Center** to receive time through NTP from the Defense Center. See [Serving Time from the Defense Center](#) on page 1993 for more information.
 - Select **Via NTP from** to receive time through NTP from different servers. In the text box, type a comma-separated list of IP addresses of the NTP servers or, if DNS is enabled, type the fully qualified host and domain names.

IMPORTANT! It may take a few minutes for the managed device to synchronize with the configured NTP servers. In addition, if you are synchronizing the managed device to a Defense Center that is configured as an NTP server, and the Defense Center itself is configured to use an NTP server, it may take some time for the time to synchronize. This is because the Defense Center must first synchronize with its configured NTP server before it can serve time to the managed device.

7. 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](#) on page 1965 for more information.

Serving Time from the Defense Center

LICENSE: Any

You can configure the Defense Center as a time server using NTP and then use it to synchronize time between the Defense Center and managed devices.

Note that you cannot set the time manually after configuring the Defense Center to serve time using NTP. If you need to manually change the time, you should do so **before** configuring the Defense Center to serve time using NTP. If you need to change the time manually **after** configuring the Defense Center as an NTP server, disable the **Via NTP** option and click **Save**, change the time manually and click **Save**, and then enable **Via NTP** and click **Save**.

IMPORTANT! If you configure the Defense Center to serve time using NTP, and then later disable it, the NTP service on managed devices still attempts to synchronize time with the Defense Center. You must disable NTP from the managed devices' web interfaces to stop the synchronization attempts.

To configure the Defense Center as an NTP server:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

2. You have the following options:

- To modify the NTP server settings in an existing system policy, click the edit icon (✎) next to the system policy.
- To configure the NTP server 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](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Time Synchronization**.

The Time Synchronization page appears.

4. From the **Serve Time via NTP** drop-down list, select **Enabled**.

5. In the **Set My Clock** option for the managed device, select **Via NTP from Defense Center**.

6. Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy to the Defense Center and its managed devices. See [Applying a System Policy](#) on page 1965 for more information.

IMPORTANT! It may take a few minutes for the Defense Center to synchronize with its managed devices.

Configuring User Interface Settings

LICENSE: Any

Unattended login sessions of the Sourcefire 3D System web interface or command line interface may be security risks. You can configure, in minutes, the amount of idle time before a user's login session times out due to inactivity. You can also set a similar timeout for shell (command line) sessions.

Your deployment may have users who plan to passively, securely monitor the web interface for long periods of time. You can exempt users from the web interface session timeout with a user configuration option. (Users with the Administrator role, whose complete access to menu options poses an extra risk if compromised, cannot be made exempt from session timeouts.) For more information, see [Managing User Login Settings](#) on page 1903.

For cases in which you must restrict shell access to the system, a third option allows you to permanently disable the `expert` command in the command line. Disabling expert mode on an appliance prevents any user, even users with Configuration shell access, from going into expert mode in the shell. When a user goes into expert mode on the command line, the user can run any Linux command appropriate to the shell. When not in expert mode, command line users can only run the commands provided by the command line interface. Note that the command line interface is not supported for Series 2 appliances.

For more information on command line interface commands, see [Command Line Reference](#) on page 2242. For information on setting up users for command line access, see [Managing Command Line Access](#) on page 1901 and [Command Line Reference](#) on page 2242 (for virtual device CLI user management).

To configure user interface settings:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

2. You have the following options:

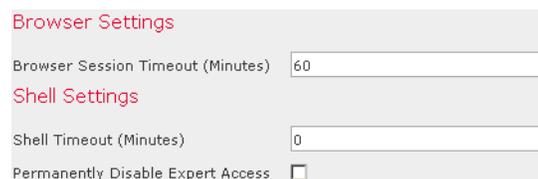
- To modify user interface settings in an existing system policy, click the edit icon (✎) next to the system policy.
- To configure user interface settings as part of a new system policy, click **Create Policy**.

Provide a name and description for the system policy as described in [Creating a System Policy](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **User Interface**.

The User Interface page appears.



The screenshot shows a configuration form with two sections: "Browser Settings" and "Shell Settings". Under "Browser Settings", there is a text input field for "Browser Session Timeout (Minutes)" containing the value "60". Under "Shell Settings", there is a text input field for "Shell Timeout (Minutes)" containing the value "0", and a checkbox for "Permanently Disable Expert Access" which is currently unchecked.

4. You have the following options:

- To configure session timeout for the web interface, type a number (of minutes) in the **Browser Session Timeout (Minutes)** field. The default value is 60; the maximum value is 1440 (24 hours).

For information on how to exempt users from this session timeout, see [Managing User Login Settings](#) on page 1903.

- To configure session timeout for the command line interface, type a number (of minutes) in the **Shell Timeout (Minutes)** field. The default value is 0; the maximum value is 1440 (24 hours).
- To permanently disable the **expert** command in the command line interface, select the **Permanently Disable Expert Access** check box.

WARNING! After you apply a system policy with expert mode disabled to an appliance, you cannot restore the ability to access expert mode through the web interface or the command line. You must contact Sourcefire Support to restore the expert mode capability.

5. Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy to the Defense Center and its managed devices. Changes to session timeout intervals do not take effect until the next login session.

Mapping Vulnerabilities for Servers

LICENSE: Protection

The Sourcefire 3D System automatically maps vulnerabilities to a host IP address for any application protocol traffic received or sent from that address, when the server has an application ID in the discovery event database and the packet header for the traffic includes a vendor and version.

However, many servers do not include vendor and version information. For the server listed in the system policy, you can configure whether the system associates vulnerabilities with server traffic for vendor and versionless servers.

For example, a host serves SMTP traffic that does not have a vendor or version in the header. If you enable the SMTP server on the Vulnerability Mapping page of a system policy, then apply that policy to the Defense Center managing the device that detects the traffic, all vulnerabilities associated with SMTP servers are added to the host profile for the host.

Although detectors collect server information and add it to host profiles, the application protocol detectors will not be used for vulnerability mapping, because you cannot specify a vendor or version for a custom application protocol detector and cannot select the server for vulnerability mapping in the system policy.

To configure vulnerability mapping for servers:

ACCESS: Admin

1. Select **System > Local > System Policy**.

The System Policy page appears.

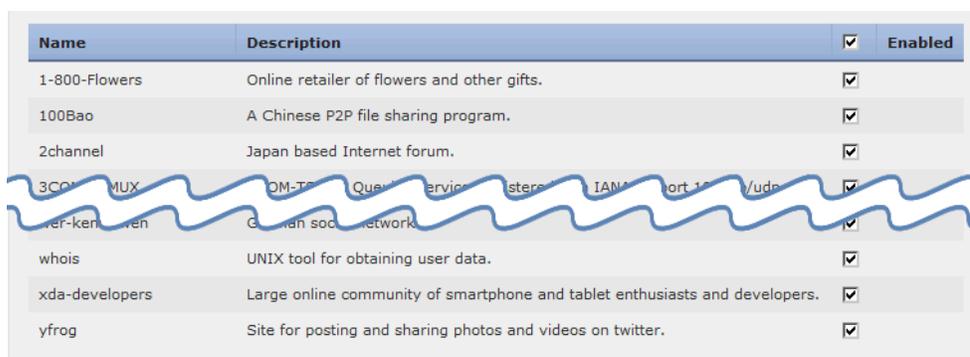
2. You have the following options:
 - To modify vulnerability mapping settings in an existing system policy, click the edit icon (✎) next to the system policy.
 - To configure vulnerability mapping settings as part of a new system policy, click **Create Policy**.

Provide a name and description for the system policy as described in [Creating a System Policy](#) on page 1962, and click **Save**.

In either case, the Access List page appears.

3. Click **Vulnerability Mapping**.

The Vulnerability Mapping page appears.



Name	Description	<input checked="" type="checkbox"/>	Enabled
1-800-Flowers	Online retailer of flowers and other gifts.	<input checked="" type="checkbox"/>	
100Bao	A Chinese P2P file sharing program.	<input checked="" type="checkbox"/>	
2channel	Japan based Internet forum.	<input checked="" type="checkbox"/>	
3COM MUX	OM-TE... Queue service... sters... IANA port 1... /udp	<input checked="" type="checkbox"/>	
er-ken... en	German social network	<input checked="" type="checkbox"/>	
whois	UNIX tool for obtaining user data.	<input checked="" type="checkbox"/>	
xda-developers	Large online community of smartphone and tablet enthusiasts and developers.	<input checked="" type="checkbox"/>	
yfrog	Site for posting and sharing photos and videos on twitter.	<input checked="" type="checkbox"/>	

4. You have the following options:
 - To prevent vulnerabilities for a server from being mapped to hosts that receive application protocol traffic without vendor or version information, clear the check box for that server.
 - To cause vulnerabilities for a server to be mapped to hosts that receive application protocol traffic without vendor or version information, select the check box for that server.

TIP! You can select or clear all check boxes at once using the check box next to **Enabled**.

5. Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy to the Defense Center and its managed devices. See [Applying a System Policy](#) on page 1965 for more information.

CHAPTER 49

CONFIGURING APPLIANCE SETTINGS

A Sourcefire 3D System appliance's *local configuration* (**System > Local > Configuration**) is a group of settings that is likely to be specific to a single appliance. Contrast the local configuration with the system policy ([Managing System Policies](#) on page 1961), which controls appliance settings that are likely to be similar across a deployment.

The following table summarizes an appliance's local configuration.

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 on page 1999
HTTPS Certificate	Allows you to request an HTTPS server certificate, if needed, from a trusted authority and upload certificates to your appliance.	Using Custom HTTPS Certificates on page 2002
Database	Lets you enable external read-only access to the appliance database, and provides a client driver for you to download.	Enabling Access to the Database on page 2007
Network	Enables you to change options such as the IP address, hostname, and proxy settings of the appliance that were initially set up as part of the installation.	Configuring Network Settings on page 2009

Local Configuration Options (Continued)

OPTION	DESCRIPTION	FOR MORE INFORMATION, SEE...
Management Interface	Allows you to view and modify the settings for the management interfaces on your appliance.	Editing Management Interface Configurations on page 2012
Process	Provides options that you can use to shut down or reboot the appliance, and restart Sourcefire 3D System-related processes.	Shutting Down and Restarting the System on page 2014
Time	Displays the current time. If the time synchronization settings in the current system policy for the appliance is set to Manually in Local Configuration , then you can use this page to change the time.	Setting the Time Manually on page 2015
Remote Storage Device	On Defense Centers, allows you to configure remote storage for backups and reports.	Managing Remote Storage on page 2017
Change Reconciliation	Allows you to receive, via email, a detailed report of changes to your system over the last 24 hours.	Understanding Change Reconciliation on page 2024
Console Configuration	Allows you configure console access to Sourcefire appliances via VGA or serial port, or via Lights-Out Management (LOM), which allows you to perform limited monitoring and management tasks without being physically near the appliance.	Managing Remote Console Access on page 2025
Cloud Services	On Defense Centers, allows you to download URL filtering data from the Sourcefire cloud, perform lookups for uncategorized URLs, and send diagnostic information on detected files to Sourcefire.	Enabling Sourcefire Cloud Communications on page 2032

Viewing and Modifying the Appliance Information

LICENSE: Any

The Information page provides you with information about your appliances. The information includes read-only information, such as the product name and model number, the operating system and version, and the current appliance-level policies. The page also provides you with an option to change the name of the appliance.

The [Appliance Information](#) table describes each field.

Appliance Information

FIELD	DESCRIPTION
Name	A name you assign to the appliance. Note that this name is only used within the context of the Sourcefire 3D System. 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.
Software Version	The version of the software currently installed.
Serial Number	The chassis serial number of the appliance.
Store Events Only on Defense Center	Select this check box on the managed device to store event data on the Defense Center, but not the managed device. Clear this check box to store event data on both appliances.
Prohibit Packet Transfer to the Defense Center	Select this check box on the managed device to prevent the managed device from sending packet data with the events. Clear this check box to allow packet data to be stored on the Defense Center with events.
Operating System	The operating system currently running on the appliance.
Operating System Version	The version of the operating system currently running on the appliance.
IPv4Address	The IPv4 address of the appliance. If IPv4 management is disabled for the appliance, this field indicates that.
IPv6 Address	The IPv6 address 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:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears. The Defense Center version of the page is shown below.

The screenshot shows the 'Appliance Information' page in Defense Center. The 'Name' field contains 'katsura'. Other fields include Product Model (Defense Center 3500), Serial Number (.....), Software Version (5.0.0), Operating System (Sourcefire Linux OS), Operating System Version (5.0.0), IPv4 Address (10.10.10.2), and IPv6 Address (Disabled). Under 'Current Policies', 'Health Policy' is 'Blacklisted Power Supply' and 'System Policy' is 'katsura system policy'. The 'Model Number' is '00'. At the bottom are 'Save' and 'Refresh' buttons.

For comparison, the managed device version of the page is shown below.

The screenshot shows the 'Appliance Information' page for a managed device. The 'Name' field contains 'linden'. Other fields include Product Model (3D8250), Serial Number (00000000000000-A), Software Version (5.0.0), Store Events Only on DC (checked), Prohibit Packet Transfer to the Defense Center (unchecked), Operating System (Sourcefire Linux OS), Operating System Version (5.0.0), IPv4 Address (10.10.10.6), and IPv6 Address (Disabled). Under 'Current Policies', 'Health Policy' is 'Blacklisted Power Supply (Remotely Authored by katsura)' and 'System Policy' is 'katsura system policy (Remotely Authored by katsura)'. The 'Model Number' is '00'. At the bottom are 'Save' and 'Refresh' buttons.

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.
3. To save your changes, click **Save**.
The page refreshes and your changes are saved.

Using Custom HTTPS Certificates

LICENSE: Any

Sourcefire Defense Centers and managed devices that support web-based user interfaces include default SSL (Secure Sockets Layer) certificates that you can use to initiate an encrypted communication channel between your web browser and the appliance. However, because the default certificate for an appliance is not generated by a certificate authority (CA) trusted by any globally known CA, you can replace it with a custom certificate signed by a globally known CA.

You can manage certificates through the local configuration for your appliance. For more information, see the following:

- [Viewing the Current HTTPS Server Certificate](#) on page 2002
- [Generating a Server Certificate Request](#) on page 2003
- [Uploading Server Certificates](#) on page 2004
- [Configuring User Certificates](#) on page 2006

Viewing the Current HTTPS Server Certificate

LICENSE: Any

You can view details from the server certificate currently in place for your appliance. The certificate provides the following information:

HTTPS Server Certificate Information

FIELD	DESCRIPTION
Subject	For the appliance where the certificate is installed, provides the commonName, countryName, organizationName, and organizationalUnitName.
Issuer	For the appliance that issued the certificate, provides the commonName, countryName, organizationName, and organizationalUnitName.
Validity	Indicates the timeframe during which the certificate is valid.
Version	Indicates the certificate version.
Serial Number	Indicates the certificate serial number.
Signature Algorithm	Indicates the algorithm used to sign the certificate.

To view the certificate details:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **HTTPS Certificate**.

The HTTPS Certificate page appears, with the details of the current certificate for the appliance.

Current HTTPS Certificate				
Subject	commonName Sourcefire3D	countryName US	organizationName Sourcefire, Inc.	organizationalUnitName Intrusion Management System
Issuer	commonName Sourcefire3D	countryName US	organizationName Sourcefire, Inc.	organizationalUnitName Intrusion Management System
Validity	Not Before Nov 9 21:22:18 2011 GMT	Not After Nov 9 21:22:18 2031 GMT		
Version	02			
Serial Number	FFFFFFFFFFFFFFFF			
Signature Algorithm	sha1WithRSAEncryption			

HTTPS User Certificate Settings

Enable User Certificates

Save

Generating a Server Certificate Request

LICENSE: Any

You can generate a certificate request based on your appliance information and the identification information you supply. You can send the resulting request to a certificate authority to request a server certificate. You can also use it to self-sign a certificate if you have an internal certificate authority (CA) installed that is trusted by your browser. The generated key is in Base-64 encoded PEM format.

Note that when you generate a certificate request through the local configuration HTTPS Certificate page, you can only generate a certificate for a single server. You must type the fully qualified domain name of the server exactly as it should appear in the certificate in the **Common Name** field. If the common name and the DNS host name do not match, you receive a warning when connecting to the appliance. Similarly, if you install a certificate that is not signed by a globally known authority, you receive a security warning when you connect to the appliance.

To generate a certificate request:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **HTTPS Certificate**.

The Current HTTPS Certificate page appears.

3. Click **Generate New CSR**.

The Generate Certificate Signing Request pop-up window appears.



Generate Certificate Signing Request

Country Name (two-letter code)

State or Province

Locality or City

Organization

Organizational Unit (Department)

Common Name

4. Type the two-letter country code for your country into the **Country Name (two-letter code)** field.
5. Type the postal abbreviation for your state or province in the **State or Province** field.
6. Type the name of your **Locality or City**.
7. Type your **Organization** name.
8. Type an **Organizational Unit (Department)** name.
9. Type the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field, exactly as you want it to appear in the certificate.
10. Click **Generate**.
The Certificate Signing Request pop-up window appears.
11. Open a text editor.
12. Copy the entire block of text in the certificate request, including the **BEGIN CERTIFICATE REQUEST** and **END CERTIFICATE REQUEST** lines, and paste it into a blank text file.
13. Save the file as *servername.csr*, where *servername* is the name of the server where you plan to use the certificate.
14. Upload the CSR file to the certificate authority where you want to request a certificate or use the CSR to create a self-signed certificate.

Uploading Server Certificates

LICENSE: Any

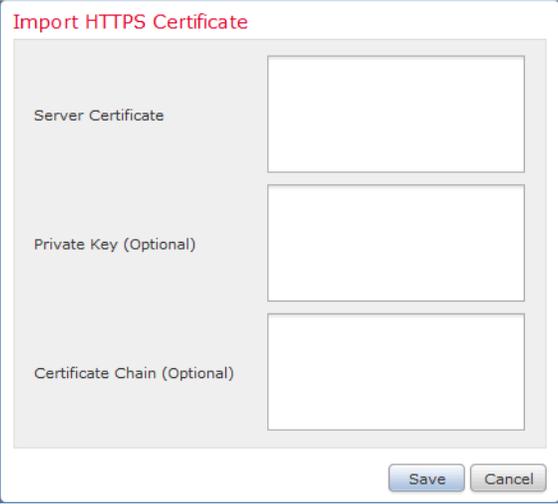
After you have a signed certificate from a certificate authority (CA), you can upload it. If the signing authority that generated the certificate requires you to

trust an intermediate CA, you must also supply a certificate chain, sometimes referred to as a certificate path. If you require user certificates, they must be generated by a certificate authority whose intermediate authority is included in the certificate chain.

To import a certificate:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.
2. Click **HTTPS Certificate**.
The Current HTTPS Certificate page appears.
3. Click **Import HTTPS Certificate**.
The Import HTTPS Certificate pop-up window appears.



The screenshot shows a dialog box titled "Import HTTPS Certificate". It contains three text input fields stacked vertically. The first field is labeled "Server Certificate", the second is "Private Key (Optional)", and the third is "Certificate Chain (Optional)". At the bottom right of the dialog, there are two buttons: "Save" and "Cancel".

4. Open the server certificate in a text editor, copy the entire block of text, including the **BEGIN CERTIFICATE** and **END CERTIFICATE** lines, and paste it into the **Server Certificate** field.
5. Optionally, open the private key file, copy the entire block of text, including the **BEGIN RSA PRIVATE KEY** and **END RSA PRIVATE KEY** lines, and paste it into the **Private Key** field.
6. Open any intermediate certificates you need to provide, copy the entire block of text, for each, and paste it into the **Certificate Chain** field.
7. Click **Save** to import the certificate.
The certificate imports and the Current Certificate View updates to reflect the new certificate.

Configuring User Certificates

LICENSE: Any

You can restrict access to the Sourcefire 3D System web server using client browser certificate checking. When you enable user certificates, the web server checks that a user's browser client has a valid user certificate selected. That user certificate must be generated by the same trusted certificate authority used for the server certificate. If the user selects a certificate in the browser that is not valid or not generated by a certificate authority in the certificate chain on the device, the browser cannot load the web interface.

You can also load a certificate revocation list (CRL) for the server. The CRL lists any certificates that have been revoked by the certificate authority, so the web server can verify that the client browser certificate has not been revoked. If the user selects a certificate that is listed in the CRL as a revoked certificate, the browser cannot load the web interface. The appliance supports upload of CRLs in Distinguished Encoding Rules (DER) format. You can only load one CRL for a server.

To ensure that the list of revoked certificates stays current, you can create a scheduled task to update the CRL. The most recent refresh of the CRL is listed in the interface.

Make sure you use the same certificate authority used for the server certificate and that you have uploaded the intermediate certificate for the certificates. For more information, see [Uploading Server Certificates](#) on page 2004.

To require valid user certificates:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.
2. Click **HTTPS Certificate**.
The Current HTTPS Certificate page appears.
3. Select **Enable User Certificates**.
The Enable Fetching of CRL option appears.
4. Optionally, select **Enable Fetching of CRL**.
The remaining CRL configuration options appear.

The screenshot shows the 'HTTPS User Certificate Settings' configuration page. It includes the following elements:

- Enable User Certificates:** A checkbox that is checked.
- Enable Fetching of CRL:** A checkbox that is checked.
- Enter URL for CRL:** A text input field with a 'Refresh CRL' button to its right.
- Scheduled CRL Download:** A text area containing the message: 'A scheduled task for downloading the CRL has been added, visit the scheduled task page to edit the task'.
- Last time CRL was fetched:** A text area containing the message: 'Not available'.
- Save:** A button at the bottom of the form.

5. Type a valid URL to an existing CRL file and click **Refresh CRL**.
The current CRL at the supplied URL loads to the server.

IMPORTANT! Enabling fetching of the CRL creates a scheduled task to update the CRL on a regular basis. Edit the task to set the frequency of the update. For more information, see [Automating Certificate Revocation List Downloads](#) on page 1934.

6. Verify that you have a valid user certificate generated by the same certificate authority that created the server certificate.

WARNING! When you save a configuration with enabled user certificates, if you do not have a valid user certificate installed in your browser certificate store, you disable all web server access to the appliance. Make sure you have a valid certificate installed before saving settings.

7. To apply the user certificate configuration to the web server, click **Save**.
Note that you can disable user certificate enforcement via the command line if you enable certificates and find that your user certificate does not enable access. For more information, see [disable-http-user-cert](#) on page 2280.

Enabling Access to the Database

LICENSE: Any

You can configure the Defense Center to allow read-only access to its database by a third-party client. This allows you to query the database using SQL using any of the following:

- industry-standard reporting tools such as Actuate BIRT, JasperSoft iReport, or Crystal Reports
- any other reporting application (including a custom application) that supports JDBC SSL connections
- the Sourcefire-provided command-line Java application called RunQuery, which you can either run interactively or use to obtain comma-separated results for a single query

From the Database Settings local configuration page, you can enable database access and create an access list that allows selected hosts to query the database. Note that this access list does not also control appliance access. For more information on appliance access lists, see [Configuring the Access List for Your Appliance](#) on page 1971.

You can also download a package that contains the following:

- RunQuery, the Sourcefire-provided database query tool
- InstallCert, a tool you can use to retrieve and accept the SSL certificate from the Defense Center you want to access
- the JDBC driver you must use to connect to the database

Note that when you connect to the database from an external client you must provide a username and password that match those for an Administrator or External Database user on the Defense Center. For more information, see [Adding New User Accounts](#) on page 1899.

For detailed information on configuring external access to the Sourcefire 3D System database, including information on the database schema and supported queries, see the *Sourcefire 3D System Database Access Guide*.

To enable database access:

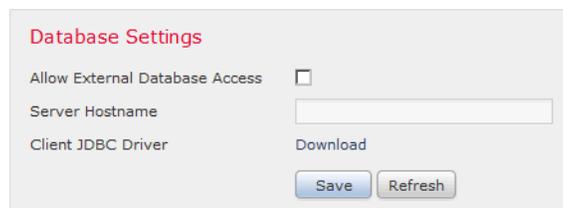
ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **Database**.

The Database Settings page appears.



3. Select the **Allow External Database Access** check box.
The **Access List** field appears. See step 6 for more information.

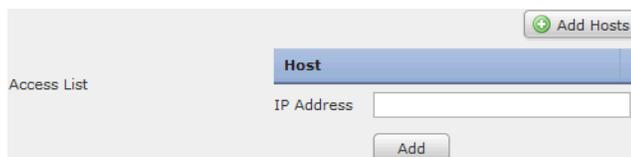
4. Type the fully qualified domain name (FQDN), IPv4 address, or IPv6 address of the Defense Center in the **Server Hostname** field, depending on your third-party application requirements.

If you type a FQDN, you must make sure that the client can resolve the FQDN of the Defense Center. If you type an IP address, you must make sure that the client can connect to the Defense Center using the IP address.

5. Next to **Client JDBC Driver**, click **Download** and follow your browser's prompts to download the `client.zip` package.

See the *Sourcefire 3D System Database Access Guide* for information on using the tools in the package you downloaded to configure database access.

- To add database access for one or more IP addresses, click **Add Hosts**.
An **IP Address** field appears in the **Access List** field.



The screenshot shows a web interface for configuring an 'Access List'. At the top right, there is a green button labeled 'Add Hosts'. Below it, a table with a blue header 'Host' is visible. The table has one column labeled 'IP Address' with an empty text input field. Below the table, there is a grey button labeled 'Add'.

- 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 range using CIDR notation (for example, 192.168.1.1/24)
For information on using CIDR in the Sourcefire 3D System, see [IP Address Conventions](#) on page 57.
 - any, to designate any IP address
- Click **Add**.
The IP address is added to the database access list.
- Optionally, to remove an entry in the database access list, click the delete icon ().
- Click **Save**.
Your database access settings are saved.

TIP! Click **Refresh** to revert to the last saved database settings.

Configuring Network Settings

LICENSE: Any

With some exceptions, your Sourcefire 3D System provides a dual stack implementation so that you can choose IPv4, IPv6, or IPv4 and IPv6 network settings.

TIP! For information on configuring network settings for virtual managed devices, which do not have a web interface, see the *Sourcefire 3D System Virtual Installation Guide*.

You can either manually configure network settings, or retrieve settings from a local DHCP server. For IPv6 deployments you can also retrieve settings from a local router. When you manually configure network settings, you must configure:

- the IP address of the management interface. In most deployments, the management interface is connected to an internal, protected network. This is the network through which Defense Centers and their managed devices communicate. For IPv4, you must set the address and netmask in dotted decimal form (for example: a netmask of 255.255.0.0). For IPv6, you must set the address in colon-separated hexadecimal form and the number of bits in the prefix (for example: a prefix length of 112).
- the IP address of the gateway device for your network.
- the DNS-resolvable name for the appliance. Note that if you change the hostname, the new name is not reflected in the syslog until after you reboot the appliance.
- the IP address of the primary DNS server for the network where the appliance resides. Optionally, specify a secondary and tertiary DNS servers.
- for Series 3 devices, whether you want to allow users to reconfigure network settings using the LCD panel.

By default, Sourcefire appliances are configured to directly connect to the internet; see [Security, Internet Access, and Communication Ports](#) on page 50. Specific features of the Sourcefire 3D System require this direct connection, and others support use of a proxy server, to which you can authenticate via HTTP Digest.

To configure network settings:

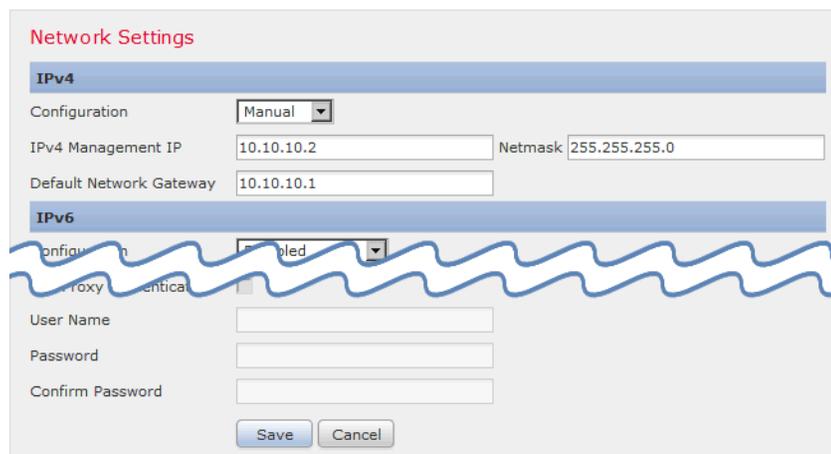
ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **Network**.

The following graphic shows the Network page for most Sourcefire appliances.



On a Series 3 device, you can also reconfigure the network using the LCD panel. Note that allowing reconfiguration using the LCD panel is a potential security issue because it provides the opportunity for someone to physically access the appliance and change the network configuration.



3. Specify whether you want to use IPv4, IPv6, or both.
- Select **Disabled** to use only the alternative IP version (for example, if your network uses only IPv6, in the IPv4 section select **Disabled**).
 - Select **DHCP** to allow DHCP server network setting resolution.
 - Select **Router assigned** (an IPv6-only configuration) to allow router assigned network setting resolution.
 - Select **Manual** to manually specify network settings, then specify settings as described above.
4. Optionally, in the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet, in bytes, that can pass through the appliance's management interface. The default value is 1500.
5. Optionally, identify a proxy server by selecting **Manual proxy configuration**. Enter the IP address or fully qualified domain name of your proxy server in the **HTTP Proxy** field and the port in the **Port** field. To supply authentication credentials for the connection to the proxy server, click **Use Proxy Authentication**. Supply a user name in the **User Name** field. Type a password in the **Password** field. Confirm the password in the **Confirm Password** field.

6. Optionally, if you want to use the LCD panel on your Series 3 device to reconfigure its network settings, select the **Allow reconfiguration of network settings** check box.
If you enable this, the web interface displays **Potential security issue** as a warning.
7. Click **Save**.
The network settings are changed.

Editing Management Interface Configurations

LICENSE: Any

You can use the Management Interface page to modify the default settings for each management interface on your Defense Center. Any changes you make to the Auto Negotiate value are ignored for Gigabit interfaces. Note that you can also configure management interfaces for a managed device from the managing device; see [Configuring the Management Interface](#) on page 260.

WARNING! Do not modify the settings for the management interface unless you have physical access to the appliance. It is possible to select a setting that makes it difficult to access the web interface.

To edit a management interface:

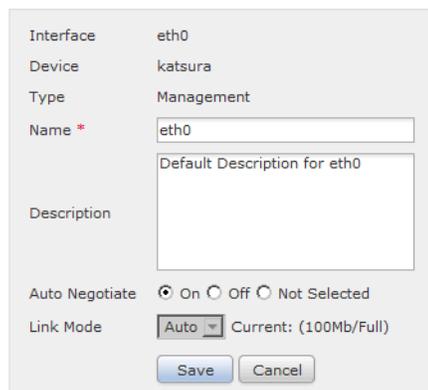
ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.
2. Click **Management Interface**.
The Management Interface page appears, listing the current settings for each interface on your Defense Center.

Management Interface Settings

Name	Description	Interface	Type	Link Mode	Device	
eth0	Default Description for eth0	eth0	Management	1Gb/Full (Auto)	bob.englab.sourcefire.com	 E

3. Click **Edit** next to the interface that you want to modify.
The current settings for the interface appear.



The screenshot shows a configuration window for the interface 'eth0'. The settings are as follows:

Interface	eth0
Device	katsura
Type	Management
Name *	eth0
Description	Default Description for eth0
Auto Negotiate	<input checked="" type="radio"/> On <input type="radio"/> Off <input type="radio"/> Not Selected
Link Mode	Auto (Current: (100Mb/Full))

At the bottom of the window are 'Save' and 'Cancel' buttons.

These settings include:

- interface
- device name
- interface type: Management.
- interface name
- interface description
- whether the interface is configured to auto-negotiate speed and duplex settings
- the current link mode, including the bandwidth and duplex setting (Full or Half); N/A indicates that there is no link for the interface

You can modify the interface name and description, MDI/MDIX settings, and the link mode as needed. However, keep the following in mind:

- In the **Auto Negotiate** field, select **Off** only if you require a specific link mode setting. You cannot change the Auto Negotiate setting for 10Gb interfaces.
- If **Auto Negotiate** is disabled and you need to specify a link mode, select it in the **Link Mode** field.
- Any changes you make to the **Auto Negotiate** value are ignored for Gigabit interfaces.

4. Click **Save**.
The Management Interface page appears again.

Shutting Down and Restarting the System

LICENSE: Any

You have several options for controlling the processes on your appliance. You can:

- shut down the appliance
- reboot the appliance
- restart communications, database, and HTTP server processes on the appliance (this is typically used during troubleshooting)
- restart the Snort process

IMPORTANT! If you shut down the appliance, the process shuts down the operating system on the appliance, but does not physically shut off power. To shut off power to the appliance, you must press the power button on the appliance, or, for an appliance without a power button, unplug it.

To shut down or restart your appliance:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **Process**.

The Appliance Process page appears. The Defense Center version of the page is shown below.

Name	
Shutdown Defense Center	Run Command
Reboot Defense Center	Run Command
Restart Defense Center Console	Run Command

3. Specify the command you want to perform:

On Defense Centers:

- To shut down the appliance, click **Run Command** next to **Shutdown Defense Center**.
- To reboot the appliance, click **Run Command** next to **Reboot Defense Center**. Note that this logs you out of the Defense Center.
- To restart the appliance, click **Run Command** next to **Restart Defense Center Console**. Note that restarting the Defense Center may cause deleted hosts to reappear in the network map.

On managed devices:

- To shut down the appliance, click **Run Command** next to **Shutdown Appliance**.
- To reboot the appliance, click **Run Command** next to **Reboot Appliance**. Note that this logs you out of the device.
- To restart the appliance, click **Run Command** next to **Restart Appliance Console**.
- To restart the Snort process, click **Run Command** next to **Restart Snort**.

Setting the Time Manually

LICENSE: Any

If the Time Synchronization setting in the currently applied system policy is set to **Manually in Local Configuration**, then you can manually set the time for the appliance using the Time page in the local configuration.

If the appliance is synchronizing its time based on NTP, you cannot change the time manually. Instead, the NTP Status section on the Time page provides the following information:

NTP Status

COLUMN	DESCRIPTION
NTP Server	The IP address and name of the configured NTP server.
Status	The status of the NTP server time synchronization. The following states may appear: <ul style="list-style-type: none">• Being Used indicates that the appliance is synchronized with the NTP server.• Available indicates that the NTP server is available for use, but time is not yet synchronized.• Not Available indicates that the NTP server is in your configuration, but the NTP daemon is unable to use it.• Pending indicates that the NTP server is new or the NTP daemon was recently restarted. Over time, its value should change to Being Used, Available, or Not Available.• Unknown indicates that the status of the NTP server is unknown.
Offset	The number of milliseconds of difference between the time on the appliance and the configured NTP server. Negative values indicate that the appliance is behind the NTP server, and positive values indicate that it is ahead.
Last Update	The number of seconds that have elapsed since the time was last synchronized with the NTP server. The NTP daemon automatically adjusts the synchronization times based on a number of conditions. For example, if you see larger update times such as 300 seconds, that indicates that the time is relatively stable and the NTP daemon has determined that it does not need to use a lower update increment.

See [Synchronizing Time](#) on page 1991 for more information about the time settings in the system policy.

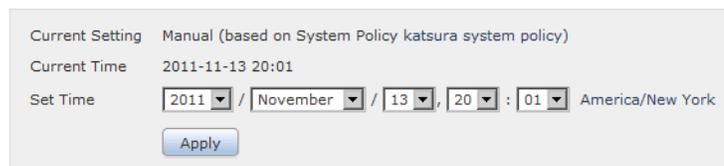
To manually configure the time:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.

2. Click **Time**.

The Time page appears.



Current Setting Manual (based on System Policy katsura system policy)
Current Time 2011-11-13 20:01
Set Time 2011 / November / 13, 20 : 01 America/New York
Apply

3. Select the following from the **Set Time** drop-down lists:

- year
- month
- day
- hour
- minute

4. Click **Apply**.

The time is updated.

5. If you want to change the time zone, click the time zone link located next to the date and time.

A pop-up window appears.

6. From the left list, select the continent or area that contains the time zone you want to use.

For example, if you want to use a time zone standard to North America, South America, or Canada, select **America**.

7. From the right list, select the zone (city name) that corresponds with the time zone you want to use.

For example, if you want to use Eastern Standard Time, you would select **New York** after selecting **America** in the first time zone box.

8. Click **Save** and, after the time zone setting is saved, click **Done** to close the pop-up window.

For more information about using the time zone page, see [Setting Your Default Time Zone](#) on page 2222.

Managing Remote Storage

LICENSE: Any

On Defense Centers, you can use local or remote storage for backups and reports. You can use Network File System (NFS), Secure Shell (SSH), or Server Message Block (SMB)/Common Internet File System (CIFS) for backup and report remote storage. You cannot send backups to one remote system and reports to another, but you can choose to send either to a remote system and store the

other on the local Defense Center. For information on backup and restore, see [Using Backup and Restore](#) on page 2204.

TIP! After configuring and selecting remote storage, you can switch back to local storage **only** if you **have not** increased the connection database limit.

You must ensure that your external remote storage system is functional and accessible from the Defense Center.

Select one of the backup and report storage options:

- To disable external remote storage and use the local Defense Center for backup and report storage, see [Using Local Storage](#) on page 2018.
- To use NFS for backup and report storage, see [Using NFS for Remote Storage](#) on page 2019.
- To use SSH for backup and report storage, see [Using SSH for Remote Storage](#) on page 2020.
- To use SMB for backup and report storage, see [Using SMB for Remote Storage](#) on page 2022.

Using Local Storage

LICENSE: Any

You can store backups and reports on the local Defense Center.

To store backups and reports locally:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.
2. Click **Remote Storage Device**.
The Remote Storage Device page appears.



3. Select **Local (No Remote Storage)** from the **Storage Type** drop-down list.
4. Click **Save**.
Your storage location choice is saved.

TIP! You do not use the **Test** button with local storage.

Using NFS for Remote Storage

LICENSE: Any

You can select Network File System (NFS) protocol to store your reports and backups. Optionally, select the **Use Advanced Options** check box to use one of the mount binary options as documented in an NFS mount man page.

To store backups and reports using NFS:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.
2. Click **Remote Storage Device**.
The Remote Storage Device page appears.
3. Select **NFS** from the **Storage Type** drop-down list.
The page refreshes to display the NFS storage configuration options.

The screenshot shows the NFS storage configuration interface. At the top, 'Storage Type' is set to 'NFS'. Below this, the 'Connection' section contains three input fields: 'Host' (labeled 'IP or hostname'), 'Directory', and 'Domain' (labeled 'MSHOME'). The 'Advanced' section has a 'Use Advanced Options' checkbox. The 'System Usage' section includes 'Use for Backups' and 'Use for Reports' checkboxes, and a 'Disk Space Threshold' field set to '85 %'. 'Save' and 'Test' buttons are located at the bottom of the form.

4. Add the connection information:
 - Enter the IPv4 address or hostname of the storage system in the **Host** field.
 - Enter the path to your storage area in the **Directory** field.
 - Optionally, enter the domain name for the remote storage system in the **Domain** field.
5. If there are any required command line options, select **Use Advanced Options**.
A **Command Line Options** field appears where you can enter mount binary options.

6. Under **System Usage**, select either or both of the following:
 - Select **Use for Backups** to store backups on the designated host.
 - Select **Use for Reports** to store reports on the designated host.
 - Enter **Disk Space Threshold** for backup to remote storage. Default is 85%.
7. Optionally, click **Test**.

The test ensures that the Defense Center can access the designated host and directory.
8. Click **Save**.

Your remote storage configuration is saved.

Using SSH for Remote Storage

LICENSE: Any

You can select Secure Shell (SSH) protocol to store your reports and backups. Optionally, select the **Use Advanced Options** check box to use one of the mount binary options as documented in a SSH mount man page.

To store backups and reports using SSH:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.
2. Click **Remote Storage Device**.

The Remote Storage Device page appears.

3. At **Storage Type**, select **SSH**.

The page refreshes to display the SSH storage configuration options.

Storage Type: SSH

Connection

Host: [] IP or hostname

Directory: []

Domain: [] MSHOME

Username: []

Password: []

SSH Public Key: ssh-rsa AAAAB3NzaC1yc2EAAAAD
To use ssh keys place this public key in your authorized_keys file.

Advanced

Use Advanced Options:

System Usage

Use for Backups:

Use for Reports:

Disk Space Threshold: 85 %

Save Test

4. Add the connection information:

- Enter the IP address or hostname of the storage system in the **Host** field.
- Enter the path to your storage area in the **Directory** field.
- Optionally, enter the domain name for the remote storage system in the **Domain** field.
- Enter the storage system's user name in the **Username** field and the password for that user in the **Password** field.
- To use SSH keys, copy the content of the **SSH Public Key** field and place it in your authorized_keys file.

5. If there are any required command line options, select **Use Advanced Options**.

A **Command Line Options** field appears where you can enter mount binary options.

6. Under System Usage, select either or both of the following:

- Select **Use for Backups** to store backups on the designated host.
- Select **Use for Reports** to store reports on the designated host.

7. Optionally, click **Test**.
The test ensures that the Defense Center can access the designated host and directory.
8. Click **Save**.
Your remote storage configuration is saved.

Using SMB for Remote Storage

LICENSE: Any

You can select Server Message Block (SMB) protocol to store your reports and backups. Optionally, select the **Use Advanced Options** check box to use one of the mount binary options, as documented in an SMB mount man page. For example, using SMB, you can enter the security mode in the **Command Line Options** field using the following format:

sec=mode

where *mode* is the security mode you want to use for remote storage. See the [Security Mode Settings](#) table for setting options.

Security Mode Settings

MODE	DESCRIPTION
[none]	Attempt to connect as null user (no name).
krb5	Use Kerberos version 5 authentication.
krb5i	Use Kerberos authentication and packet signing.
ntlm	Use NTLM password hashing. (Default)
ntlmi	Use NTLM password hashing with signing (may be Default if <code>/proc/fs/cifs/PacketSigningEnabled</code> is on or if server requires signing).
ntlmv2	Use NTLMv2 password hashing.
ntlmv2i	Use NTLMv2 password hashing with packet signing.

To store backups and reports using SMB:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.

2. Click **Remote Storage Device**.

The Remote Storage Device page appears.

3. Under **Storage Type**, select **SMB**.

The page refreshes to display the SMB storage configuration options.

The screenshot shows the SMB storage configuration interface. At the top, 'Storage Type' is set to 'SMB'. Below this, the 'Connection' section contains five input fields: 'Host' (with a hint 'IP or hostname'), 'Share', 'Domain' (with a hint 'MSHOME'), 'Username', and 'Password'. The 'Advanced' section features a 'Use Advanced Options' checkbox. The 'System Usage' section includes 'Use for Backups' and 'Use for Reports' checkboxes, and a 'Disk Space Threshold' field set to '85 %'. At the bottom, there are 'Save' and 'Test' buttons.

4. Add the connection information:

- Enter the IPv4 address or hostname of the storage system in the **Host** field.
- Enter the share of your storage area in the **Share** field. Note that the system only recognizes top-level shares and not full file paths. To use the specified Share directory as a remote backup destination, it must be shared on the Windows system.
- Optionally, enter the domain name for the remote storage system in the **Domain** field.
- Enter the user name for the storage system in the **Username** field and the password for that user in the **Password** field.

5. If there are any required command line options, select **Use Advanced Options**.

A **Command Line Options** field appears where you can enter the mount binary commands, such as security modes. See [Security Mode Settings](#) on page 2022 for more information.

6. Under System Usage, select either or both of the following:

- Select **Use for Backups** to store backups on the designated host.
- Select **Use for Reports** to store reports on the designated host.

7. Optionally, click **Test**.
The test ensures that the Defense Center can access the designated host and directory.
8. Click **Save**.
Your remote storage configuration is saved.

Understanding Change Reconciliation

LICENSE: Any

To monitor the changes that users make and ensure that they follow your organization's preferred standard, you can configure your system to send, via email, a detailed report of changes made to your system over the past 24 hours. Whenever a user saves changes to the system configuration, a snapshot is taken of the changes. The change reconciliation report combines information from these snapshots to present a clear summary of recent system changes.

The following sample graphic displays a User section of an example change reconciliation report and lists both the previous value for each configuration and the value after changes. When users make multiple changes to the same configuration, the report lists summaries of each distinct change in chronological order, beginning with the most recent.

6 User - SampleUser

6.1 User (2011-03-29 12:42:17 by admin from 10.4.4.4)

Field	Previous Value	Current Value
Name	SampleUser	
Active	Enabled	
Authentication	SHA512	
Password	*****	
Maximum Number of Failed Logins	5	
Days Until Password Expiration	Unlimited	
Days Until Expiration Warning	0	
Check Password Strength	No	
Roles		
	Administrator	

6.2 User (2011-03-29 12:42:12 by admin from 10.4.4.4)

Field	Previous Value	Current Value
Name		SampleUser
Active		Enabled

You can view changes made during the previous 24 hours. However, to view prior changes, you must view the audit log. See [Using the Audit Log to Examine Changes](#) on page 2196 for more information.

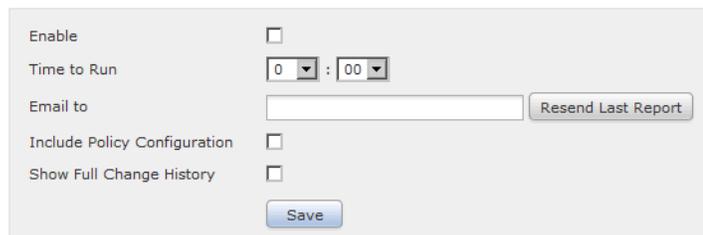
To use the change reconciliation feature:

ACCESS: Admin

1. Select **System > Local > Configuration**.
The Information page appears.

2. Click **Change Reconciliation**.

The Change Reconciliation page appears.



3. Select the **Enable** check box.
4. Select the time of day you want the system to send out the change reconciliation report from the **Time to Run** drop-down lists.
5. In the **Email to** field, enter the email addresses of report recipients. At any time, you can click **Resend Last Report** to send recipients another copy of the most recent change reconciliation report.

IMPORTANT! To receive change reconciliation reports, you must first configure a mail relay host and notification address. For more information, see [Configuring a Mail Relay Host and Notification Address](#) on page 1983.

6. Optionally, select **Include Policy Configuration** to include records of policy changes in the change reconciliation report. This includes changes to access control, intrusion, system, health, and network discovery policies. If you do not select this option, the report will not show changes to any policies.

IMPORTANT! This option is not available on managed devices.

7. Optionally, select **Show Full Change History** to include records of all changes over the past 24 hours in the change reconciliation report. If you do not select this option, the report includes only a consolidated view of changes for each category.
8. Click **Save**.

Your changes are saved. The report runs daily at the time you selected.

Managing Remote Console Access

LICENSE: Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

You can use a Linux system console for remote access on any appliance via either the VGA port (which is the default) or the serial port on the physical appliance.

Choose the option most suitable to the physical layout of your organization's Sourcefire deployment.

You can use Lights-Out Management (LOM) on a Serial Over LAN (SOL) connection to remotely monitor or manage Series 3 appliances without logging into the management interface of the appliance. You can perform limited tasks, such as viewing the chassis serial number or monitoring such conditions as fan speed and temperature, using a command line interface on an out-of-band management connection. Series 2 and virtual appliances do not support LOM.

You must enable LOM for both the appliance and the user you want to manage the appliance. After you enable the appliance and the user, you use a third-party Intelligent Platform Management Interface (IPMI) utility to access and manage your appliance.

For more information, see the following topics:

- [Configuring Remote Console Settings on the Appliance](#) on page 2026
- [Enabling Lights-Out Management User Access](#) on page 2028
- [Using a Serial Over LAN Connection](#) on page 2029
- [Using Lights-Out Management](#) on page 2030

Configuring Remote Console Settings on the Appliance

LICENSE: Any

SUPPORTED DEVICES: feature dependent

SUPPORTED DEFENSE CENTERS: feature dependent

Use the web interface of the appliance you want to remotely manage to select and configure the remote console access option you want to use.

IMPORTANT! Before you can connect to a 7000 Series device using LOM/SOL, you must disable Spanning Tree Protocol (STP) on any third-party switching equipment connected to the device's management interface.

To configure remote console settings:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Select **Console Configuration**.

The Console Configuration page appears.

Console Configuration

VGA
 Physical Serial Port

Lights-Out Management Settings

IPv4 Settings

Configuration: Disabled

IP Address:

Netmask:

Default Gateway:

Lights-Out Management Users

Username	Status	Action
admin	Access Granted	Edit
jkpolk	Access Granted	Edit

Save Refresh

3. Select a remote console access option:

- Select **VGA** to use the appliance's VGA port. This is the default option.
- Select **Physical Serial Port** to use the appliance's serial port, or to use LOM/SOL on a Series 3 Defense Center or 8000 Series device.
- Select **Lights-Out Management** to use LOM/SOL on a 7000 Series device. On these devices, you cannot use SOL and a regular serial connection at the same time.

If you selected **Physical Serial Port** or **Lights-Out Management**, the LOM settings appear.

IMPORTANT! When you change your remote console from **Physical Serial Port** to **Lights-Out Management** or from **Lights-Out Management** to **Physical Serial Port** on the 70xx Family of devices, you may have to reboot the appliance twice to see the expected boot prompt.

4. To configure LOM via SOL, enter the appropriate settings:

- DHCP **Configuration** for the appliance (**DHCP** or **Static**)
- **IP Address** to be used for LOM

IMPORTANT! The LOM IP address must be different from the management interface IP address of the appliance.

- the **Netmask** for the appliance
 - the **Default Gateway** for the appliance
5. Click **Save**.

Remote console configuration for the appliance is saved. If you configured Lights-Out Management, you must enable it for at least one user; see [Enabling Lights-Out Management User Access](#) on page 2028.

Enabling Lights-Out Management User Access

LICENSE: Any

SUPPORTED DEVICES: Series 3

SUPPORTED DEFENSE CENTERS: Series 3

You must explicitly grant Lights-Out Management permissions to users who will use the feature. You configure LOM and LOM users on a per-appliance basis using each appliance's local web interface. That is, you cannot use the Defense Center to configure LOM on a managed device. Similarly, because users are managed independently per appliance, enabling or creating a LOM-enabled user on the Defense Center does not transfer that capability to users on managed devices.

LOM users also have the following restrictions:

- You must assign the Administrator role to the user.
- The username may have up to 16 alphanumeric characters. Hyphens and longer user names are not supported for LOM users.
- The password may have up to 20 alphanumeric characters. Longer passwords are not supported for LOM users. A user's LOM password is the same as that user's system password.
- Series 3 Defense Centers and 8000 Series devices can have up to 13 LOM users. 7000 Series devices can have up to eight LOM users.

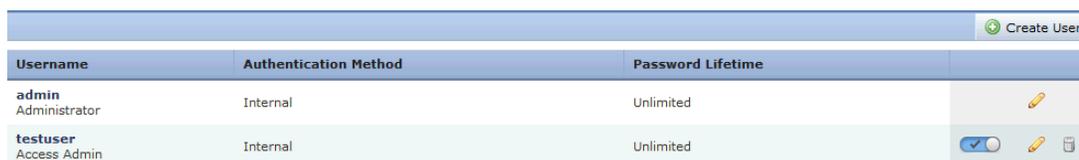
Note that if you deactivate, then reactivate, a role with LOM while a user with that role is logged in, or restore a user or user role from a backup during that user's login session, that user must log back into the web interface to regain access to IPMItool commands. For more information, see [Managing Predefined User Roles](#) on page 1906.

To enable or view Lights-Out Management user access:

ACCESS: Admin

1. Select **System > Local > User Management**.

The User Management page appears.



Username	Authentication Method	Password Lifetime	
admin Administrator	Internal	Unlimited	
testuser Access Admin	Internal	Unlimited	  

2. You have the following options:
 - To grant LOM user access to an existing user, click the edit icon () next to a user name in the list.
 - To grant LOM user access to a new user, click **Create User**.
3. Under User Configuration, enable the Administrator role.
Administrator Options appear.

Administrator Options Allow Lights-Out Management Access
4. Select the **Allow Lights-Out Management Access** check box.
5. Click **Save**.
The user has LOM access for this appliance.

Using a Serial Over LAN Connection

LICENSE: Any

SUPPORTED DEVICES: Series 3

SUPPORTED DEFENSE CENTERS: Series 3

You use a third-party IPMI utility on your computer to create a Serial Over LAN connection to the appliance. If your computer uses a Linux-like or Mac environment, use IPMItool; for Windows environments, use IPMIutil.

Linux

IPMItool is standard with many distributions and is ready to use.

Mac

You must install IPMItool on a Mac. First, confirm that your Mac has Apple's XCode Developer tools installed, making sure that the optional components for command line development are installed ("UNIX Development" and "System Tools" in newer versions, or "Command Line Support" in older versions). Then you can install macports and the IPMItool. Use your favorite search engine for more information or try these sites:

<https://developer.apple.com/technologies/tools/>

<http://www.macports.org/>

Windows

You must compile IPMIutil on Windows. If you do not have access to a compiler, you can use IPMIutil itself to compile. Use your favorite search engine for more information or try this site:

<http://ipmiutil.sourceforge.net/>

Understanding IPMI Utility Commands

Commands used for IPMI utilities are composed of segments as in the following IPMItool example:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

where:

- `ipmitool` invokes the utility
- `-I lanplus` enables encryption for the session
- `-H IP_address` indicates the IP address of the appliance you want to access
- `-U user_name` is the name of an authorized user
- `-command` is the name of the command you want to give

The same command for Windows looks like this:

```
ipmiutil command -v 4 -j 3 -N IP_address -U user_name
```

This command connects you to the command line on the appliance where you can log in as if you were physically present at the appliance. You may be prompted to enter a password.

To create a Serial Over LAN connection:

ACCESS: Admin with LOM access

► Enter the following command:

For IPMItool:

```
ipmitool -I lanplus -H IP_address -U user_name sol activate
```

For IPMIutil:

```
ipmiutil -j 3 -H IP_address -U username sol -a
```

The command line login for the appliance appears. You may be prompted to enter a password.

Using Lights-Out Management

LICENSE: Any

SUPPORTED DEVICES: Series 3

SUPPORTED DEFENSE CENTERS: Series 3

Lights-Out Management provides the ability to perform a limited set of actions over a SOL connection without the need to log into the appliance. You use the `sol activate` command to create a SOL connection followed by one of the commands listed in the [Lights-Out Management Commands](#) table. After the command is completed,

the connection ends.

WARNING! In rare cases, if your computer is on a different subnet than the appliance's management interface and the appliance is configured for DHCP, attempting to access LOM features on a Series 3 appliance can fail. If this occurs, you can either disable and then re-enable LOM on the appliance, or use a computer on the same subnet as the appliance to ping its management interface. You should then be able to use LOM.

If all attempts to access your appliance have failed, you can use LOM to restart your appliance remotely. Note that if a system is restarted while the SOL connection is active, the LOM session may disconnect or time out.

WARNING! Do **not** restart your appliance unless it does not respond to any other attempts to restart. Remotely restarting the appliance does not gracefully reboot the system and you may lose data.

Lights-Out Management Commands

IPMItool	IPMIUTIL	DESCRIPTION
(not applicable)	-V 4	Enables admin privileges for the IPMI session
-I lanplus	-J 3	Enables encryption for the IPMI session
-H	-N	Indicates the IP address of the remote appliance
-U	-U	Indicates the username of an authorized LOM account
sol activate	sol -a	Starts the SOL session
sol deactivate	sol -d	Ends the SOL session
chassis power cycle	power -c	Restarts the appliance
chassis power on	power -u	Powers up the appliance
chassis power off	power -d	Powers down the appliance
sdr	sensor	Displays appliance information, such as fan speeds and temperatures

For example, to display a list of appliance information, the IPMItool command is:

```
ipmitool -I lanplus -H IP_address -U user_name sdr
```

The same command with the IPMIutil utility is:

```
ipmiutil sensor -v 4 -j 3 -N IP_address -U user_name
```

To use Lights-Out Management:

ACCESS: Admin with LOM access

► Enter the following command:

For IPMITool:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

For IPMIutil:

```
ipmiutil -j 3 -H IP_address -U username command
```

where *command* is one of the commands from the [Lights-Out Management Commands table](#) on page 2031.

The corresponding action as noted in the table is performed. You may be prompted to enter a password.

Enabling Sourcefire Cloud Communications

LICENSE: URL Filtering or Malware

SUPPORTED DEFENSE CENTERS: Any except DC500

The Sourcefire 3D System contacts the Sourcefire cloud to obtain various types of information:

- If your organization has a FireAMP subscription, you can receive endpoint-based malware events; see [Working with Sourcefire Cloud Connections for FireAMP](#) on page 1219.
- File policies associated with access control rules allow managed devices to detect files transmitted in network traffic. The Defense Center uses data from the Sourcefire cloud to determine if the files represent malware; see [Understanding and Creating File Policies](#) on page 1206.
- When you enable URL filtering, the Defense Center 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 [Adding URL Conditions](#) on page 506.

Use the Defense Center'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 Sourcefire 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 Defense Center 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 Sourcefire cloud, for example, for privacy reasons.

Enable Automatic Updates

Allows the system to contact the Sourcefire 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 Defense Center 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](#) on page 1955.

IMPORTANT! Sourcefire 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.

Share IP Address and URI Information of malware events with Sourcefire

Optionally, Defense Centers can send information about the files detected by their managed devices to the Sourcefire cloud. This information includes IP addresses and URI information associated with detected files, as well as their SHA-256 hash values. Although sharing is opt-in, transmitting this information to Sourcefire will help with future efforts to identify and track malware.

Licensing

Performing category and reputation-based URL filtering and device-based malware detection require that you enable the appropriate licenses on your managed devices; see [Licensing the Sourcefire 3D System](#) on page 2036.

You **cannot** configure cloud connection options if you have no URL Filtering or Malware licenses on the Defense Center. If you have one license but not the other, the Cloud Services local configuration page displays only the options for which you are licensed. Defense Centers with expired licenses cannot contact the cloud.

Note that receiving endpoint-based malware events using a FireAMP subscription does not require a license, nor does specifying individual URLs or groups of URLs to allow or block. For more information, see [Understanding Malware Protection and File Control](#) on page 1199 and [Adding URL Conditions](#) on page 506.

Internet Access and High Availability

The system uses ports 80/HTTP and 443/HTTPS to contact the Sourcefire cloud for URL filtering data, and port 32137 to share malware data with the cloud. You must open the appropriate ports, both inbound and outbound for 80/HTTP and 443/HTTPS and outbound only for 32137, on the Defense Center.

To share malware data with the cloud, the Defense Center must have direct access to the internet, without a proxy. However, you can configure a proxy to download URL filtering data; see [Configuring Network Settings](#) on page 2009.

Although all URL filtering configurations and information are synchronized between Defense Centers in a high availability deployment, only the primary Defense Center downloads URL filtering data. If the primary Defense Center fails, you must not only make sure that the secondary Defense Center has direct access to the internet, but also must use the web interface on the secondary Defense Center to switch its role to **Active & Primary**. For more information, see [Monitoring and Changing High Availability Status](#) on page 202.

On the other hand, although they share file policies and related configurations, Defense Centers in a high availability pair share neither cloud connections nor malware dispositions. To ensure continuity of operations, and to ensure that detected files' malware dispositions are the same on both Defense Centers, both primary and secondary Defense Centers must have access to the cloud.

Health Monitoring

The default health policy includes the following modules that track the state and stability of the Defense Center's cloud connections:

- URL Filtering Monitor, which also warns you if the Defense Center fails to push category and reputation updates to its managed devices
- Advanced Malware Protection

TIP! Another module, the FireAMP Status Monitor, tracks the Defense Center's connection to the Sourcefire cloud for FireAMP subscription holders. For more information on health monitoring, see [Using the Health Monitor](#) on page 2163.

The following procedures explain how to enable communications the Sourcefire 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 Sourcefire cloud:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **Cloud Services**.

The Cloud Services page appears. If you have a URL Filtering license, the page displays the last time URL data was updated.



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**.

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 Defense Center retrieves URL filtering data from the cloud.

To perform an on-demand update of the system's URL data:

ACCESS: Admin

1. Select **System > Local > Configuration**.

The Information page appears.

2. Click **URL Filtering**.

The URL Filtering page appears.

3. Click **Update Now**.

The Defense Center contacts the cloud and updates its URL filtering data if an update is available.

CHAPTER 50

LICENSING THE SOURCEFIRE 3D SYSTEM

You can license a variety of features to create an optimal Sourcefire 3D System deployment for your organization. You use the Defense Center to manage licenses for itself and the devices it manages.

For more information, see:

- [Understanding Licensing](#) on page 2036
- [Viewing Your Licenses](#) on page 2048
- [Adding a License to the Defense Center](#) on page 2050
- [Deleting a License](#) on page 2052
- [Changing a Device's Licensed Capabilities](#) on page 2052

Understanding Licensing

LICENSE: Any

You can license a variety of features to create an optimal Sourcefire 3D System deployment for your organization. A FireSIGHT license is included with your Defense Center and is required to perform host, application, and user discovery.

Additional model-specific licenses allow your managed devices 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

- switching and routing
- device clustering
- network address translation (NAT)
- virtual private network (VPN) deployments

There are a few ways you may lose access to licensed features in the Sourcefire 3D System. You can remove licenses from the Defense Center, which affects all of its managed devices. You can also disable licensed capabilities on specific managed devices. Finally, some licenses may expire. Though there are some exceptions, you cannot use the features associated with an expired or deleted license.

For more information, see:

- [License Types and Restrictions](#) on page 2037
- [Licensing High Availability Pairs](#) on page 2045
- [Licensing Stacked and Clustered Devices](#) on page 2045
- [Licensing Series 2 Appliances](#) on page 2045
- [Understanding FireSIGHT Host and User License Limits](#) on page 2045

License Types and Restrictions

LICENSE: Any

This section describes the types of licenses available in a Sourcefire 3D System deployment. The licenses you can enable on an appliance depend on its model, version, and (for managed devices) the other licenses enabled.

For virtual and Series 3 devices, licenses are model specific; you cannot enable a license on a managed device unless the license exactly matches the device's model. For example, you cannot use a 3D8250 Protection license to enable Protection capabilities on a 3D8140 device. As your organization and deployment grow, you can purchase additional licenses for additional managed devices.

Series 2 devices automatically have Protection capabilities (with the exception of Security Intelligence filtering). Although you do not need to explicitly enable Protection on Series 2 devices, you also cannot enable any other licenses.

Also note that although you can enable Control on a virtual device to perform user and application control, virtual devices do not support switching, routing, or clustering.

The following table summarizes Sourcefire 3D System licenses.

Sourcefire 3D System Licenses

LICENSE	PLATFORMS	GRANTED CAPABILITIES	REQUIRES
FireSIGHT	Defense Center	discovery	none
RNA Host and RUA User (legacy)	Defense Center	discovery	none
Protection (licensed)	Series 3 and virtual devices	intrusion detection and prevention file control Security Intelligence filtering	none
Protection (automatic)	Series 2 devices	intrusion detection and prevention file control	none
Control	virtual devices	user and application control	Protection
Control	Series 3 devices	user and application control switching and routing clustering	Protection
Malware	Series 3 and virtual devices	advanced malware protection (network-based malware detection and blocking)	Protection
URL Filtering	Series 3 and virtual devices	category and reputation-based URL filtering	Control
VPN	Series 3 devices	deploying virtual private networks	Control

Note that the DC500 Defense Center does not support the capabilities provided by a URL Filtering or Malware license; note also that on a Series 2 device, you can only enable Protection capabilities.

For more information, see:

- [FireSIGHT](#) on page 2039
- [RNA Host and RUA User](#) on page 2039
- [Protection](#) on page 2041
- [Control](#) on page 2042
- [Malware](#) on page 2043

- [URL Filtering](#) on page 2043
- [VPN](#) on page 2044

FireSIGHT

LICENSE: FireSIGHT

A FireSIGHT license is included with your Defense Center and allows you to perform host, application, and user discovery. Discovery data allows the system to create a complete, up-to-the-minute profile of your network, and correlate threat, endpoint, and network intelligence with user identity information. You can use discovery data to perform traffic profiling, assess network compliance, and implement correlation policies.

Your FireSIGHT license also determines how many individual hosts and users you can monitor with the Defense Center and its managed devices. Note that the user limit applies *independently* to the following:

- the Users database, which contains a record for each user detected by the Sourcefire 3D System
- the number of users you can use in access control rules to perform user control, also called *access-controlled users*

For information on the consequences of reaching the licensed limit, see [Understanding FireSIGHT Host and User License Limits](#) on page 2045.

Without a FireSIGHT license, you can still perform basic system configuration, monitoring, network-based access control (zone, network, VLAN, and port rule conditions), connection logging, and reporting. Additionally, you can receive endpoint-based malware events from the Sourcefire cloud without a FireSIGHT license, although your organization does need a FireAMP subscription.

TIP! The License statements in this guide assume your Defense Center has a FireSIGHT license. However, if the Defense Center was previously running Version 4.10.x, you may be able to use legacy RNA Host and RUA User licenses instead of a FireSIGHT license. For more information, see the next section, [RNA Host and RUA User](#).

RNA Host and RUA User

LICENSE: Custom

In Version 4.10.x of the Sourcefire 3D System, RNA Host and RUA User feature licenses determined your monitored host and user limits, respectively. If your Defense Center was previously running Version 4.10.x, you may be able to use your legacy host and user licenses instead of a FireSIGHT license.

Version 5.2 Defense Centers using legacy licenses use the RNA Host limit as the FireSIGHT host limit and the RUA User limit as both the FireSIGHT user and access-controlled user limit. The FireSIGHT Host License Limit health module alerts appropriately for your licensed limit; see [Understanding FireSIGHT Host and](#)

[User License Limits](#) on page 2045.

Note that RNA Host and RUA User limits are cumulative. That is, you can add multiple licenses of each type to the Defense Center to monitor the total number of hosts or users allowed by the licenses.

If you later add a FireSIGHT license, the Defense Center uses the higher of the limits. For example, the FireSIGHT license on the DC1500 supports up to 50,000 hosts and users. If the RNA Host limit on your Version 4.10.x DC1500 was higher than 50,000, using that legacy host license on the same Defense Center running Version 5.2 gives you the higher limit. For your convenience, the web interface displays only the licenses that represent the higher limits.

IMPORTANT! Because FireSIGHT licensed limits are matched to the hardware capabilities of Defense Centers, Sourcefire does **not** recommend exceeding them. For guidance, contact Sourcefire Support.

Because there is no update path from Version 4.10.x to Version 5.2, you must use an ISO file to “restore” a physical Defense Center. Similarly, you must install new versions of any virtual Defense Centers you want to use with legacy licenses. Note that Version 5.2 Defense Centers cannot manage Version 4.10.x devices. You can, however, restore and update Version 4.10.x devices to the latest version.

During the restore process on a physical Defense Center, you are prompted to delete license and network settings. Keep these settings, but if you accidentally delete them you can re-add them later. Because you reinstall rather than restore virtual Defense Centers, you cannot keep these settings.

Note that you restore or reinstall appliances to major versions of the Sourcefire 3D System. After you complete this process, Sourcefire recommends you also install any available patches or feature updates. For detailed information on the restore and reinstall processes, see the *Sourcefire 3D System Installation Guide* and the *Sourcefire 3D System Virtual Installation Guide*.

WARNING! Restoring or reinstalling an appliance results in the loss of all configuration and event data on the appliance. Consider backing up your appliance before you begin the process. Also, you must make sure the appliance supports the new version. The release notes list supported hardware and hosting environments.

Protection

LICENSE: Protection

SUPPORTED DEVICES: Series 3, virtual

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](#) on page 2043), you can also inspect and block a restricted set of those file types based on their malware dispositions.

IMPORTANT! Managed devices **must** be running at least Version 5.1.1 of the Sourcefire 3D System to perform file control.

- *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 Defense Center, then enable it on the devices targeted by the policy.

If you delete your Protection license from the Defense Center or disable Protection on managed devices, the Defense Center stops acknowledging intrusion and file events from the affected devices. As a consequence, correlation rules that use those events as a trigger criteria stop firing. Additionally, the Defense Center will not contact the internet for either Sourcefire-provided or third-party Security Intelligence information. You cannot reapply existing policies until you re-enable Protection.

Because a Protection license is required for Malware and Control licenses, deleting or disabling a Protection license has the same effect as deleting or disabling your Malware or Control license.

IMPORTANT! Series 2 devices automatically have most Protection capabilities; you do not have to purchase or enable Protection licenses for these devices. However, Series 2 devices cannot perform Security Intelligence filtering.

Control

LICENSE: Control

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: feature dependent

A Control license allows you to implement user and application control by adding user and application conditions to access control rules. It also allows you to configure your Series 3 managed devices to perform switching and routing (including DHCP relay and NAT), as well as cluster managed devices. To enable Control on a managed device, you must also enable Protection.

IMPORTANT! Although you can enable a Control license on a virtual device, virtual devices do **not** support switching, routing, or clustering.

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 Defense Center, then enable it on the devices targeted by the policy.

Note that the DC500 Defense Center does not support adding user conditions in access control rules.

Without a Control license, you cannot create switched, routed, or hybrid interfaces on your managed devices; create NAT entries; or configure DHCP relay for virtual routers. Although you can create virtual switches and routers, they are not useful without switched and routed interfaces to populate them. Further, you cannot apply a device configuration that includes switching or routing to a managed device where you have not enabled Control. Additionally, establishing clustering between managed devices requires that the devices are enabled for Control.

If you delete your Control license from the Defense Center or disable Control on individual devices, the affected devices do **not** stop performing switching or routing, nor do device clusters break. Although you can edit and delete existing configurations, you cannot apply your changes to the affected devices. You cannot add new switched, routed, or hybrid interfaces, nor can you add new NAT entries, configure DHCP relay, or establish device clustering. Finally, you cannot reapply existing access control policies if they include rules with user or application conditions.

Because a Control license is required for URL filtering, deleting or disabling a Control license has the same effect as deleting or disabling your URL Filtering license.

URL Filtering

LICENSE: URL Filtering

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

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 Sourcefire cloud by the Defense Center. To enable URL Filtering, you must also enable Protection and Control licenses.

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 Defense Center 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 Defense Center, then enable it on the devices targeted by the policy.

You may lose access to URL filtering if you delete the license from the Defense Center or disable URL Filtering on managed devices. Also, URL Filtering licenses may expire. If your license expires or if you delete or disable it, access control rules with URL conditions immediately stop filtering URLs, and your Defense Center 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

SUPPORTED DEVICES: Series 3, virtual

SUPPORTED DEFENSE CENTERS: Any except DC500

A Malware license allows you to perform advanced malware protection, that is, use managed devices to detect and block malware in files transmitted over your network. To enable Malware on a managed device, you must also enable Protection.

IMPORTANT! Managed devices **must** be running at least Version 5.1.1 of the Sourcefire 3D System to perform malware detection, and Version 5.2 to perform malware blocking. You cannot enable a Malware license on a pre-Version 5.1.1 device.

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.

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, then enable it on the devices targeted by the policy. If you later disable the license on the devices, you cannot reapply an existing access control policy to those devices if it includes file policies that perform malware detection.

If you delete all your Malware licenses or they all expire, the Defense Center stops performing malware cloud lookups, and also stops acknowledging retrospective events sent from the Sourcefire 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.

Note that a Malware license is only required if you want the system to detect malware in network traffic. Without a Malware license, the Defense Center can receive endpoint-based malware events from the Sourcefire cloud if your organization has a FireAMP subscription. For more information, see [Understanding Malware Protection and File Control](#) on page 1199.

VPN

LICENSE: VPN

SUPPORTED DEVICES: Series 3

VPN allows you to establish secure tunnels between endpoints via a public source, such as the Internet or other network. You can configure the Sourcefire 3D System to build secure VPN tunnels among the virtual routers of Sourcefire managed devices or from managed devices to remote devices or other third-party VPN endpoints. To enable VPN, you must also enable Protection and Control licenses.

Without a VPN license, you cannot configure a VPN deployment with your managed devices. Although you can create deployments, they are not useful without at least one VPN-enabled routed interface to populate them.

If you delete your VPN license from the Defense Center or disable VPN on individual devices, the affected devices do **not** break the current VPN deployments. Although you can edit and delete existing deployments, you cannot apply your changes to the affected devices.

Licensing High Availability Pairs

LICENSE: Any

SUPPORTED DEFENSE CENTERS: DC1000, DC1500, DC3000, DC3500

Defense Centers in a high availability pair do **not** share licenses. You must apply equivalent licenses to each member of the pair. Because Sourcefire generates licenses based on each Defense Center's unique license key, you cannot use the same licenses on different Defense Centers.

Licensing Stacked and Clustered Devices

LICENSE: Any

SUPPORTED DEVICES: feature dependent

Individual devices must have equivalent licenses before they can be stacked or clustered. After you stack devices, you can change the licenses for the entire stack. However, you cannot change the enabled licenses on a device cluster.

You can stack 3D8140, 3D8200 family, and 3D9900 devices of the same model that meet the requirements described in [Managing Stacked Devices](#) on page 237. You can cluster two devices of the same Series 3 model that meet the requirements described in [Clustering Devices](#) on page 220.

Licensing Series 2 Appliances

LICENSE: Protection

SUPPORTED DEVICES: Series 2

With the exception of the DC500, Series 2 and Series 3 Defense Center licensing is identical. Because the DC500 does not support URL filtering or network-based malware detection, it cannot take advantage of URL Filtering or Malware licenses.

Series 2 devices automatically have the capabilities, except for Security Intelligence, enabled by a Protection license. You cannot disable the Protection license on Series 2 devices, and you cannot enable other licenses.

See the following sections for more information:

- [License Types and Restrictions](#) on page 2037 describes the types of licenses available in a Sourcefire 3D System deployment.
- [Supported Capabilities by Appliance Model](#) on page 42 summarizes supported and unsupported features on Series 2 appliances.

Understanding FireSIGHT Host and User License Limits

LICENSE: FireSIGHT

The FireSIGHT license on your Defense Center determines how many individual hosts and users you can monitor with the Defense Center and its managed devices, as well as how many users you can use to perform user control.

FireSIGHT host and user license limits are model specific, as listed in the following table.

FireSIGHT Limits by Defense Center Model

DEFENSE CENTER MODEL	FIRE SIGHT HOST AND USER LIMIT
DC500	1000
DC750	2000
DC1000	20,000
DC1500	50,000
DC3000	100,000
DC3500	300,000
virtual	50,000

For example, you can monitor 1000 hosts and 1000 users with the DC500.

If your Defense Center was previously running Version 4.10.x of the Sourcefire 3D System and you used an ISO file to “restore” the appliance to Version 5.x factory defaults, you may be able to use your legacy RNA Host and RUA User licenses instead of a FireSIGHT license.

For more information, see the following sections:

- [Understanding the FireSIGHT Host Limit](#) on page 2046
- [Understanding the FireSIGHT User Limit](#) on page 2047
- [Understanding the Access-Controlled User Limit](#) on page 2048
- [RNA Host and RUA User](#) on page 2039

Understanding the FireSIGHT Host Limit

LICENSE: FireSIGHT

The FireSIGHT license on your Defense Center determines how many individual hosts you can monitor with the Defense Center and its managed devices, and therefore how many hosts you can store in your network map.

Note that the system counts MAC-only hosts separately from hosts identified by both IP addresses and MAC addresses. All IP addresses associated with a host are counted together as one host.

When the system detects activity associated with a host with an IP address in your monitored network (as defined by your network discovery policy), that host is added to the network map.

If you reach the host limit and the system detects a new host, whether the new host is added to the network map depends on the **When Host Limit Reached** setting in your network discovery policy. You can configure the system either to stop adding new hosts to the database, or to replace the hosts that have remained inactive for the longest time.

IMPORTANT! Even if you cannot add a new host to the network map, the system still performs access control on that host's network traffic. Although reaching the FireSIGHT host limit does not prevent you from performing access control on hosts discovered after you reached your licensed limit, you cannot view or perform analysis on those hosts using host profile data. For example, you cannot use compliance white lists to monitor network compliance for those hosts, or use those hosts in host profile qualifications, and so on.

You can also manually delete a host, an entire subnet, or all of your hosts from the network map. Keep in mind, however, that if the system detects activity associated with a deleted host, it re-adds the host to the network map.

Note also that if the system has not detected network traffic from a host in the last **Host Timeout** period specified in your network discovery policy, the host is removed from the network map. The default setting is 10080 minutes (7 days).

To help you track your host license use, the FireSIGHT Host License Limit health module warns you if you have fewer than a configurable number of host licenses left.

Understanding the FireSIGHT User Limit

LICENSE: FireSIGHT

The FireSIGHT license on your Defense Center determines how many individual users you can monitor. When the system detects activity from a new user, that user is added to the Users database. You can detect users in the following ways:

- You can use the network discovery policy to configure managed devices to passively detect logins for LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), and SMTP users.
- You can install Sourcefire User Agents on your Microsoft Active Directory LDAP servers to detect authentications against Active Directory credentials.

After you reach the licensed limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete users from the database, or purge all users from the database.

However, the system favors authoritative user logins. If you have reached the licensed limit and the system detects an authoritative user login for a previously

undetected user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new user.

TIP! Note that if you are using managed devices to detect user activity, you can restrict user logging by protocol to help minimize username clutter and preserve FireSIGHT user licenses. For example, monitoring users discovered via AIM, POP3, and IMAP may add users not relevant to your organization due to network access from contractors, visitors, and other guests. For more information, see [Restricting User Logging](#) on page 1285.

Understanding the Access-Controlled User Limit

LICENSE: Control

SUPPORTED DEVICES: Series 3, virtual

The FireSIGHT license on your Defense Center determines not only how many individual users you can monitor, but also how many users you can use in access control rules to perform user control. These users are called *access-controlled users*.

IMPORTANT! To perform user control, your organization **must** use Microsoft Active Directory. The system uses Sourcefire User Agents running on Active Directory servers to associate access-controlled users with IP addresses, which is what allows access control rules to trigger.

You specify the groups that access-controlled users must belong to by configuring a connection (called a *user awareness authentication object*) between the Defense Center and an Active Directory server. Then, on a regular basis, the Defense Center queries the server and retrieves a list of the users in the groups you specified in the authentication object. You can then use these users to perform access control.

You **must** make sure the total number of users in the groups you specify in the authentication object is less than your FireSIGHT user license. If your parameters are too broad, the Defense Center obtains information on as many users as it can and reports the number of users it failed to retrieve in the task queue. For performance and licensing reasons, Sourcefire recommends that you specify only the groups that represent the users you want to use in access control.

Viewing Your Licenses

LICENSE: Any

Use the Licenses page to view the licenses for a Defense Center and its managed devices. For each type of appliance in your deployment, the page lists

the total number of licenses you have as well as the portion of those licenses that are in use.

Maximum 3D8130 Licenses

Protection (Used)	5 (1)
Control (Used)	5 (1)
URL Filtering (Used)	5 (1)
Malware (Used)	5 (1)
VPN (Used)	5 (1)

Maximum DC1500 Licenses

FireSIGHT Host (Used)	50000 (898)
FireSIGHT User (Used)	50000 (129)

Keep in mind that on this page, the number of FireSIGHT User licenses in use represents the number of users detected by the Sourcefire 3D System, that is, the number of users in the Users database. It does not represent the number of access-controlled users you are using for access control. For more information, see [Understanding FireSIGHT Host and User License Limits](#) on page 2045.

The Licenses page also provides details on each of your licenses. For each model, you can see how many licenses of each type you have, and how many managed devices you can license with each type of license. For licenses that expire, the page provides you with the expiration date.

3D8250

License Type	Status	Number of Licenses	Expires	
URL Filtering	Valid License	5	2013-11-19 14:33:09	
Protection	Valid License	5	Never	
VPN	Valid License	5	Never	
Malware	Valid License	5	2013-11-19 14:32:03	
Control	Valid License	5	Never	

FireSIGHT DC3500

License Type	Status	Number of Licenses	Expires	
FireSIGHT Host FireSIGHT User	Valid License	300000 300000	Never	

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 Management page (**Devices > Device Management**) lists the licenses applied to each of your managed devices.
- Two health modules, License Monitor and FireSIGHT Host License Limit, communicate license status when used in a health policy.

To view your licenses:

ACCESS: Admin

- ▶ Select **System > Licenses**.
The Licenses page appears.

Adding a License to the Defense Center

LICENSE: Any

Before you add a license to the Defense Center, make sure you have the activation key provided by Sourcefire when you purchased the license. If you do not have the key, log on to the Sourcefire Support Site (<https://support.sourcefire.com/>) and check the **Entitlements** tab.

With the exception of FireSIGHT, you **must** enable licenses on your managed devices before you can use licensed features. You can enable a license either when you add a device to the Defense Center, or by editing the device's general properties after you add the device. Note that because Series 2 devices automatically have Protection capabilities, with the exception of Security Intelligence filtering, you cannot disable these capabilities, nor can you apply other licenses to a Series 2 device. See [Changing a Device's Licensed Capabilities](#) on page 2052.

To add a license:

ACCESS: Admin

1. Select **System > Licenses**.
The Licenses page appears.

2. Click **Add New License**.

The Add License page appears.

Add Feature License

License Key **66:00:00:77:FF:CC:88**

License

If your web browser cannot access the Internet, you must switch to a host with Internet access and navigate to <https://keyserver.sourcefire.com>.

Using the license key, **66:00:00:77:FF:CC:88**, follow the on-screen instructions to generate a license.

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 Licensing Center web site 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://keyserver.sourcefire.com/>.

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.

5. Copy the license from the email, paste it into the **License** field in the Defense Center's web interface, and click **Submit License**.

If the license is valid, it is added. You can now enable the license's capabilities on your managed devices, as described in [Changing a Device's Licensed Capabilities](#) on page 2052.

Deleting a License

LICENSE: Any

Use the following procedure if you need to delete a license for any reason. Keep in mind that because Sourcefire generates licenses based on each Defense Center's unique license key, you cannot delete a license from one Defense Center and then reuse it on a different Defense Center.

In most cases, deleting a license removes your ability to use features enabled by that license. For more information, see [License Types and Restrictions](#) on page 2037.

To delete a license:

ACCESS: Admin

1. Select **System > Licenses**.

The Licenses page appears.

2. Next to the license you want to delete, click the delete icon ().

Deleting a license removes the licensed capability from all devices using that license. For example, if your Protection license is valid for and enabled on 100 managed devices, deleting the license removes Protection capabilities from all 100 devices.

3. Confirm that you want to delete the license.

The license is deleted.

Changing a Device's Licensed Capabilities

LICENSE: Any

SUPPORTED DEVICES: Series 3, virtual

To change the licensed capabilities of a Series 3 or virtual device, edit the device's general properties on the Device Management page. Although there are some exceptions, you cannot use the features associated with a license if you disable it on a managed device.

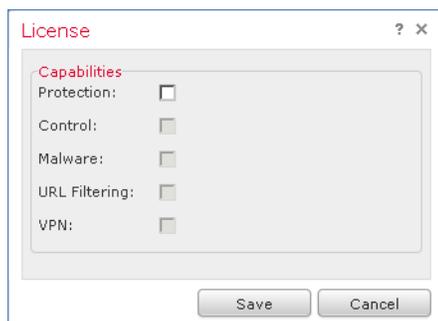
Series 2 devices automatically have Protection capabilities, with the exception of Security Intelligence filtering. You cannot disable these capabilities, nor can you apply other licenses to a Series 2 device. Note that, although you cannot use a Malware or URL Filtering license with a DC500 Defense Center, you can use a DC500 to enable or change these and other licensed capabilities of a Series 3 or virtual device.

For detailed information on the licenses you can enable, including version, model, and other requirements, see [License Types and Restrictions](#) on page 2037.

To enable or disable a device's licensed capabilities:

ACCESS: Admin/Network Admin

1. Select **Devices > Device Management**.
The Device Management page appears.
2. Next to the device where you want to enable or disable a license, click the edit icon (✎).
The Interfaces tab for that device appears.
3. Click **Device**.
The Device tab appears.
4. Next to the License section, click the edit icon (✎).
The License pop-up window appears.



5. Enable or disable the licensed capabilities of the device by clearing or selecting the appropriate check boxes.
6. Click **Save**.
The changes are saved but do not take effect until you apply the device configuration; see [Applying Changes to Devices](#) on page 211.

CHAPTER 51

UPDATING SYSTEM SOFTWARE

Sourcefire electronically distributes several different types of updates, including major and minor updates to the system software itself, as well as rule updates, geolocation database (GeoDB) updates, and Sourcefire Vulnerability Database (VDB) updates.

WARNING! This chapter contains general information on updating the Sourcefire 3D System. Before you update any part of the Sourcefire 3D System, 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 supported platforms, compatibility, prerequisites, warnings, and specific installation and uninstallation instructions.

Unless otherwise documented in the release notes or advisory text, updating an appliance does not modify its configuration; the settings on the appliance remain intact.

See the following sections for more information:

- [Understanding Update Types](#) on page 2055
- [Performing Software Updates](#) on page 2056
- [Uninstalling Software Updates](#) on page 2068
- [Updating the Vulnerability Database](#) on page 2070
- [Importing Rule Updates and Local Rule Files](#) on page 2072
- [Updating the Geolocation Database](#) on page 2092

Understanding Update Types

LICENSE: Any

Sourcefire electronically distributes several different types of updates, including major and minor updates to the system software itself, as well as intrusion rule updates and VDB updates.

The following table describes the types of updates provided by Sourcefire. For most update types, you can schedule their download and installation; see [Scheduling Tasks](#) on page 1929 and [Using Recurring Rule Updates](#) on page 2077.

Sourcefire 3D System Update Types

UPDATE TYPE	DESCRIPTION	SCHEDULE?	UNINSTALL?
patches to the Sourcefire 3D System	Patches include a limited range of fixes (and usually change the fourth digit in the version number; for example, 5.0.0.1).	X	X
feature updates to the Sourcefire 3D System	Feature updates are more comprehensive than patches and generally include new features (and usually change the third digit in the version number; for example, 5.0.1).	X	X
major updates (major and minor version releases) to the Sourcefire 3D System	Major updates, sometimes referred to as upgrades, include new features and functionality and may entail large-scale changes to the product (and usually change the first or second digit in the version number; for example, 5.0 or 5.1).		
VDB	VDB updates affect the vulnerabilities reported by the Sourcefire 3D System as well as the detected operating systems, applications, and clients.	X	
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 and provide new rule categories and system variables.	X	
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 must install the GeoDB to view geolocation details. The DC500 Defense Center does not support this feature.	X	

Note that while you can uninstall patches and other minor updates to the Sourcefire 3D System, you cannot uninstall major updates or return to previous versions of the VDB, GeoDB, or intrusion rules. If you updated your appliance to a new major version of the Sourcefire 3D System, and you need to revert to an older version, contact Support.

Performing Software Updates

LICENSE: Any

There are a few basic steps to updating your Sourcefire 3D System deployment. First, you must prepare for the update, including reading the release notes and completing any required pre-update tasks. Then, you can begin the update — first update your Defense Centers, then the devices they manage. You must monitor the update's progress until it completes, then verify the update's success. Finally, complete any required post-update steps.

For more information, see the following sections:

- [Planning for the Update](#) on page 2056
- [Understanding the Update Process](#) on page 2058
- [Updating a Defense Center](#) on page 2062
- [Updating Managed Devices](#) on page 2064
- [Monitoring the Status of Major Updates](#) on page 2066

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 [Sourcefire Support Site](#). The release notes describe supported platforms, new features and functionality, known and resolved issues, and product compatibility. 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.

Sourcefire 3D System Version Requirements

You must make sure your appliances (including software-based devices) are running the correct version of the Sourcefire 3D System. The release notes indicate the required version. If you are running an earlier version, you can obtain updates from the [Sourcefire Support Site](#).

Operating System Requirements

Make sure the computers where you installed software-based devices are running the correct versions of their operating systems. The release notes indicate the required versions. For information on supported operating systems for virtual devices, see the *Sourcefire 3D System Virtual Installation Guide*.

Time and Disk Space Requirements

Make sure you have enough free disk space and allow enough time for the update. When you update a managed device, the update requires additional disk space on the Defense Center. The release notes indicate space and time requirements.

Configuration and Event Backup Guidelines

Before you begin a major update, Sourcefire recommends that you delete any backups that reside on the appliance after copying them to an external location. Regardless of the update type, you should also back up current event and configuration data to an external location. Event data is **not** backed up as part of the update process.

You can use the Defense Center to back up event and configuration data for itself and the devices it manages; see [Using Backup and Restore](#) on page 2204.

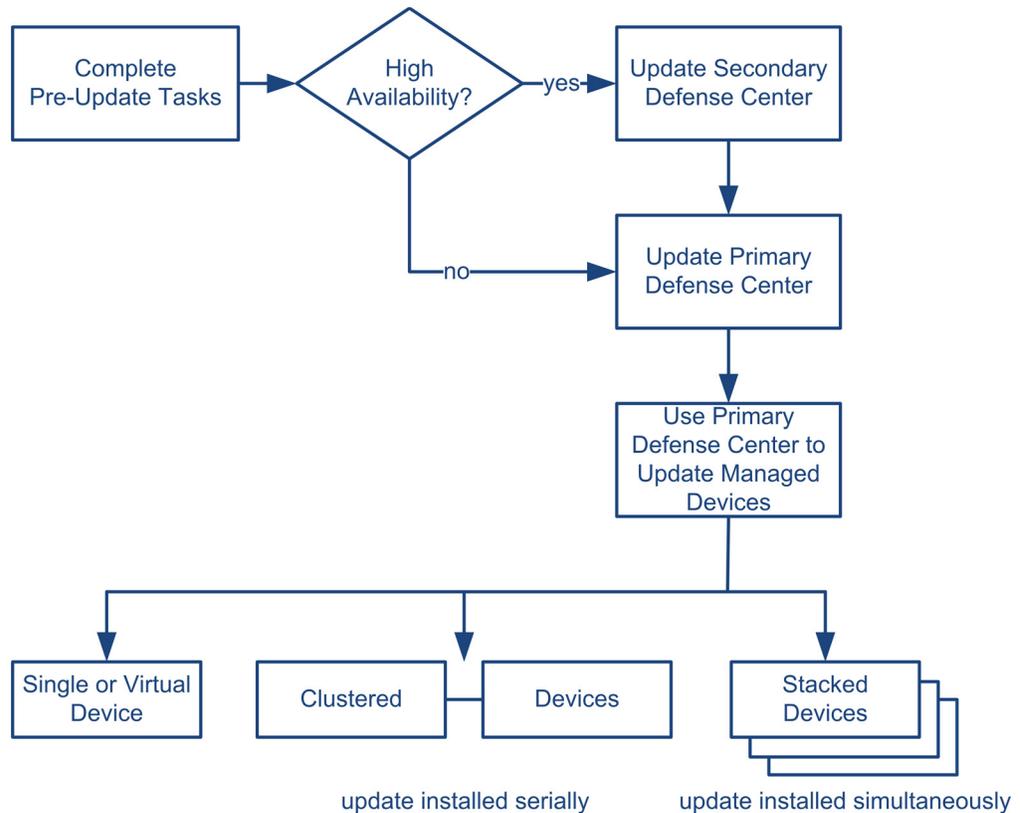
When to Perform the Update

Because the update process may affect traffic inspection, traffic flow, and link state, and because the Data Correlator is disabled while an update is in progress, Sourcefire recommends you perform the update in a maintenance window or at a time when the interruption will have the least impact on your deployment.

Understanding the Update Process

LICENSE: Any

The following diagram summarizes the update process.



Order of Update

You **must** update your Defense Centers before you can update the devices they manage.

Use the Defense Center to Perform the Update

Sourcefire recommends that you use the Defense Center's web interface to update not only itself, but also the devices it manages. You **must** use the Defense Center to update managed devices that do not have a web interface, such as virtual devices.

The Product Updates page (**System > Updates**) shows the version of each update, as well as the date and time it was generated. It also indicates whether a reboot is required as part of the update.

Product Updates Rule Updates Geolocation Updates Upload Update

Defense Center running software version: **5.0.0**

Updates

Type	Version	Date	Release Notes	Reboot
Sourcefire Vulnerability And Fingerprint Database Updates	72	Mon Nov 14 21:45:31 UTC 2011		No
Sourcefire 3D DC Sample Patch	5.1	Thu Nov 17 13:46:30 UTC 2011		Yes
Sourcefire 3D Device Sample Patch	5.1	Thu Nov 17 13:47:02 UTC 2011		Yes
Sourcefire 3D DC Sample Patch	5.0.1	Thu Nov 17 13:46:19 UTC 2011		Yes
Sourcefire 3D Device Sample Patch	5.0.1	Thu Nov 17 13:46:51 UTC 2011		Yes

Download updates

When you upload updates obtained from Sourcefire Support to your appliance, they appear on the page. Uninstallers for patch and feature updates also appear; see [Uninstalling Software Updates](#) on page 2068. On the Defense Center, the page can list VDB updates.

TIP! For patches and feature updates, you can take advantage of the automated update feature; see [Automating Software Updates](#) on page 1945.

Updating Paired Defense Centers

When you begin to update one Defense Center in a high availability pair, the other Defense Center in the pair becomes the primary, if it is not already. In addition, the paired Defense Centers stop sharing configuration information; paired Defense Centers do not receive software updates as part of the regular synchronization process.

To ensure continuity of operations, do **not** update paired Defense Centers at the same time. First, complete the update procedure for the secondary Defense Centers, then update the primary.

Updating Clustered Devices

When you install an update on clustered devices or clustered stacks, the system performs the update on the devices or stacks one at a time. When the update starts, the system first applies it to the backup device or stack, which goes into maintenance mode until any necessary processes restart and the device or stack is processing traffic again. The system then applies the update to the active device or stack, which follows the same process.

To update devices in a clustered stack, you must perform the update from the managing Defense Center on all members of a cluster at once; you cannot perform the upgrade directly from the devices.

Updating Stacked Devices

When you install an update on stacked devices, the system performs the updates simultaneously. Each device resumes normal operation when the update completes. Note that:

- If the primary device completes the update *before* all of the secondary devices, the stack operates in a limited, mixed-version state until all devices have completed the update.
- If the primary device completes the upgrade *after* all of the secondary devices, the stack resumes normal operation when the update completes on the primary device.

Traffic Flow and Inspection

When you install or uninstall updates from a managed device, 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, including switching, routing, and related functionality
- link state

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 the components of the Sourcefire 3D System that the update affects, how your devices are configured and deployed, and whether the update reboots the device. For specific information on how and when network traffic is affected for a particular update, see the release notes.

TIP! When you update clustered devices, the system performs the updates one at a time to avoid traffic interruption.

Using the Web Interface During the Update

Regardless of the type of update, do **not** use the web interface of the appliance you are updating to perform tasks other than monitoring the update.

To prevent you from using an appliance during a major update, and to allow you to easily monitor a major update's progress, the system streamlines the appliance's web interface. You can monitor a minor update's progress in the task queue (**System > Monitoring > Task Status**). Although you are not prohibited from using the

web interface during a minor update, Sourcefire recommends against it.

TIP! To monitor updates to its managed devices, use the task queue on the Defense Center.

Even for minor updates, the web interface on the updating appliance may become unavailable during the update process, or the appliance may log you out. This is expected behavior. If this occurs, log in again to view the task queue. If the update is still running, you **must** continue to refrain from using the web interface until the update has completed. Note that while updating, managed devices may reboot a second time; this is also expected behavior.

WARNING! If you encounter issues with the update (for example, if the web interface indicates that the update has failed or if a manual refresh of the task queue or 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, both after you update the Defense Center and then again after you update its managed devices. 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 [Applying an Access Control Policy](#) on page 466.

Additionally, you should:

- verify that the update succeeded
- make sure that all appliances in your deployment are communicating successfully
- 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 a Defense Center

LICENSE: Any

Update the Defense Center in one of two ways, depending on the type of update and whether your Defense Center has access to the Internet:

- You can use the Defense Center to obtain the update directly from the [Sourcefire Support Site](#), if your Defense Center has access to the Internet. This option is **not** supported for major updates.
- You can manually download the update from the [Sourcefire Support Site](#) and then upload it to the Defense Center. Choose this option if your Defense Center does not have access to the Internet or if you are performing a major update.

WARNING! To ensure continuity of operations, do **not** update paired Defense Centers at the same time; see [Updating Paired Defense Centers](#) on page 2059.

For major updates, updating the Defense Center removes uninstallers for previous updates.

To update the Defense Center:

ACCESS: Admin

1. Read the release notes and complete any required pre-update tasks.
Pre-update tasks may include making sure that: the Defense Center is running the correct version of the Sourcefire software, you have enough free disk space to perform the update, you set aside adequate time to perform the update, you backed up event and configuration data, and so on.
2. Upload the update to the Defense Center. You have two options, depending on the type of update and whether your Defense Center has access to the Internet:
 - For all except major updates, and if your Defense Center has access to the Internet, select **System > Updates**, then click **Download Updates** to check for the latest updates on the [Sourcefire Support Site](#).
 - For major updates, or if your Defense Center does not have access to the Internet, you must first manually download the update from the [Sourcefire Support Site](#). Select **System > Updates**, then click **Upload Update**. Browse to the update and click **Upload**.

IMPORTANT! Download the update directly from the [Sourcefire 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 to the Defense Center.

3. Make sure that the appliances in your deployment are successfully communicating and that there are no issues being reported by the health monitor.
4. Select **System > 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.

5. Select **System > Updates**.
The Product Updates page appears.
6. Click the install icon next to the update you uploaded.
The Install Update page appears.
7. Select the Defense Center and click **Install**. If prompted, confirm that you want to install the update and reboot the Defense Center.

The update process begins. How you monitor the update depends on whether the update is a major or minor update. See the [Sourcefire 3D System Update Types table](#) on page 2055 and the release notes to determine your update type:

- For minor updates, you can monitor the update's progress in the task queue (**System > Monitoring > Task Status**).
- For major updates, you can begin monitoring the update's progress in the task queue. However, after the Defense Center completes its necessary pre-update checks, you are logged out. When you log back in, the Upgrade Status page appears. See [Monitoring the Status of Major Updates](#) on page 2066 for information.

WARNING! Regardless of the update type, do **not** use the web interface to perform tasks other than monitoring the update until the update has completed and, if necessary, the Defense Center reboots. For more information, see [Using the Web Interface During the Update](#) on page 2060.

8. After the update finishes, if necessary, log into the Defense Center.
If you are the first user to log in after a major update, the End User License Agreement (EULA) may appear. You must review and accept the EULA to continue.
9. Clear your browser cache and force a reload of the browser. Otherwise, the user interface may exhibit unexpected behavior.

10. Select **Help > About** and confirm that the software version is listed correctly. Also note the versions of the rule update and VDB on the Defense Center; you will need this information later.
11. Verify that all managed devices are successfully communicating with the Defense Center.
12. If the rule update available on the [Sourcefire Support Site](#) is newer than the rules on your Defense Center, import the newer rules.
For more information, see [Importing Rule Updates and Local Rule Files](#) on page 2072.
13. 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 [Applying an Access Control Policy](#) on page 466.
14. If the VDB available on the [Sourcefire Support Site](#) is newer than the VDB on your Defense Center, 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](#) on page 2070.
15. Continue with the next section, [Updating Managed Devices](#), to update Sourcefire software on the devices that the Defense Center manages.

Updating Managed Devices

LICENSE: Any

After you update your Defense Centers, Sourcefire recommends that you use them to update the devices they manage. You **must** use the Defense Center to update managed devices that do not have a web interface, such as virtual devices.

Updating managed devices is a two-step process. First, download the update from the [Sourcefire Support Site](#) and upload it to the managing Defense Center. Next, install the software.

IMPORTANT! Traffic inspection, traffic flow, and link state may be affected during the update, depending on how your devices are configured and deployed, the components that the update affects, and whether the update reboots the devices. For specific information on how and when network traffic is affected for a particular update, see the release notes for that update.

To update managed devices:

ACCESS: Admin

1. Read the release notes and complete any required pre-update tasks.
Pre-update tasks may include updating your managing Defense Center, backing up event and configuration data, and making sure that the devices are running the correct version of the Sourcefire software, that computers where you installed software-based devices are running the correct version of their operating systems, that you have enough free disk space to perform the update, that you have set aside adequate time to perform the update, and so on.
2. Update the Sourcefire software on the devices' managing Defense Center; see [Updating a Defense Center](#) on page 2062.
3. Download the update from the [Sourcefire Support Site](#).
Different device models may use different updates. For information on the updates you can download, see the release notes.

IMPORTANT! Download the update directly from the [Sourcefire Support Site](#). If you transfer an update file by email, it may become corrupted.

4. Make sure that the appliances in your deployment are successfully communicating and that there are no issues being reported by the health monitor.
5. On the managing Defense Center, select **System > Updates**.
The Product Updates page appears.
6. Click **Upload Update** to browse to the update you downloaded, then click **Upload**.
The update is uploaded to the Defense Center. The Product Updates tab shows the type of update you just uploaded, its version number, and the date and time when it was generated. The page also indicates whether a reboot is required as part of the update.
7. Click the install icon next to the update you are installing.
The Install Update page appears.

8. Select the devices where you want to install the update, then click **Install**; you can update multiple devices at once if they use the same update. If prompted, confirm that you want to install the update and reboot the devices. The update process begins. Depending on the size of the file, it may take some time to install the update on all devices. You can monitor the update's progress in the Defense Center's task queue (**System > Monitoring > Task Status**). Note that managed devices may reboot twice during the update; this is normal.

WARNING! If you encounter issues with the update (for example, if the task queue indicates that the update has failed or if a manual refresh of the task queue shows no progress), do **not** restart the update. Instead, contact Support.

9. Optionally, after a major update, log in to the device's local web interface. If you are the first user to log in after a major update, the End User License Agreement (EULA) may appear. You must review and accept the EULA to continue. Note that the EULA also appears, and must be accepted, if your first login is via the command line interface rather than the web interface.
10. On the Defense Center, select **Devices > Device Management** and confirm that the devices you updated have the correct version listed.
11. Verify that the devices you updated are successfully communicating with the Defense Center.
12. 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 [Applying an Access Control Policy](#) on page 466.

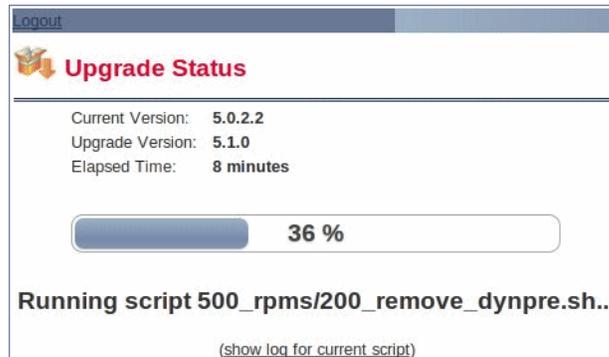
Monitoring the Status of Major Updates

LICENSE: Any

For major updates, the Sourcefire 3D System provides you with a streamlined web interface so that you can easily monitor the update process. The streamlined interface also prevents you from using the web interface to perform tasks other than monitoring the update.

You can begin monitoring the update's progress in the task queue (**System > Monitoring > Task Queue**). However, after the appliance completes its necessary pre-update checks, you and all other users are logged out of the web interface. Unless you are an administrator or a maintenance user, you cannot log back in until the update is complete.

For administrators, when you log back in, the streamlined update page appears. The following graphic shows the Defense Center version of the page.



When using a Defense Center to update a managed device, Sourcefire recommends that you monitor the update's progress from the Defense Center's task queue. Note, however, that if you attempt to log into the device's local web interface after the appliance finishes its pre-update checks, the streamlined update page appears and you can use it to monitor the update's progress.

The page displays the version of the Sourcefire 3D System 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.

WARNING! 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 appliance displays a success message and reboots. After the appliance finishes rebooting, refresh the page to log in and complete any required post-update steps.

Uninstalling Software Updates

LICENSE: Any

When you apply a patch or feature update to a Sourcefire appliance, the update process creates an uninstaller that allows you to remove the update from that appliance, using its web interface.

When you uninstall an update, the resulting Sourcefire software version depends on the update path for your appliance. For example, consider a scenario where you updated an appliance directly from Version 5.0 to Version 5.0.0.2. Uninstalling the Version 5.0.0.2 patch might result in an appliance running Version 5.0.0.1, even though you never installed the Version 5.0.0.1 update. For information on the resulting Sourcefire software version when you uninstall an update, see the release notes.

IMPORTANT! Uninstalling from the web interface is not supported for major updates. If you updated your appliance to a new major version of the Sourcefire 3D System and you need to revert to an older version, contact Support.

Order of Uninstallation

Uninstall the update in the reverse order that you installed it. That is, first uninstall the update from managed devices, then from Defense Centers.

Use the Local Web Interface to Perform the Update

You must use the local web interface to uninstall updates; you cannot use the Defense Center to uninstall updates from managed devices. For information on uninstalling a patch from a device that does not have a local web interface, see the release notes.

Uninstalling the Update from Clustered or Paired Appliances

Clustered devices and Defense Centers in high availability pairs must run the same version of the Sourcefire 3D System. Although the uninstallation process triggers an automatic failover, appliances in mismatched pairs or clusters do not share configuration information, nor do they install or uninstall updates as part of their synchronization. If you need to uninstall an update from redundant appliances, plan to perform the uninstallations in immediate succession.

You cannot uninstall an update from devices in a clustered stack if uninstalling would revert these devices to a version in which clustered stacking is not supported.

To ensure continuity of operations, uninstall the update from clustered devices and paired Defense Centers one at a time. First, uninstall the update from the

secondary appliance. Wait until the uninstallation process completes, then immediately uninstall the update from the primary appliance.

WARNING! If the uninstallation process on a clustered device or paired Defense Center fails, do **not** restart the uninstall or change configurations on its peer. Instead, contact Support.

Uninstalling the Update from Stacked Devices

All devices in a stack must run the same version of the Sourcefire 3D System. Uninstalling the update from any of the stacked devices causes the devices in that stack to enter a limited, mixed-version state.

To minimize impact on your deployment, Sourcefire recommends that you uninstall an update from stacked devices simultaneously. The stack resumes normal operation when the update completes on all devices in the stack.

You cannot uninstall an update from devices in a clustered stack if uninstalling would revert these devices to a version in which clustered stacking is not supported.

Traffic Flow and Inspection

Uninstalling an update from managed devices may affect traffic inspection, traffic flow, and link state. 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, there are several steps you should take to ensure that your deployment is performing properly. These include verifying that the uninstall succeeded and that all appliances in your deployment are communicating successfully. For specific information for each update, see the release notes.

To uninstall a patch or feature update using the local web interface:

ACCESS: Admin

1. Select **System > Updates**.
The Product Updates page appears.

2. Click the install icon next to the uninstaller for the update you want to remove.
 - On the Defense Center, the Install Update page appears. Select the Defense Center and click **Install**.
 - On a managed device, there is no intervening page.

In either case, if prompted, confirm that you want to uninstall the update and reboot the appliance.

The uninstall process begins. You can monitor its progress in the task queue (**System > Monitoring > Task Status**).

WARNING! Do **not** use the web interface to perform tasks other than monitoring the update until the uninstall has completed and, if necessary, the appliance reboots. For more information, see [Using the Web Interface During the Update](#) on page 2060.

3. After the uninstall finishes, if necessary, log into the appliance.
4. Clear your browser cache and force a reload of the browser. Otherwise, the user interface may exhibit unexpected behavior.
5. Select **Help > About** and confirm that the software version is listed correctly.
6. Verify that the appliance where you uninstalled the patch is successfully communicating with its managed devices (for the Defense Center) or its managing Defense Center (for managed devices).

Updating the Vulnerability Database

LICENSE: Any

The Sourcefire Vulnerability Database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The Sourcefire 3D System correlates the fingerprints with the vulnerabilities to help you determine whether a particular host increases your risk of network compromise. The Sourcefire Vulnerability Research Team (VRT) issues periodic updates to the VDB.

To update the VDB, use the Product Updates page on the Defense Center.

Type	Version	Date	Release Notes	Reboot
Sourcefire Vulnerability And Fingerprint Database Updates	72	Mon Nov 14 21:45:31 UTC 2011		No
Sourcefire 3D DC Sample Patch	5.1	Thu Nov 17 13:46:30 UTC 2011		Yes
Sourcefire 3D Device Sample Patch	5.1	Thu Nov 17 13:47:02 UTC 2011		Yes
Sourcefire 3D DC Sample Patch	5.0.1	Thu Nov 17 13:46:19 UTC 2011		Yes
Sourcefire 3D Device Sample Patch	5.0.1	Thu Nov 17 13:46:51 UTC 2011		Yes

When you upload VDB updates obtained from Sourcefire Support to your appliance, they appear on the page along with updates and uninstaller updates for the Sourcefire 3D System.

The time it takes to update vulnerability mappings depends on the number of hosts in your network map. You may want to schedule the update during low system usage times to minimize the impact of any system downtime. As a rule of thumb, divide the number of hosts on your network by 1000 to determine the approximate number of minutes to perform the update.

IMPORTANT! When you install a VDB update with changes to application detectors or operating system fingerprints, Sourcefire recommends that you check whether any of your managed devices are out-of-date and need to be reapplied. Installing a VDB update with detection updates may cause a short pause in traffic flow and processing on your managed devices, and may also cause a few packets to pass uninspected.

This section explains how to plan for and perform manual VDB updates. You can take advantage of the automated update feature to schedule VDB updates; see [Automating Vulnerability Database Updates](#) on page 1951.

To update the vulnerability database:

ACCESS: Admin

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, as well as product compatibility information.
2. Select **System > Updates**.
The Product Updates page appears.

3. Upload the update to the Defense Center:
 - If your Defense Center has access to the Internet, click **Download Updates** to check for the latest updates on the [Sourcefire Support Site](#).
 - If your Defense Center does not have access to the Internet, manually download the update from the [Sourcefire Support Site](#), then click **Upload Update**. Browse to the update and click **Upload**.

IMPORTANT! Download the update directly from the [Sourcefire 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 to the Defense Center.

4. Click the install icon next to the VDB update.
The Install Update page appears.
5. Select the Defense Center then click **Install**.
The update process begins. Depending on the number of hosts in your network map, installing the update may take some time. You can monitor the update's progress in the task queue (**System > Monitoring > Task Status**).

WARNING! Do **not** use the web interface to perform tasks related to mapped vulnerabilities until the update has completed. If you encounter issues with the update (for example, if the task queue indicates that the update has failed or if a manual refresh of the task queue shows no progress) **do not** restart the update. Instead, contact Support.

6. After the update finishes, select **Help > About** to confirm that the VDB build number matches the update you installed.

Importing Rule Updates and Local Rule Files

LICENSE: Any

As new vulnerabilities become known, the Sourcefire Vulnerability Research Team (VRT) releases rule updates. 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 and provide new rule categories and system variables.

IMPORTANT! Rule updates may contain new binaries. Make sure your process for downloading and installing rule updates complies with your security policies. In addition, rule updates may be quite large, so make sure to import rules during periods of low network use.

The following are additional important points you should keep in mind when you import rules:

- For new rules in rule updates, the rule state may be different in each default policy. For example, a new rule may be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. See [Using Default Intrusion Policies](#) on page 692 for more information.

Rule updates may also change the default state of existing rules. For information on choosing whether to allow rule updates to change the default states of existing rules in intrusion policies you create, see [Allowing Rule Updates to Modify the Base Policy](#) on page 694.
- Rule updates are cumulative, so the newest rule update contains the intrusion rules of all previous updates. You cannot import a rule update that either matches or predates the version of the currently installed rules.
- When you use a default policy provided by Sourcefire as your base policy, you can choose whether to allow rule updates to modify your base policy with any changes to intrusion rules, preprocessor rules, and advanced settings. See [Understanding the Base Policy](#) on page 691 and [Allowing Rule Updates to Modify the Base Policy](#) on page 694 for more information.
- Rule updates may include new system variables and modified values for existing system variables. New variables are always added to your system. Your existing variable values are updated only if you have not modified them. See [Managing Variables](#) on page 698 for more information.
- Rule updates may include new rule categories. New rule categories in rule updates are always added to your system. See [Understanding Rule Categories](#) on page 739 for more information.
- The Rule Updates page lists intrusion policies with cached changes and the users who made those changes. Importing a rule update discards all cached changes. See [Committing Intrusion Policy Changes](#) on page 679 for more information.
- When a rule update includes shared object rules, applying an access control policy for this first time after the rule import causes a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected.
- If your Sourcefire 3D System deployment includes two Defense Centers configured as a high availability pair, you only need to update rules on one of the Defense Centers. The second Defense Center receives the rule update as part of the regular synchronization process.
- Optionally, when the import completes, you can automatically reapply intrusion policies owned by the appliance where you import the rule update.

See the following sections for more information:

- [Using One-Time Rule Updates](#) on page 2074 explains how to import a single rule update from the [Sourcefire Support Site](#).
- [Using Recurring Rule Updates](#) on page 2077 explains how to use an automated feature on the web interface to download and install rule updates from the [Sourcefire Support Site](#).
- [Importing Local Rule Files](#) on page 2080 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](#) on page 2082 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](#) on page 2074 explains how to manually download a rule update from the [Sourcefire Support Site](#) to your local machine and then manually install the rule update.
- [Using Automatic One-Time Rule Updates](#) on page 2076 explains how to use an automated feature on the web interface to search the [Sourcefire 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 Defense Center does not have Internet access.

To manually import a rule update:

ACCESS: Admin

1. From a computer that can access the Internet, access and log into the [Sourcefire Support Site](#).
2. Click **Download**, then click **Rules**.
3. Navigate to the latest rule update.

TIP! Rule updates are cumulative, so the newest rule update contains the intrusion rules and new features of all previous rule updates. You cannot import a rule update with a version number lower than the version of the currently installed update.

4. Click the rule update file that you want to download and save it to your computer.

5. Log into your appliance's web interface.
6. Select **System > Updates**, then select the **Rule Updates** tab.
The Rule Updates page appears.

One-Time Rule Update/Rules Import

Source	<input type="radio"/> Rule Update or text rule file to upload and install <input type="text"/> <input type="button" value="Browse..."/>
	<input type="radio"/> Download new Rule Update from the support site
Policy Reapply	<input type="checkbox"/> Reapply intrusion policies after the Rule Update import completes
<input type="button" value="Import"/>	

Recurring Rule Update Imports
The scheduled Rule Update feature is not enabled

TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policies > Intrusion > Rule Editor**.

7. 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](#) on page 1188 for more information.
8. Select **Rule Update or text rule file to upload and install** and click **Browse** to navigate to and select the rule update file.
9. Optionally, select **Reapply intrusion policies after the Rule Update import completes** to automatically reapply intrusion policies currently applied from this appliance when the rule update import completes.

Note that you cannot apply access control policies to stacked devices that are running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). See [Managing Stacked Devices](#) on page 237 for more information.

10. Click **Import**.

The rule update is installed and the system displays the Rule Update Log detailed view. See [Understanding the Rule Update Import Log Detailed View](#) on page 2088 for more information.

If you selected **Reapply intrusion policies after the Rule Update import completes** in step 9, the system applies only the intrusion policies in the currently applied access control policy but does not apply the access control policy. See [Applying an Access Control Policy](#) on page 466 for more information.

If you did not select **Reapply intrusion policies after the Rule Update import completes**, changes in the rule update are not implemented until the next time you apply the affected intrusion policies. See [Reapplying an Intrusion Policy](#) on page 680 for more information.

IMPORTANT! 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 [Sourcefire Support Site](#). You can use this procedure only if the appliance has Internet access.

To automatically import a rule update:

ACCESS: Admin

1. Select **System > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

The screenshot shows two sections of the web interface. The top section is titled "One-Time Rule Update/Rules Import" and contains the following elements: a "Source" label, a radio button selected for "Rule Update or text rule file to upload and install" with a text input field and a "Browse..." button, a radio button for "Download new Rule Update from the support site", a "Policy Reapply" label, and a checkbox for "Reapply intrusion policies after the Rule Update import completes" which is currently unchecked. Below these options is an "Import" button. The bottom section is titled "Recurring Rule Update Imports" and contains the text "The scheduled Rule Update feature is not enabled" and a checkbox for "Enable Recurring Rule Update Imports" which is also unchecked. Below this checkbox are "Save" and "Cancel" buttons.

TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policies > Intrusion > Rule Editor**.

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](#) on page 1188 for more information.
3. Select **Download new Rule Update from the support site**.
4. Optionally, select **Reapply intrusion policies after the Rule Update import completes** to automatically reapply intrusion policies currently applied from this appliance when the rule update completes.

Note that you cannot apply intrusion policies to stacked devices that are running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). See [Managing Stacked Devices](#) on page 237 and [Reapplying an Intrusion Policy](#) on page 680 for more information.

5. Click **Import**.

The rule update is installed and the system displays the Rule Update Log detailed view workflow. See [Rule Update Import Log Detailed View Fields](#) on page 2088 for more information.

If you selected **Reapply intrusion policies after the Rule Update import completes** in step 4, the system applies only the intrusion policies in the currently applied access control policy but does not apply the access control policy. See [Applying an Access Control Policy](#) on page 466 for more information.

If you did not select **Reapply intrusion policies after the Rule Update import completes**, changes in the rule update are not implemented until the next time you apply the affected intrusion policies. See [Reapplying an Intrusion Policy](#) on page 680 for more information.

IMPORTANT! 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. If your Sourcefire 3D System deployment includes two Defense Centers configured as a high availability pair, you only need to update rules on one of the Defense Centers. The second Defense Center receives the rule update as part of the regular synchronization process.

To schedule recurring rule updates:

ACCESS: Admin

1. Select **System > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

One-Time Rule Update/Rules Import

Source

Rule Update or text rule file to upload and install

Download new Rule Update from the support site

Policy Reapply

Reapply intrusion policies after the Rule Update import completes

Import

Recurring Rule Update Imports

The scheduled Rule Update feature is not enabled

Enable Recurring Rule Update Imports

Save Cancel

TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policies > Intrusion > Rule Editor**.

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](#) on page 1188 for more information.
3. Select **Enable Recurring Rule Update Imports**.

The page expands to display options for configuring recurring imports.

Recurring Rule Update Imports

The scheduled SRU import feature is not enabled

Enable Recurring Rule Update Imports

Import Frequency

Daily at 12 :00 AM

America/New York

Policy Reapply

Reapply intrusion policies after the Rule Update import completes

Save Cancel

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**.

4. In the **Import Frequency** field, select **Daily**, **Weekly**, or **Monthly** from the drop-down list.

TIP! You can select from a recurring task drop-down list either by clicking on your selection or by typing the first letter or number in the selection one or more times and pressing Enter.

5. If you selected **Weekly** in the **Import Frequency** field, use the drop-down list that appears to select the day of the week when you want to import rule updates.
6. If you selected **Monthly** in the **Import Frequency** field, use the drop-down list that appears to select the day of the month when you want to import rule updates.
7. In the **Import Frequency** field, specify the time when you want to start your recurring rule update import.
8. Optionally, select **Reapply intrusion policies after the Rule Update import completes** to automatically reapply intrusion policies currently applied from this appliance when the rule update import completes.

Note that you cannot apply intrusion policies to stacked devices that are running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the devices fails). See [Managing Stacked Devices](#) on page 237 for more information.

9. 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.
The rule update is installed at the scheduled time and the rules are updated. You can log off or use the web interface to perform other tasks before or during the import. When accessed during an import, the Rule Update Log displays a red status icon (). See [Viewing the Rule Update Log](#) on page 2082 for more information. During an import, you can also view

messages as they occur in the Rule Update Log detailed view. See [Rule Update Import Log Detailed View Fields](#) on page 2088 for more information.

IMPORTANT! Depending on rule update size and content, several minutes may pass before status messages appear in the Rule Update Log or Rule Update Log detailed view.

If you selected **Reapply intrusion policies after the Rule Update import completes** in step 8, the system applies only the intrusion policies in the currently applied access control policy but does not apply the access control policy. See [Applying an Access Control Policy](#) on page 466 for more information.

If you did not select **Reapply intrusion policies after the Rule Update import completes**, changes in the rule update are not implemented until the next time you apply the affected intrusion policies. See [Reapplying an Intrusion Policy](#) on page 680 for more information.

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 appliance where the recurring import is configured.

IMPORTANT! 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 in an ASCII text file from a local machine. 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 local rules:

- 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 > 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 > 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](#) on page 743 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.
- Sourcefire strongly recommends that you import local rules on the primary Defense Center in a High Availability Pair to avoid SID numbering issues.
- 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](#) on page 747 for more information.

To import local rule files:

ACCESS: Admin

1. Select **Policies > Intrusion > Rule Editor**.

The Rule Editor page appears.

2. Click **Import Rules**.

The Import Rules page appears.

One-Time Rule Update/Rules Import

Source

Rule Update or text rule file to upload and install

Download new Rule Update from the support site

Policy Reapply

Reapply intrusion policies after the Rule Update import completes

Import

Recurring Rule Update Imports

The scheduled Rule Update feature is not enabled

TIP! You can also select **System > Updates**, then select the **Rule Updates** tab.

3. Select **Rule Update or text rule file to upload and install** and click **Browse** to navigate to the rule file. Note that all rules uploaded in this manner are saved in the local rule category.

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.

IMPORTANT! Managed devices do **not** use the new rule set for inspection until after you apply their intrusion policies. See [Applying an Access Control Policy](#) on page 466 for procedures.

Viewing the Rule Update Log

LICENSE: Any

The Defense Center 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. The fields in the Rule Update Log are described in the [Rule](#)

[Update Log Actions](#) table.

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 on page 2084.
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. IMPORTANT! 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](#) on page 2084 describes the fields in the list of rule updates and local rule files that you import.
- [Viewing Rule Update Import Log Details](#) on page 2085 describes the detailed record for each object imported in a rule update or local rule file.
- [Understanding the Rule Update Import Log Detailed View](#) on page 2088 describes each field in the Rule Update Log detailed view.
- [Searching the Rule Update Import Log](#) on page 2089 explains how you can search the import log for specific records or for all records matching the search criteria.

To view the Rule Update Log:

ACCESS: Admin

1. Select **System > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

One-Time Rule Update/Rules Import

Source Rule Update or text rule file to upload and install

Download new Rule Update from the support site

Policy Reapply Reapply intrusion policies after the Rule Update import completes

Recurring Rule Update Imports
The scheduled Rule Update feature is not enabled

TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policies > Intrusion > Rule Editor**.

2. Click **Rule Update Log**.

The Rule Update Log page appears. This page lists each imported rule update and local rule file.

Product Updates		Rule Updates			
Summary	Time	User ID	Status		
Sourcefire Rule Update 2011 11 06 001 dev Completed install of Sourcefire Rule Update 2011-11-06-001-dev	2011-11-10 10:59:43	admin	✓		
Sourcefire Rule Update 2011 11 03 002 dev Completed install of Sourcefire Rule Update 2011-11-03-002-dev	2011-11-09 16:17:48	admin	✓		

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 [Rule Update Log Fields](#) table.

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.

Rule Update Log Fields (Continued)

FIELD	DESCRIPTION
User ID	The user name of the user that triggered the import.
Status	Whether the import: <ul style="list-style-type: none"> • succeeded () • failed or is currently in progress () <p>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.</p>

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 [Rule Update Import Log Detailed View Actions](#) table below describes specific actions you can perform on a Rule Update Import Log detailed view workflow page.

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 on page 2088.
sort and constrain records on the current workflow page	find more information in Sorting Drill-Down Workflow Pages on page 1835.

Rule Update Import Log Detailed View Actions (Continued)

To...	You CAN...
temporarily use a different workflow	click (switch workflows) . For information on selecting workflows, see Selecting Workflows on page 1809. For information on creating custom workflows, see Creating Custom Workflows on page 1841.
bookmark the current page so that you can quickly return to it	click Bookmark This Page . For more information, see Using Bookmarks on page 1838.
navigate to the bookmark management page	click View Bookmarks . For more information, see Using Bookmarks on page 1838.
generate a report based on the data in the current view	click Report Designer . For more information, see Creating a Report Template from an Event View on page 1725.
search the entire Rule Update Import Log database for rule update import records	click Search . For more information, see Searching the Rule Update Import Log on page 2089.
open a search page prepopulated with the current single constraint	select Edit Search or Save Search next to Search Constraints. For more information, see the Table View and Drill-Down Page Features table on page 1814.

To view the Rule Update Import Log Detailed View:

ACCESS: Admin

1. Select **System > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

One-Time Rule Update/Rules Import

Source	<input type="radio"/> Rule Update or text rule file to upload and install <input type="text"/> <input type="button" value="Browse..."/>
	<input type="radio"/> Download new Rule Update from the support site
Policy Reapply	<input type="checkbox"/> Reapply intrusion policies after the Rule Update import completes
<input type="button" value="Import"/>	

Recurring Rule Update Imports

The scheduled Rule Update feature is not enabled

TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policies > Intrusion > Rule Editor**.

2. Click **Rule Update Log**.

The Rule Update Log page appears.

Product Updates	Rule Updates			
Summary	Time	User ID	Status	
Sourcefire Rule Update 2011 11 06 001 dev Completed install of Sourcefire Rule Update 2011-11-06-001-dev	2011-11-10 10:59:43	admin	✓	 
Sourcefire Rule Update 2011 11 03 002 dev Completed install of Sourcefire Rule Update 2011-11-03-002-dev	2011-11-09 16:17:48	admin	✓	 

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 [Rule Update Import Log Detailed View Fields](#) table.

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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • new (for a rule, this is the first time the rule has been stored on this appliance) • 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 on the appliance) • 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 default policy provided by Sourcefire) • disabled (for rules, the rule has been disabled in a default policy provided by Sourcefire) • drop (for rules, the rule has been set to Drop and Generate Events in a default policy provided by Sourcefire) • 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). See the Generator IDs table on page 785 for more information.
SID	The SID for a rule.

Rule Update Import Log Detailed View Fields (Continued)

FIELD	DESCRIPTION
Rev	The revision number for a rule.
Policy	For imported rules, this field displays All , which indicates that the imported rule was included in all default 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.

Searching the Rule Update Import Log

LICENSE: Any

You can search the import log for specific records or for all records matching the search criteria. You may want to create customized searches and save them to reuse later.

TIP! You search the entire Rule Update Import Log database even when you initiate a search by clicking **Search** on the toolbar from the Rule Update Import Log detailed view with only the records for a single import file displayed. Make sure you set your time constraints to include all objects you want to include in the search. See [Specifying Time Constraints in Searches](#) on page 1775 for more information.

The search criteria you can use are described in the [Rule Update Import Log Search Criteria](#) table. Note that record searches are case-insensitive. For example, searching for `RULE` or `rule` yields the same results.

Rule Update Import Log Search Criteria

SEARCH FIELD	DESCRIPTION	EXAMPLE
Time	Specify the date and time the record was generated. See Specifying Time Constraints in Searches on page 1775 for the syntax for entering time.	<code>> 2006-01-15 13:30:00</code> returns all rule records imported after January 15, 2006 at 1:30 PM.
Name	Specify all or part of the content of the rule Message field. You can use an asterisk (*) as a wildcard character in this field.	<code>*dhcp*</code> returns all rule records with DHCP in the Message field.
Type	Specify the type of record, which can be <code>SRU component</code> , <code>rule</code> , or <code>policy apply</code> . Note that you can use the <code>update</code> search value to search for rules imported prior to Version 5.0.1.	<code>SRU component</code> returns imported components such as a rule pack or policy pack; <code>rule</code> returns rule updates, including new rules; <code>policy apply</code> returns a table row of information for rule updates where intrusion policies were automatically reapplied following the update.
Action	Specify an action for the object you want to view. See the Rule Update Import Log Detailed View Fields table on page 2088 for a list of actions you can specify.	When the type is <code>rule</code> , <code>new</code> returns all rules imported for the first time on the appliance.
GID	Specify the generator ID for the rule.	<code>3</code> returns all shared object rules.
SID	Specify a signature ID or a range of SIDs for a rule.	<code>923</code> returns the record for the rule with the SID 923.
Rev	Specify the revision number for the rule.	<code>3</code> returns rules with the revision number 3.
Policy	Specify the default policy the rule is imported into.	<code>A11</code> returns rules imported into all default policies.
Rule Update	Specify the Rule Update filename.	<code>filename</code> returns all records for the specified import file.
Details	Specify details on the imported object.	<code>previously*</code> returns the record for all rules that have changed.

For more information on searching, including how to load and delete saved searches, see [Searching for Events](#) on page 1770.

To search the Rule Update Import Log:

ACCESS: Admin/Intrusion Admin

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Rule Update Import Log**.
The page reloads with the appropriate constraints.

The screenshot shows a web interface for searching the Rule Update Import Log. At the top, it says "Search Information" and includes a note: "Note: If a search name is not specified, an automatically generated name will be used." Below this, there are several fields for defining the search criteria:

- Table:** A dropdown menu set to "Rule Update Import Log".
- Name:** An empty text input field. To its right, the text "Search 1, My Search" is displayed.
- Save As Private:** A checkbox that is checked.
- Constraint:** A section with multiple rows of search criteria, each with a text input field and a corresponding value:
 - Time:** Input field is empty; value is "> 2009-07-16 13:00:31, < today at 4:30pm".
 - Name:** Input field is empty; value is "P2P WinNY connection attempt".
 - Type:** Input field is empty; value is "SRU component, update".
 - Action:** Input field is empty; value is "new, changed, deleted, collision, error, apply, disabled, enabled, drop".
 - GID:** Input field is empty; value is "1".
 - SID:** Input field is empty; value is "923".
 - Rev:** Input field is empty; value is "2".
 - Policy:** Input field is empty; value is "All".
 - Rule Update:** Input field is empty; value is "Your Rule Update".
 - Details:** Input field is empty; value is "previously (1:494:10)".

At the bottom of the form, there are two buttons: "Search" and "Save As New Search".

TIP! You can also click **Search** on the Rule Update Log detailed view; see [Viewing Rule Update Import Log Details](#) on page 2085.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, the web interface automatically creates one when you save it.
4. Enter your search criteria in the appropriate fields, as described in the [Rule Update Import Log Search Criteria table](#) on page 2090. If you enter multiple criteria, the search returns the records that match all the criteria.

5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default Rule Update Import Log detailed view workflow. To use a different workflow, including a custom workflow, click (**switch workflows**). For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**) so that you can run it at a later time.

Updating the Geolocation Database

LICENSE: FireSIGHT

SUPPORTED DEFENSE CENTERS: Any except DC500

The Sourcefire Geolocation Database (GeoDB) is a database of geographical data (such as country, city, coordinates, and so on) and connection-related data (such as internet service provider, domain name, connection type, and so on) associated with routable IP addresses. When your system detects GeoDB information that matches a detected IP address, you can view the geolocation information associated with that IP address. You must install the GeoDB on your system to view any geolocation details other than country or continent. Sourcefire issues periodic updates to the GeoDB.

To update the GeoDB, use the Geolocation Updates page (**System > Updates > Geolocation Updates**) on the Defense Center. When you upload GeoDB updates

you obtained from Sourcefire Support or from your appliance, they appear on this page.

The screenshot shows a web interface with three tabs: 'Product Updates', 'Rule Updates', and 'Geolocation Updates'. The 'Geolocation Updates' tab is active. Below the tabs, it says 'Defense Center running geolocation update version: 2013-03-12-001'. There are two main sections: 'One-Time Geolocation Update' and 'Recurring Geolocation Updates'. The 'One-Time' section has a note: 'Note that updates may be large and can take up to 45 minutes.' It has two radio buttons: 'Upload and install geolocation update' (selected) and 'Download and install geolocation update from the Support Site'. The first option has a text input field and a 'Browse...' button. Below it is an 'Import' button. The 'Recurring Geolocation Updates' section has a checkbox 'Enable Recurring Weekly Updates' which is checked. Below it is 'Update Start Time' with dropdown menus for 'Monday', '01:00', 'Pm', and 'America/New York'. At the bottom are 'Save' and 'Cancel' buttons.

Time needed to update the GeoDB depends on your appliance; 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](#) on page 1942.

To update the geolocation database:

ACCESS: Admin

1. Select **System > Updates**.
The Product Updates page appears.
2. Click the **Geolocation Updates** tab.
The Geolocation Updates page appears.

3. Upload the update to the Defense Center.
 - If your Defense Center has access to the Internet, click **Download and install geolocation update from the Support Site** to check for the latest updates on the [Sourcefire Support Site](#).
 - If your Defense Center does not have access to the Internet, manually download the update from the [Sourcefire Support Site](#), then click **Upload and install geolocation update**. Browse to the update and click **Import**.

IMPORTANT! Download the update directly from the [Sourcefire 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; this may vary depending on your appliance hardware. You can monitor the update's progress in the task queue (**System > Monitoring > Task Status**).

4. After the update finishes, return to the Geolocation Updates page or select **Help > About** 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.

CHAPTER 52

MONITORING THE SYSTEM

The Sourcefire 3D System 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 and intrusion event information, as well as statistics for the Data Correlator and network discovery processes for the current day. You can also monitor both summary and detailed information on all processes that are currently running on the Defense Center or managed device. The following sections provide more information about the monitoring features that the system provides:

- [Viewing Host Statistics](#) on page 2096 describes how to view host information such as:
 - system uptime
 - disk and memory usage
 - Data Correlator statistics
 - system processes
 - intrusion event information
- On the Defense Center, you can also use the health monitor to monitor disk usage and alert on low disk space conditions. For more information, see [Understanding Health Monitoring](#) on page 2110.
- [Monitoring System Status and Disk Space Usage](#) on page 2099 describes how to view basic event and disk partition information.
- [Viewing System Process Status](#) on page 2100 describes how to view basic process status.
- [Understanding Running Processes](#) on page 2103 describes the basic system processes that run on the appliance.

You can use the options in **Overview > Summary** to view and graph statistics for intrusion and discovery events. For more information, see:

- [Viewing Intrusion Event Statistics](#) on page 595
- [Viewing Intrusion Event Graphs](#) on page 601
- [Viewing Discovery Event Statistics](#) on page 1378
- [Viewing Discovery Performance Graphs](#) on page 1384

Viewing Host Statistics

LICENSE: Any

The Statistics page lists the current status of the following:

- general host statistics; see the [Host Statistics table](#) on page 2096 for details
- Data Correlator statistics (Defense Center only — requires FireSIGHT); see the [Data Correlator Process Statistics table](#) on page 2097 for details
- intrusion event information (requires Protection); see the [Intrusion Event Information table](#) on page 2098 for details

The [Host Statistics](#) table describes the host statistics listed on the Statistics page.

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.
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 on page 2099 for more information.
Processes	A summary of the processes running on the system. See Viewing System Process Status on page 2100 for more information.

If your Sourcefire 3D System deployment includes a Defense Center with a FireSIGHT license, you can also view statistics about the Data Correlator and network discovery processes for the current day. As the managed devices perform data acquisition, decoding, and analysis, the network discovery process correlates the data with the fingerprint and vulnerability databases, then produces binary files that are processed by the Data Correlator running on the Defense Center. The Data Correlator analyzes the information from the binary files, generates events, and creates the discovery network map.

The statistics that appear for network discovery and the Data Correlator are averages for the current day, using statistics gathered between 12:00 AM and 11:59 PM for each device.

The [Data Correlator Process Statistics](#) table describes the statistics displayed for the Data Correlator process.

Data Correlator Process Statistics

CATEGORY	DESCRIPTION
Events/Sec	Number of discovery events that the Data Correlator receives and processes per second
Connections/Sec	Number of connections that the Data Correlator receives and processes per second
CPU Usage — User (%)	Average percentage of CPU time spent on user processes for the current day
CPU Usage — System (%)	Average percentage of CPU time spent on system processes for the current day
VmSize (KB)	Average size of memory allocated to the Data Correlator for the current day, in kilobytes
VmRSS (KB)	Average amount of memory used by the Data Correlator for the current day, in kilobytes

On managed devices and on Defense Centers that manage devices, you can also view the date and time of the last intrusion event, the total number of events that

have occurred in the past hour and the past day, and the total number of events in the database.

IMPORTANT! The information in the Intrusion Event Information section of the Statistics page is based on intrusion events stored on the managed device rather than those sent to the Defense Center. If you manage your device so that intrusion events are not stored locally, no intrusion event information is listed on this page. This is also the case for managed devices that cannot store events locally.

The [Intrusion Event Information](#) table describes the statistics displayed in the Intrusion Event Information section of the Statistics page.

Intrusion Event Information

STATISTIC	DESCRIPTION
Last Alert Was	The date and time that the last event occurred
Total Events Last Hour	The total number of events that occurred in the past hour
Total Events Last Day	The total number of events that occurred in the past twenty-four hours
Total Events in Database	The total number of events in the events database

To view the Statistics page:

ACCESS: Admin/Maint

1. Select **System > Monitoring > Statistics**.

The Statistics page appears. The Defense Center version of the page is shown below.



Statistics for katsura

Time	2012-06-04 14:27:50
Uptime	4 days 4:32
Memory Usage	 25.54% (3.0 GB / 11.7 GB)
Load Average	0.31, 0.23, 0.17
Disk Usage	▶ Total  1.2% used (9.9 GB / 819.4 GB)
Processes	▶ Tasks: 172 total, 1 running, 171 sleeping, 0 stopped, 0 zombie

SFDataCorrelator Process Statistics

Events/Sec	0.00
Connections/Sec	408.55
CPU Usage - User	1.21%
CPU Usage - System	0.24%
VmSize (KB)	922724
VmRSS (KB)	619732

2. On the Defense Center, you can also list statistics for managed devices. From the **Select Device(s)** box, click **Select Devices**. You can use the Shift and Ctrl keys to select multiple devices at once.

The Statistics page is updated with statistics for the devices you selected.

Monitoring System Status and Disk Space Usage

LICENSE: Any

The Disk Usage section of the Statistics page provides a quick synopsis of 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.

TIP! On the Defense Center, you can also use the health monitor to monitor disk usage and alert on low disk space conditions. For more information, see [Understanding Health Monitoring](#) on page 2110.

To access disk usage information:

ACCESS: Admin/Maint

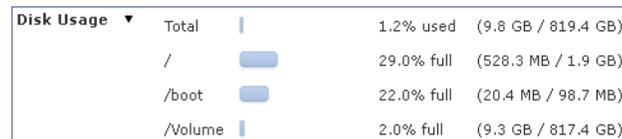
1. Select **System > Monitoring > Statistics**.

The Statistics page appears.

2. Click the down arrow next to **Disk Usage** to expand it.



The Disk Usage section expands.



If your deployment includes multiple managed devices, you may want to constrain disk usage data by specific devices.

On the Defense Center, to view disk usage information for a specific device:

ACCESS: Admin/Maint

1. Select the device name from the **Select Device(s)** box, and click **Select Devices**.

The page reloads, listing host statistics for each device you selected.

2. Click the down arrow next to **Disk Usage** to expand it.

The Disk Usage section expands.

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. If you are managing devices with a Defense Center, you can use the Defense Center's web interface to view the process status for any managed device.

The [Process Status](#) table describes each column that appears in the process list.

Process Status

COLUMN	DESCRIPTION
Pid	The process ID number
Username	The name of the user or group running the process
Pri	The process priority

Process Status (Continued)

COLUMN	DESCRIPTION
Nice	The <i>nice</i> 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 <i>m</i> , which indicates megabytes)
Res	The amount of resident paging files in memory (in kilobytes unless the value is followed by <i>m</i> , which indicates megabytes)
State	The process state: <ul style="list-style-type: none"> • 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:

ACCESS: Admin/Maint

1. Select **System > Monitoring > Statistics**.
The Statistics page appears.
2. On the Defense Center, select the device or devices you want to view process statistics for from the **Select Device(s)** box and click **Select Devices**.

3. 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.

Processes ▼ Tasks: 167 total, 1 running, 166 sleeping, 0 stopped, 0 zombie 11:50:12 up 7 days, 2:34, 0 users, load average: 0.10, 0.10, 0.08 Cpu(s): 0.6%us, 0.1%sy, 0.0%ni, 99.3%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st Mem: 12272868k total, 7067576k used, 5205292k free, 93336k buffers Swap: 5238520k total, 0k used, 5238520k free, 5538428k cached									
Pid	Username	Pri	Nice	Size	Res	State	Time	Cpu	Command
1	root	19	0	949	608	Ss	00:00:14	0.0	init [3]
2	root	19	0	0	0	S	00:00:00	0.0	[kthreadd]
3	root	19	0	0	0	S	00:00:02	0.0	[ksoftirqd/0]
3	root	19	0	0	0	S	00:00:02	0.0	[ksoftirqd/0]
3	root	19	0	0	0	S	00:00:02	0.0	[ksoftirqd/0]
6	root	19	0	0	0	S	00:00:02	0.0	[ksoftirqd/1]

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

IMPORTANT! For more information about the types of processes that run on the appliance, see [Understanding Running Processes](#) on page 2103.

To collapse the process list:

ACCESS: Admin/Maint

- ▶ 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](#) on page 2103
- [Understanding Executables and System Utilities](#) on page 2105

Understanding System Daemons

LICENSE: Any

Daemons continually run on an appliance. They ensure that services are available and spawn processes when required. The [System Daemons](#) table lists daemons that you may see on the Process Status page and provides a brief description of their functionality.

IMPORTANT! The table below is not an exhaustive list of all processes that may run on an appliance.

System Daemons

DAEMON	DESCRIPTION
crond	Manages the execution of scheduled commands (cron jobs)
dhclient	Manages dynamic host IP addressing
fpcollect	Manages the collection of client and server fingerprints
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

System Daemons (Continued)

DAEMON	DESCRIPTION
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 Sourcefire 3D System database processes
ntpd	Manages the Network Time Protocol (NTP) process
pm	Manages all Sourcefire 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
SFDataCorrelator	Manages data transmission
sfstreamer (Defense Center only)	Manages connections to third-party client applications that use the Event Streamer
sfmgr	Provides the RPC service for remotely managing and configuring an appliance using an sftunnel connection to the appliance
SFRemediateD (Defense Center only — requires FireSIGHT)	Manages remediation responses
sftimeserviced (Defense Center only)	Forwards time synchronization messages to managed devices
sfmbservice (requires Protection)	Provides access to the sfmb message broker process running on a remote appliance, using an sftunnel connection to the appliance. Currently used only by health monitoring to send health events and alerts from a managed device to a Defense Center or, in a high availability environment, between Defense Centers

System Daemons (Continued)

DAEMON	DESCRIPTION
sftroughd	Listens for connections on incoming sockets and then invokes the correct executable (typically the Sourcefire 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 [System Executables and Utilities](#) table describes the executables that you may see on the Process Status page.

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
SFDataCorrelator (Defense Center only — requires FireSIGHT)	Analyzes binary files created by FireSIGHT to generate events, connection data, and the network map
cp	Utility that copies files
df	Utility that lists the amount of free space on the appliance

System Executables and Utilities (Continued)

EXECUTABLE	DESCRIPTION
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 on page 1971 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

System Executables and Utilities (Continued)

EXECUTABLE	DESCRIPTION
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
sfheartbeat	Identifies a heartbeat broadcast, indicating that the appliance is active; heartbeat used to maintain contact between a device and Defense Center
sfmb	Indicates a message broker process; handles communication between Defense Centers and device.
sh	Public domain version of the Korn shell
shutdown	Utility that shuts down the appliance
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

System Executables and Utilities (Continued)

EXECUTABLE	DESCRIPTION
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

CHAPTER 53

USING HEALTH MONITORING

The health monitor provides numerous tests for determining the health of an appliance from the Defense Center. You can use the health monitor to create a collection of tests, referred to as a *health policy*, and apply the health policy to one or more appliances. You can create one health policy for every appliance in your system, customize a health policy for the specific appliance where you plan to apply it, or use the default health policy. You can also import a health policy exported from another Defense Center.

The tests, referred to as *health modules*, are scripts that test for criteria you specify. You can modify a health policy by enabling or disabling tests or by changing test settings, and you can delete health policies that you no longer need. You can also suppress messages from selected appliances by blacklisting them.

The tests in a health policy run automatically at the interval you configure. You can also run all tests, or a specific test, on demand. The health monitor collects health events based on the test conditions configured. Optionally, you can also configure email, SNMP, or syslog alerting in response to health events.

On the Defense Center, you can view health status information for the entire system or for a particular appliance. Fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. These event views allow you to search and view event data and to access other information that may be related to the events you are investigating.

You can also generate troubleshooting files for an appliance if you are asked to do so by Support.

See the following sections for more information:

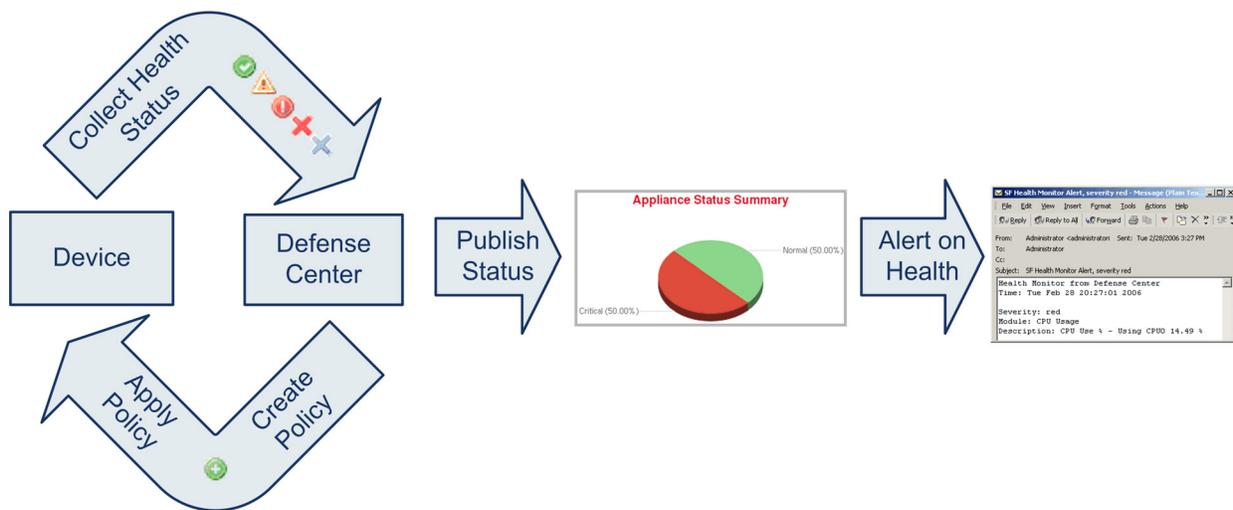
- [Understanding Health Monitoring](#) on page 2110
- [Configuring Health Policies](#) on page 2116

- [Using the Health Monitor Blacklist](#) on page 2155
- [Configuring Health Monitor Alerts](#) on page 2159
- [Using the Health Monitor](#) on page 2163
- [Using Appliance Health Monitors](#) on page 2166
- [Working with Health Events](#) on page 2174

Understanding Health Monitoring

LICENSE: Any

You can use the health monitor to check the status of critical functionality across your Sourcefire 3D System deployment. Monitor the health of your entire Sourcefire 3D System through the Defense Center by applying health policies to each of the managed devices and collecting the resulting health data at the Defense Center. Pie charts and status tables on the Health Monitor page visually represent the health status for monitored appliances, so you can check status at a glance, then drill down into status details if needed.



You can use the health monitor to access health status information for the entire system or for a particular appliance. The Health Monitor page provides a visual summary of the status of all appliances on your system. Individual appliance health monitors let you drill down into health details for a specific appliance.

You can also view health events in the standard Sourcefire 3D System table view. From an individual appliance's health monitor, you can open a table view of occurrences of a specific event, or you can retrieve all the health events for that appliance. You can also search for specific health events. For example, if you want to see all the occurrences of CPU usage with a certain percentage, you can search for the CPU usage module and enter the percentage value.

You can also configure email, SNMP, or syslog alerting in response to health events. A *health alert* is an association between a standard alert and a health status level. For example, if you need to make sure an appliance never fails due to hardware overload, you can set up an email alert. You can then create a health alert that triggers that email alert whenever CPU, disk, or memory usage reaches the Warning level you configure in the health policy applied to that appliance. You can set alerting thresholds to minimize the number of repeating alerts you receive.

Because health monitoring is an administrative activity, only users with administrator user role privileges can access system health data. For more information on assigning user privileges, see [Modifying User Privileges and Options](#) on page 1912.

IMPORTANT! Except for the Defense Center, Sourcefire 3D System devices do not have health monitoring policies applied to them by default. Managed devices report hardware status automatically via the Hardware Alarms health module; if you want to use other modules to monitor a managed device, you must apply a health policy to that device. For more information on the Sourcefire-provided default health policy for your appliances, see [Understanding the Default Health Policy](#) on page 2117. For more information on creating customized health policies, see [Creating Health Policies](#) on page 2118. For details on applying policies, see [Applying Health Policies](#) on page 2145.

For more information on health policies and the health modules you can run to test system health, see the following topics:

- [Understanding Health Policies](#) on page 2111
- [Understanding Health Modules](#) on page 2112
- [Understanding Health Monitoring Configuration](#) on page 2115

Understanding Health Policies

LICENSE: Any

A *health policy* is a collection of health module settings you apply to an appliance to define the criteria that the Defense Center uses when checking the health of the appliance. The health monitor tracks a variety of health indicators to ensure that your Sourcefire 3D System hardware and software are working correctly.

When you create health policies, you choose which tests to run to determine appliance health. You can also apply the default health policy to any appliance.

Understanding Health Modules

LICENSE: Any

Health modules, also sometimes referred to as *health tests*, are scripts that test for the criteria you specify in a health policy. The available health modules are described in the [Health Modules](#) table.

Health Modules

MODULE	DESCRIPTION
Advanced Malware Protection	This module alerts if the Defense Center cannot contact the Sourcefire cloud to retrieve file disposition information for files detected in network traffic. This module runs on all Defense Centers except the DC500, which does not support advanced malware protection.
Appliance Heartbeat	This module determines if an appliance heartbeat is being heard from the appliance and alerts based on the appliance heartbeat status.
Automatic Application Bypass Status	This module determines if an appliance has been bypassed because it did not respond within the number of seconds set in the bypass threshold, and alerts when a bypass occurs.
CPU Usage	This module checks that the CPU on the appliance is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. This module is not available for health policies applied to 3D9900 devices.
Card Reset	This module checks for network cards which have restarted due to hardware failure and alerts when a reset occurs.
Discovery Event Status	This module indicates whether a specified period of time has passed since any discovery events have been detected by a device.
Disk Status	This module examines performance of the hard disk on the appliance and alerts when the hard disks and RAID controller (if installed) are in danger of failing.
Disk Usage	This module compares disk usage on the appliance to the limits configured for the module and alerts when usage exceeds the percentages configured for the module.
FireAMP Status Monitor	The module alerts if the Defense Center cannot connect to the Sourcefire cloud after an initial successful connection, or if you deregister a cloud connection using the FireAMP portal. This module only runs on Defense Centers.

Health Modules (Continued)

MODULE	DESCRIPTION
FireSIGHT Host License Limit	<p>This module determines if sufficient FireSIGHT host licenses remain and alerts based on the warning level configured for the module.</p> <p>This module only runs on Defense Centers.</p>
Hardware Alarms	<p>This module determines if hardware needs to be replaced on a Series 3 or 3D9900 device and alerts based on the hardware status. The module also reports on the status of hardware-related daemons and on the status of clustered appliances.</p> <p>For more information on the details reported for these devices, see Interpreting Hardware Alert Details for 3D9900 Devices on page 2179 and Interpreting Hardware Alert Details for Series 3 Devices on page 2181.</p>
Health Monitor Process	<p>This module monitors the status of the health monitor itself and alerts if the number of minutes since the last health event received by the Defense Center exceeds the Warning or Critical limits.</p> <p>This module only runs on Defense Centers.</p>
Intrusion Event Rate	<p>This module compares the number of intrusion events per second to the limits configured for this module and alerts if the limits are exceeded. If the Intrusion Event Rate is zero, the intrusion process may be down or the managed device may not be sending events. Select Analysis > Intrusion > Events to check if events are being received from the device.</p>
License Monitor	<p>This module determines if sufficient licenses for Control, Protection, URL Filtering, Malware, and VPN remain. It also alerts when devices in a stack have mismatched license sets. It alerts based on a warning level automatically configured for the module. You cannot change the configuration of this module.</p> <p>This module only runs on Defense Centers.</p>
Link State Propagation	<p>This module determines when a link in a paired inline set fails and triggers the link state propagation mode.</p>
Memory Usage	<p>This module compares memory usage on the appliance to the limits configured for the module and alerts when usage exceeds the levels configured for the module.</p>
Power Supply	<p>This module determines if power supplies on the device require replacement and alerts based on the power supply status.</p> <p>This module runs on these Defense Centers: DC1500, DC3500.</p> <p>This module runs on these devices: 3D3500, 3D4500, 3D6500, 3D9900, and Series 3.</p>

Health Modules (Continued)

MODULE	DESCRIPTION
Process Status	<p>This module determines if processes on the appliance exit or terminate outside of the process manager. If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited, until the module runs again and the process has restarted. If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process, until the module runs again and the process has restarted.</p>
RRD Server Process	<p>This module determines if the round robin data server that stores time series data is running properly and alerts based on the number of recent RRD server restarts.</p> <p>This module only runs on Defense Centers.</p>
Security Intelligence	<p>This module alerts in a variety of situations involving Security Intelligence filtering, including feed update, feed corruption, and memory issues.</p> <p>This module runs on all Defense Centers except the DC500, which does not support Security Intelligence filtering.</p>
Time Series Data Monitor	<p>This module tracks the presence of corrupt files in the directory where time series data (such as compliance event counts) are stored and alerts when files are flagged as corrupt and removed.</p> <p>This module only runs on Defense Centers.</p>
Time Synchronization Status	<p>This module tracks the synchronization of a device clock that obtains time using NTP with the clock on the NTP server and alerts if the difference in the clocks is more than ten seconds.</p>
Traffic Status	<p>This module determines if the device currently collects traffic and alerts based on the traffic status.</p>

Health Modules (Continued)

MODULE	DESCRIPTION
URL Filtering Monitor	<p>This module tracks communication between the Defense Center and the Sourcefire cloud, where the system obtains its URL filtering (category and reputation) data for commonly visited URLs. The module alerts if the Defense Center fails to successfully communicate with or retrieve an update from the cloud.</p> <p>This module also tracks communications between the Defense Center and any managed devices where you have enabled URL filtering. The module alerts if the Defense Center cannot push URL filtering data to those devices.</p> <p>This module only runs on all Defense Centers except the DC500, which does not support URL filtering.</p>
User Agent Status Monitor	<p>This module alerts when heartbeats are not detected for any Sourcefire User Agents connected to the Defense Center.</p> <p>This module only runs on Defense Centers.</p>
VPN Status	<p>This module alerts when the system detects that the VPN feature is not functioning.</p> <p>This module only runs on Defense Centers.</p>

Understanding Health Monitoring Configuration

LICENSE: Any

There are several steps to setting up health monitoring on your Sourcefire 3D System, as indicated in the following procedure:

1. Create health policies for your appliances.

You can set up specific policies for each kind of appliance you have in your Sourcefire 3D System, enabling only the appropriate tests for that appliance.

TIP! If you want to quickly enable health monitoring without customizing the monitoring behavior, you can apply the default policy provided for that purpose.

For more information on setting up health policies, see [Configuring Health Policies](#) on page 2116.

2. Apply a health policy to each appliance where you want to track health status. For information on the default health policy available for immediate application, see [Understanding the Default Health Policy](#) on page 2117.

3. Optionally, configure health monitor alerts.

You can set up email, syslog, or SNMP alerts that trigger when the health status level reaches a particular severity level for specific health modules.

For more information on setting up health monitor alerts, see [Configuring Health Monitor Alerts](#) on page 2159.

After you set up health monitoring on your system, you can view the health status at any time on the Health Monitor page or the Health Events table view. For more information about viewing system health data, see the following topics:

- [Using the Health Monitor](#) on page 2163
- [Using Appliance Health Monitors](#) on page 2166
- [Working with Health Events](#) on page 2174

Configuring Health Policies

LICENSE: Any

A health policy contains configured health test criteria for several modules. You can control which health modules run against each of your appliances and configure the specific limits used in the tests run by each module. For more information on the health modules you can configure in a health policy, see [Understanding Health Monitoring](#) on page 2110.

You can create one health policy that can be applied to every appliance in your system, customize each health policy to the specific appliance where you plan to apply it, or use the default health policy provided for you. You can also import a health policy exported from another Defense Center.

When you configure a health policy, you decide whether to enable each health module for that policy. You also select the criteria that control which health status each enabled module reports each time it assesses the health of a process.

For more information on the default health policy, which is applied to the Defense Center automatically, see [Understanding the Default Health Policy](#) on page 2117.

For more information, see the following topics:

- [Understanding the Default Health Policy](#) on page 2117
- [Creating Health Policies](#) on page 2118
- [Applying Health Policies](#) on page 2145
- [Editing Health Policies](#) on page 2147
- [Comparing Health Policies](#) on page 2150
- [Deleting Health Policies](#) on page 2154

Understanding the Default Health Policy

LICENSE: Any

The Defense Center health monitor includes a default health policy to make it easier for you to quickly implement health monitoring for your appliances. The default health policy is automatically applied to the Defense Center. You cannot edit the default health policy, but you can copy it to create custom policies based on its configuration. For more information, see [Creating Health Policies](#) on page 2118.

To also monitor device health, you can push health policies to your managed devices.

In the default health policy, most of the health modules available on the running platform are automatically enabled. The [Default Active Health Modules](#) table details the modules activated in the default policy for Defense Centers and managed devices.

Default Active Health Modules

MODULE	DEFENSE CENTER	MANAGED DEVICE
Advanced Malware Protection	X	
Appliance Heartbeat	X	
Automatic Application Bypass		X
CPU Usage		
Card Reset		
Discovery Event Status		
Disk Status	X	X
Disk Usage	X	X
FireAMP Status Monitor	X	
FireSIGHT Host License Limit	X	
Hardware Alarm		X
Health Monitor Process		
Intrusion Event Rate		X
License Monitor	X	

Default Active Health Modules (Continued)

MODULE	DEFENSE CENTER	MANAGED DEVICE
Link State Propagation		X
Memory Usage	X	X
Power Supply		X
Process Status	X	X
RRD Server Process	X	
Security Intelligence	X	
Time Series Data Monitor	X	
Time Synchronization Status	X	X
Traffic Status		X
URL Filtering Monitor	X	
User Agent Status Monitor	X	
VPN Status	X	

Creating Health Policies

LICENSE: Any

If you want to customize a health policy to use with your appliances, you can create a new policy. The settings in the policy initially populate with the settings from the health policy you select as a basis for the new policy. You can enable or disable modules within the policy and change the alerting criteria for each module as needed.

TIP! Instead of creating a new policy, you can export a health policy from another Defense Center and then import it onto your Defense Center. You can then edit the imported policy to suit your needs before you apply it. For more information, see [Importing and Exporting Objects](#) on page 2225.

To create a health policy:

ACCESS: Admin/Maint

1. Select **Health > Health Policy**.

The Health Policy page appears.

Policy Name	Applied To	Last Modified	
Initial_Health_Policy 2012-04-06 16:37:33 Initial Health Policy	None	2012-04-06 12:37:33	   
katsura health policy	4 appliances	2012-04-06 16:11:00	   

2. Click **Create Policy**.

The Create Health Policy page appears.

Create Policy

Copy Policy Sourcefire Default Health Policy 

New Policy Name

New Policy Description

3. Select the existing policy that you want to use as the basis for the new policy from the **Copy Policy** drop-down list.
 4. Enter a name for the policy.
 5. Enter a description for the policy.
 6. Select **Save** to save the policy information.
- The Health Policy Configuration page appears, including a list of the modules.
7. Configure settings on each module you want to use to test the health status of your appliances, as described in the following sections:
 - [Configuring Policy Run Time Intervals](#) on page 2120
 - [Configuring Advanced Malware Protection Monitoring](#) on page 2121
 - [Configuring Appliance Heartbeat Monitoring](#) on page 2122
 - [Configuring Automatic Application Bypass Monitoring](#) on page 2123
 - [Configuring CPU Usage Monitoring](#) on page 2124
 - [Configuring Card Reset Monitoring](#) on page 2125
 - [Configuring Discovery Event Status Monitoring](#) on page 2126
 - [Configuring Disk Status Monitoring](#) on page 2127
 - [Configuring Disk Usage Monitoring](#) on page 2127
 - [Configuring FireAMP Status Monitoring](#) on page 2128
 - [Configuring FireSIGHT Host Usage Monitoring](#) on page 2129
 - [Configuring Hardware Alarm Monitoring](#) on page 2130
 - [Configuring Health Status Monitoring](#) on page 2131

- [Configuring Intrusion Event Rate Monitoring](#) on page 2132
- [Understanding License Monitoring](#) on page 2133
- [Configuring Link State Propagation Monitoring](#) on page 2134
- [Configuring Memory Usage Monitoring](#) on page 2134
- [Configuring Power Supply Monitoring](#) on page 2136
- [Configuring Process Status Monitoring](#) on page 2137
- [Configuring RRD Server Process Monitoring](#) on page 2138
- [Configuring Security Intelligence Monitoring](#) on page 2139
- [Configuring Time Series Data Monitoring](#) on page 2140
- [Configuring Time Synchronization Monitoring](#) on page 2141
- [Configuring Traffic Status Monitoring](#) on page 2142
- [Configuring URL Filtering Monitoring](#) on page 2142
- [Configuring User Agent Status Monitoring](#) on page 2143
- [Configuring VPN Status Monitoring](#) on page 2144

IMPORTANT! Make sure you enable each module that you want to run to test the health status on each Health Policy Configuration page as you configure the settings. Disabled modules do not produce health status feedback, even if the policy that contains the module has been applied to an appliance.

8. Click **Save Policy and Exit** to save the policy.

You must apply the policy to each appliance for it to take effect. For more information on applying health policies, see [Applying Health Policies](#) on page 2145.

Configuring Policy Run Time Intervals

LICENSE: Any

You can control how often health tests run by modifying the Policy Run Time Interval for the health policy. The maximum run time interval you can set is 99999 minutes.

WARNING! Do not set a run interval of less than five minutes.

To configure a policy run time interval:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, select **Policy Run Time Interval**.
The Health Policy Configuration — Policy Run Time Interval page appears.



Run Interval (mins)

2. In the **Run Interval (mins)** field, enter the time in minutes that you want to elapse between automatic repetitions of the test.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Advanced Malware Protection Monitoring

LICENSE: Malware

This module tracks the state and stability of the Defense Center's ability to query the Sourcefire cloud using malware cloud lookups in file policies. If the system detects that your connection with the cloud is interrupted, or if the encryption keys used for the connection are invalid, the status classification for this module changes to Warning and the module generates a health alert.

To configure Advanced Malware Protection health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select Advanced Malware Protection.
The Health Policy Configuration — Advanced Malware Protection page appears.



Description Advanced Malware Protection
Enabled On Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Appliance Heartbeat Monitoring

LICENSE: Any

The Defense Center receives heartbeats from its managed devices once every two minutes or every 200 events, whichever comes first, as an indicator that the device is running and communicating properly with the Defense Center. Use the Appliance Heartbeat health status module to track whether the Defense Center receives heartbeats from managed appliances. If the Defense Center does not detect a heartbeat from a device, the status classification for this module changes to Critical. That status data feeds into the health monitor.

To configure Appliance Heartbeat health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Appliance Heartbeat**.
The Health Policy Configuration — Appliance Heartbeat page appears.

Description	Monitors if the attached appliances are alive
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Automatic Application Bypass Monitoring

LICENSE: Any

Use this module to detect when a managed device is bypassed because it did not respond within the number of seconds configured as the bypass threshold. If a bypass occurs, this module generates an alert. That status data feeds into the health monitor.

For more information on automatic application bypass, see [Automatic Application Bypass](#) on page 251.

To configure automatic application bypass monitoring status:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Automatic Application Bypass Status**.

The Health Policy Configuration — Automatic Application Bypass Status page appears.

Description	Monitors bypassed detection applications
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate managed device if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring CPU Usage Monitoring

LICENSE: Any

SUPPORTED DEVICES: Any except 3D9900

SUPPORTED DEFENSE CENTERS: Any

Excessive CPU usage may indicate that you need to upgrade your hardware or that there are processes that are not functioning correctly. Use the CPU Usage health status module to set CPU usage limits.

If the CPU usage on the monitored appliance exceeds the Warning limit, the status classification for that module changes to Warning. If the CPU usage on the monitored appliance exceeds the Critical limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

The maximum percentage you can set for either limit is 100 percent, and the Critical limit must be higher than the Warning limit.

To configure CPU usage limits:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, select **CPU Usage**.
The Health Policy Configuration — CPU Usage page appears.

Description	Monitors CPU Usage
Enabled	<input type="radio"/> On <input checked="" type="radio"/> Off
Critical Threshold %	<input type="text" value="90"/>
Warning Threshold %	<input type="text" value="80"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Threshold %** field, enter the percentage of CPU usage that should trigger a critical health status.
4. In the **Warning Threshold %** field, enter the percentage of CPU usage that should trigger a warning health status.

5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

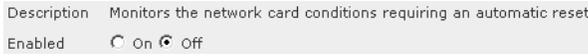
Configuring Card Reset Monitoring

LICENSE: Any

Use the card reset monitoring health status module to track when the network card restarts because of hardware failure. If a reset occurs, this module generates an alert. That status data feeds into the health monitor.

To configure card reset monitoring:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Card Reset**.
The Health Policy Configuration — Card Reset Monitoring page appears.

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate Defense Center if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Discovery Event Status Monitoring

LICENSE: FireSIGHT

Use the Discovery Event Status module to monitor the health of the discovery process on a device from the Defense Center by generating alerts when too many seconds elapse between discovery events received by the Defense Center. You can configure the elapsed duration between events, in seconds, that causes an alert to be generated. If the wait exceeds the number of seconds configured in the Warning Seconds since last event limit, the status classification for that module changes to Warning. If the wait exceeds the Critical Seconds since last event limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

The maximum number of seconds you can set for either limit is 7200, and the Critical limit must be higher than the Warning limit. The minimum number of seconds is 3600.

To configure Discovery Event Status module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Discovery Event Status**.

The Health Policy Configuration — Discovery Event Status page appears.

Description	Ensures that Discovery is reporting events
Enabled	<input type="radio"/> On <input checked="" type="radio"/> Off
Critical Seconds since last event	<input type="text" value="7200"/>
Warning Seconds since last event	<input type="text" value="3600"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Seconds since last event** field, enter the maximum number of seconds to wait between events, before triggering a critical health status.
4. In the **Warning Seconds since last event** field, enter the maximum number of seconds to wait between events, before triggering a warning health status.
5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the Defense Center for your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Disk Status Monitoring

LICENSE: Any

Use the Disk Status health module to monitor the current status of your appliance's hard disk. This module generates a Warning (yellow) health alert when the hard disk and RAID controller (if installed) are in danger of failing.

To configure Disk Status health module settings:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, click **Disk Status**.

The Health Policy Configuration — Disk Status page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Disk Usage Monitoring

LICENSE: Any

Without sufficient disk space, an appliance cannot run. The health monitor can identify low disk space conditions on your appliances before the space runs out. Use the Disk Usage health status module to set disk usage limits for the / and /volume partitions on the appliance.

IMPORTANT! Although the disk usage module lists the /boot partition as a monitored partition, the size of the partition is static so the module does not alert on the boot partition.

If the disk usage on the monitored appliance exceeds the Warning limit, the status classification for that module changes to Warning. If the disk usage on the

monitored appliance exceeds the Critical limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

The maximum percentage you can set for either limit is 100 percent, and the Critical limit must be higher than the Warning limit.

To configure Disk Usage health module settings:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, select **Disk Usage**.

The Health Policy Configuration — Disk Usage page appears.

Description	Monitors Disk Usage
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off
Critical Threshold %	<input type="text" value="90"/>
Warning Threshold %	<input type="text" value="85"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Threshold %** field, enter the percentage of disk usage that should trigger a critical health status.
4. In the **Warning Threshold %** field, enter the percentage of disk usage that should trigger a warning health status.
5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring FireAMP Status Monitoring

LICENSE: Any

Use the FireAMP Status Monitor module to alert you in the following situations:

- the Defense Center cannot connect to the Sourcefire cloud after an initial successful connection
- you deregister a cloud connection using the FireAMP portal

In these cases, the module status changes to Critical and provides the cloud name associated with the failed connection. For information on configuring a

cloud connection, see [Working with Sourcefire Cloud Connections for FireAMP](#) on page 1219.

To configure FireAMP Status Monitor module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **FireAMP Status Monitor**.
The Health Policy Configuration — FireAMP Status Monitor page appears.



Description	FireAMP Status Monitor
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for FireAMP status monitoring.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the Defense Center if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring FireSIGHT Host Usage Monitoring

LICENSE: FireSIGHT

Use the FireSIGHT Host License Limit health status module to set FireSIGHT Host amount warning limits. If the number of remaining FireSIGHT Hosts on the monitored device falls below the Warning Hosts limit, the status classification for that module changes to Warning. If the number of remaining FireSIGHT Hosts on the monitored device falls below the Critical Hosts limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

The maximum number of hosts you can set for either limit is 1000, and the Critical host limit number must be lower than the Warning limit.

To configure FireSIGHT Host License Limit health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **FireSIGHT Host License Limit**. The Health Policy Configuration — FireSIGHT Host License Limit page appears.

Description	Monitors FireSIGHT Host License Usage
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off
Critical number Hosts	<input type="text" value="10"/>
Warning number Hosts	<input type="text" value="50"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical number Hosts** field, enter the remaining number of available hosts that should trigger a critical health status.
4. In the **Warning number Hosts** field, enter the remaining number of available hosts that should trigger a warning health status.
5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Hardware Alarm Monitoring

LICENSE: Any

SUPPORTED DEVICES: Series 3, 3D9900

Use the Hardware Alarm health status module to detect hardware failure on a Series 3 or 3D9900 device. If the Hardware Alarm module finds a hardware component that has failed or clustered devices that are not communicating with each other, the status classification for that module changes to Critical. That status data feeds into the health monitor.

For more information on the hardware status conditions that can cause hardware alerts on 3D9900 devices, see [Interpreting Hardware Alert Details for 3D9900 Devices](#) on page 2179.

To configure Hardware Alarm health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Hardware Alarms**.

The Health Policy Configuration — Hardware Alarm Monitor page appears.

Description	Monitor any alarm sent by the operating system or network cards
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Health Status Monitoring

LICENSE: Any

Use the Health Monitor Process module to monitor the health of the health monitor on a Defense Center by generating alerts when too many minutes elapse between health events received from monitored appliances.

For example, if a Defense Center ([myrtle.example.com](#)) monitors a device ([dogwood.example.com](#)), you apply a health policy with the Health Monitor Process module enabled to [myrtle.example.com](#). The Health Monitor Process module then reports events that indicate how many minutes have elapsed since the last event was received from [dogwood.example.com](#).

You can configure the elapsed duration between events, in minutes, that causes an alert to be generated. If the wait exceeds the number of minutes configured in the Warning Minutes since last event limit, the status classification for that module changes to Warning. If the wait exceeds the Critical Minutes since last event limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

The maximum number of minutes you can set for either limit is 144, and the Critical limit must be higher than the Warning limit. The minimum number of minutes is 5.

To configure Health Monitor Process module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Health Monitor Process**.

The Health Policy Configuration — Health Monitor Process page appears.

Description	Monitors the status of the Health Monitor itself
Enabled	<input type="radio"/> On <input checked="" type="radio"/> Off
Critical Minutes since last event	<input type="text" value="60"/>
Warning Minutes since last event	<input type="text" value="30"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Minutes since last event** field, enter the maximum number of minutes to wait between events, before triggering a critical health status.
4. In the **Warning Minutes since last event** field, enter the maximum number of minutes to wait between events, before triggering a warning health status.
5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the Defense Center for your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Intrusion Event Rate Monitoring

LICENSE: Protection

Use the Intrusion Event Rate health status module to set limits for the number of packets per second that trigger a change in the health status. If the event rate on the monitored device exceeds the number of events per second configured in the Events per second (Warning) limit, the status classification for that module changes to Warning. If the event rate exceeds the number of events per second configured in the Events per second (Critical) limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

Typically, the event rate for a network segment averages 20 events per second. For a network segment with this average rate, Events per second (Critical) should be set to 50 and Events per second (Warning) should be set to 30. To determine limits for your system, find the Events/Sec value on the Statistics page for your

device (**System > Monitoring > Statistics**), then calculate the limits using these formulas:

- Events per second (Critical) = Events/Sec * 2.5
- Events per second (Warning) = Events/Sec * 1.5

The maximum number of events you can set for either limit is 999, and the Critical limit must be higher than the Warning limit.

To configure **Intrusion Event Rate Monitor** health module settings:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, select **Intrusion Event Rate**.

The Health Policy Configuration — Intrusion Event Rate page appears.

Description	Monitors the Events per second from Snort
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off
Events per second (Critical)	<input type="text" value="50"/>
Events per second (Warning)	<input type="text" value="30"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Events per second (Critical)** field, enter the number of events per second that should trigger a critical health status.
4. In the **Events per second (Warning)** field, enter the number of events per second that should trigger a warning health status.
5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Understanding License Monitoring

LICENSE: Any

Use the License Monitoring health status module to determine if sufficient licenses remain for Control, Protection, URL Filtering, Malware, and VPN. This module alerts if the number of remaining licenses is low or insufficient.

This module also alerts if the system detects that devices in a stacked configuration have mismatched license sets (stacked devices must have identical sets of licenses).

The License Monitoring module is automatically configured. Because you cannot change or disable this module, it does not appear on the Health Policy Configuration page.

Configuring Link State Propagation Monitoring

LICENSE: Any

Use the Link State Propagation health status module to detect the link state propagation status on an inline pair. If a link state propagates to the pair, the status classification for that module changes to Critical and the state reads:

Module Link State Propagation: ethx_ethyl is Triggered

where *x* and *y* are the paired interface numbers.

To configure Link State Propagation health module settings:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, select **Link State Propagation**.

The Health Policy Configuration — Link State Propagation monitor page appears.



Description	Monitor Link State Propagation
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Memory Usage Monitoring

LICENSE: Any

Use the Memory Usage health status module to set memory usage limits. The module calculates free memory by considering free memory, cached memory,

and swap memory. If the memory usage on the monitored appliance exceeds the Warning limit, the status classification for that module changes to Warning. If the memory usage on the monitored appliance exceeds the Critical limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

For appliances with more than 4GB of memory, the preset alert thresholds are based on a formula that accounts for proportions of available memory likely to cause system problems.

IMPORTANT! On >4GB appliances, because the interval between Warning and Critical thresholds may be very narrow, Sourcefire recommends that you manually set the **Warning Threshold %** value to 50. This will further ensure that you receive memory alerts for your appliance in time to address the issue.

The maximum percentage you can set for either limit is 100 percent, and the Critical limit must be higher than the Warning limit.

To configure Memory Usage health module settings:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, select **Memory Usage**.

The Health Policy Configuration — Memory Usage page appears.

Description	Monitors Memory Usage
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off
Critical Threshold %	<input type="text" value="90"/>
Warning Threshold %	<input type="text" value="80"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Threshold %** field, enter the percentage of memory usage that should trigger a critical health status.
4. In the **Warning Threshold %** field, enter the percentage of memory usage that should trigger a warning health status.

5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Power Supply Monitoring

LICENSE: Any

SUPPORTED DEVICES: 3D3500, 3D4500, 3D6500, 3D9900, Series 3

SUPPORTED DEFENSE CENTERS: DC1500, DC3500

Use the Power Supply health status module to detect a power supply failure on any of the supported platforms. If the module finds a power supply that has no power, the status classification for that module changes to No Power. If the module cannot detect the presence of the power supply, the status changes to Critical Error. That status data feeds into the health monitor. You can expand the Power Supply item on the Alert Detail list in the health monitor to see specific status items for each power supply.

To configure Power Supply health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Power Supply**.
The Health Policy Configuration — Power Supply page appears.



Description	Monitor Power supplies
Enabled	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Process Status Monitoring

LICENSE: Any

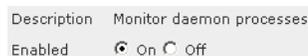
Use the Process Status health module to monitor for processes running on the appliance that exit or terminate outside of the process manager. The response of the Process Status module to a process ending depends on how the process ends:

- If the process terminates inside the process manager, the module does not report any health events.
- If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited until the module runs again and the process has restarted.
- If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process until the module runs again and the process has restarted.

To configure Process Status health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Process Status**.
The Health Policy Configuration — Process Status page appears.



Description Monitor daemon processes
Enabled On Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate appliances if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring RRD Server Process Monitoring

LICENSE: Any

Use the RRD Server Process module to see if the RRD server that stores time series data is working properly. The module will alert if the RRD server has restarted since the last time it updated; it will enter Critical or Warning status if the number of consecutive updates with an RRD server restart reaches the numbers specified in the module configuration.

To configure RRD server process monitoring settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **RRD Server Process**.
The Health Policy Configuration — RRD Server Process page appears.

Description	Monitors the status of the rrd_server process
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off
Critical Number of restarts	<input type="text" value="3"/>
Warning Number of restarts	<input type="text" value="2"/>

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Number of restarts** field, enter the number of consecutive detected RRD server resets that should trigger a critical health status.
4. In the **Warning Number of restarts** field, enter the number of consecutive detected RRD server resets that should trigger a warning health status.

5. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Security Intelligence Monitoring

LICENSE: Protection

SUPPORTED DEFENSE CENTERS: Any except DC500

Use the Security Intelligence module to warn you in a variety of situations involving Security Intelligence filtering. The module alerts if Security Intelligence is in use and:

- the Defense Center cannot update a feed, or if feed data is corrupt or contains no recognizable IP addresses
- a managed device had a problem receiving updated Security Intelligence data from the Defense Center
- a managed device cannot load all of the Security Intelligence data provided to it by the Defense Center, due to memory issues

TIP! If a Security Intelligence memory warning appears in the health monitor, you can reapply the affected device's access control policy to increase the memory allocated to Security Intelligence. See [Applying an Access Control Policy](#) on page 466 for more information.

For more information on Security Intelligence filtering, see [Filtering Traffic Based on Security Intelligence Data](#) on page 438 and [Working with Security Intelligence Lists and Feeds](#) on page 162.

To configure Security Intelligence module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Security Intelligence**. The Health Policy Configuration — Security Intelligence page appears.

Description	Security Intelligence
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for Security Intelligence monitoring.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Time Series Data Monitoring

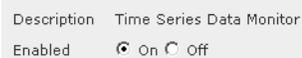
LICENSE: Any

Use the Time Series Data Monitor module to monitor the status of time series data (such as lists of compliance events) that your system has stored. This module scans your time series data storage directory for corrupt files. If the module finds corrupted data, it enters a Warning status and reports the names of all affected files.

To configure time series data monitoring settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Time Series Data Monitor**. The Health Policy Configuration — Time Series Data Monitor page appears.



Description Time Series Data Monitor
Enabled On Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Time Synchronization Monitoring

LICENSE: Any

Use the Time Synchronization Status module to detect when the time on a managed device that uses NTP to obtain time from an NTP server differs by 10 seconds or more from the time on the server.

To configure time synchronization monitoring settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Time Synchronization Status**.
The Health Policy Configuration — Time Synchronization Status page appears.

Description	Monitors the time difference of managed devices
Enabled	<input checked="" type="radio"/> On <input type="radio"/> Off

2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring Traffic Status Monitoring

LICENSE: FireSIGHT

Use the Traffic Status health status module to detect whether a device receives traffic. If the Traffic Status module determines that a device does not receive traffic, the status classification for that module changes to Critical. That status data feeds into the health monitor.

To configure Traffic Status health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **Traffic Status**.

The Health Policy Configuration — Traffic Status page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring URL Filtering Monitoring

LICENSE: URL Filtering

SUPPORTED DEFENSE CENTERS: Any except DC500

Use the URL Filtering Monitor module to track communications between the Defense Center and the Sourcefire cloud, where the system obtains its URL filtering (category and reputation) data for commonly visited URLs. If the Defense Center fails to successfully communicate with or retrieve an update from the cloud, the status classification for that module changes to Critical.

In a high availability configuration, only the primary Defense Center communicates with the URL filtering cloud; all data from this module refers only to that primary appliance.

The URL Filtering Monitor module also tracks communications between the Defense Center and any managed devices where you have enabled URL filtering.

If the Defense Center is successfully communicating with the cloud, the module status changes to Warning if the Defense Center cannot push new URL filtering data to its managed devices.

To configure URL Filtering Monitor health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **URL Filtering Monitor**.
The Health Policy Configuration — URL Filtering Monitor page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the Defense Center if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring User Agent Status Monitoring

LICENSE: FireSIGHT

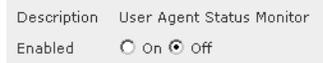
You can use the Sourcefire User Agent Status Monitor health module to monitor the heartbeat of agents connected to a Defense Center. If you enable the module in an applied health policy, the module generates a health alert if the Defense Center does not detect a heartbeat for any agent configured on the Defense Center.

If you have legacy user agents reporting to your Defense Center, the health module alerts that the agents are not issuing heartbeats. If you are using a 5.1 Defense Center with the legacy user agent, Sourcefire recommends you disable the User Agent Status Monitor.

To configure **User Agent Status Monitor** health module settings:

ACCESS: Admin/Maint

1. In the Health Policy Configuration page, select **User Agent Status Monitor**.
The Health Policy Configuration — User Agent Status Monitor page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the Defense Center if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Configuring VPN Status Monitoring

LICENSE: VPN

SUPPORTED DEFENSE CENTERS: Any except Series 2

Use the VPN Status health module to monitor the current status of your configured Gateway VPN tunnels; information for each individual tunnel is displayed. This module generates a Critical (red) health alert when any of your VPN tunnels is not working.

To configure **VPN Status** health module settings:

ACCESS: Admin/Maint

1. On the Health Policy Configuration page, click **VPN Status**.
The Health Policy Configuration — VPN Status page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

3. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.

You must apply the health policy to the appropriate devices if you want your settings to take effect. See [Applying Health Policies](#) on page 2145 for more information.

Applying Health Policies

LICENSE: Any

When you apply a health policy to an appliance, the health tests for all the modules you enabled in the policy automatically monitor the health of the processes and hardware on the appliance. Health tests then continue to run at the intervals you configured in the policy, collecting health data for the appliance and forwarding that data to the Defense Center.

If you enable a module in a health policy and then apply the policy to an appliance that does not require that health test, the health monitor reports the status for that health module as disabled.

If you apply a policy with all modules disabled to an appliance, it removes all applied health policies from the appliance so no health policy is applied.

When you apply a different policy to an appliance that already has a policy applied, expect some latency in the display of new data based on the newly applied tests.

IMPORTANT! Custom health policies created on Defense Centers in a high availability pair will be replicated between both appliances. However, changes to default health policies are not replicated; each appliance uses the local default health policy configured for that appliance.

To apply a health policy:

ACCESS: Admin/Maint

1. Select **Health > Health Policy**.

The Health Policy page appears.

Policy Name	Applied To	Last Modified	
Initial_Health_Policy 2012-04-06 16:37:33 Initial Health Policy	None	2012-04-06 12:37:33	   
katsura health policy	4 appliances	2012-04-06 16:11:00	   

2. Click the apply icon () next to the policy you want to apply.

The Health Policy Apply page appears.

By Group ▼

▼ Ungrouped (4 total)			
<input type="checkbox"/> katsura 10.10.10.3 - Defense Center 3500 v5.1.0	Health Policy katsura health policy 	System Policy katsura system policy	
<input type="checkbox"/> linden 10.10.10.4 - 3D8250 v5.0.2.1	Health Policy linden health policy 	System Policy katsura system policy	
<input type="checkbox"/> tamarix 10.10.10.5 - 3D8140 v5.1.0	 Health Policy katsura health policy 	System Policy katsura system policy	
<input type="checkbox"/> xiramat 10.10.10.6 - 3D8140 v5.1.0	 Health Policy katsura health policy 	System Policy katsura system policy	

TIP! The status icon () next to the Health Policy column indicates the current health status for the appliance. The status icon () next to the System Policy column indicates the communication status between the Defense Center and the device. Note that you can remove the currently applied policy by clicking the remove icon ().

3. Select the appliances where you want to apply the health policy.
4. Click **Apply** to apply the policy to the selected appliances.

The Health Policy page appears, with a message indicating if the application of the policy was successful. Monitoring of the appliance starts as soon as the policy is successfully applied.

To unapply a health policy:

ACCESS: Admin/Maint

1. Select **Health > Health Policy**.

The Health Policy page appears.

Policy Name	Applied To	Last Modified	
Initial_Health_Policy 2012-04-06 16:37:33 Initial Health Policy	None	2012-04-06 12:37:33	   
katsura health policy	4 appliances	2012-04-06 16:11:00	   

- Click the apply icon (✔) next to the policy you want to unapply.
The Health Policy Apply page appears.

▼ Ungrouped (4 total)		
<input type="checkbox"/> katsura 10.10.10.3 - Defense Center 3500 v5.1.0	Health Policy katsura health policy ⓘ	System Policy katsura system policy
<input type="checkbox"/> linden 10.10.10.4 - 3D8250 v5.0.2.1	Health Policy linden health policy ✔	System Policy katsura system policy ✔
<input type="checkbox"/> tamarix 10.10.10.5 - 3D8140 v5.1.0	Health Policy katsura health policy ⓘ	System Policy katsura system policy ✔
<input type="checkbox"/> xiramat 10.10.10.6 - 3D8140 v5.1.0	Health Policy katsura health policy ⓘ	System Policy katsura system policy ✔

Apply

- You have two options:
 - Apply a health policy with all modules disabled.
 - Click the x next to the health policy.

Under **Health Policy**, the status of **None** appears.

Editing Health Policies

LICENSE: Any

You can modify a health policy by enabling or disabling modules or by changing module settings. If you modify a policy that is already applied to an appliance, the changes do not take effect until you reapply the policy.

Applicable health models for various appliances are listed in the [Health Modules Applicable to Appliances](#) table.

Health Modules Applicable to Appliances

MODULE	APPLICABLE APPLIANCE
Advanced Malware Protection	Defense Centers, except DC500
Appliance Heartbeat	Defense Center
Automatic Application Bypass Status	Any managed device
CPU Usage	Any except 3D9900
Card Reset	Any managed device
Discovery Event Status	Defense Center
Disk Status	Any
Disk Usage	Any

Health Modules Applicable to Appliances (Continued)

MODULE	APPLICABLE APPLIANCE
FireAMP Status Monitor	Defense Center
FireSIGHT Host License Limit	Defense Center
Hardware Alarms	Series 3, 3D9900
Health Monitor Process	Defense Center
Intrusion Event Rate	Managed devices with Protection
License Monitor	Defense Center
Link State Propagation	Managed devices with Protection
Memory Usage	Any
Power Supply	Defense Centers: DC1500, DC3500 Devices: 3D3500, 3D4500, 3D6500, 3D9900, Series 3
Process Status	Any
RRD Server Process	Defense Center
Security Intelligence	Defense Center, except DC500
Time Series Data Monitor	Defense Center
Time Synchronization Status	Any
Traffic Status	Any managed device
URL Filtering Monitor	Defense Centers, except DC500
User Agent Status Monitor	Defense Center
VPN Status	Defense Center

To edit a health policy:

ACCESS: Admin/Maint

1. Select **Health > Health Policy**.

The Health Policy page appears.

Policy Name	Applied To	Last Modified	
Initial_Health_Policy 2012-04-06 16:37:33 Initial Health Policy	None	2012-04-06 12:37:33	   
katsura health policy	4 appliances	2012-04-06 16:11:00	   

2. Click the edit icon () next to the policy you want to modify.

The Health Policy Configuration page appears, with the Policy Run Time Interval settings selected.

3. Modify settings as needed, as described in the following sections:

- [Configuring Policy Run Time Intervals](#) on page 2120
- [Configuring Advanced Malware Protection Monitoring](#) on page 2121
- [Configuring Appliance Heartbeat Monitoring](#) on page 2122
- [Configuring Automatic Application Bypass Monitoring](#) on page 2123
- [Configuring CPU Usage Monitoring](#) on page 2124
- [Configuring Card Reset Monitoring](#) on page 2125
- [Configuring Discovery Event Status Monitoring](#) on page 2126
- [Configuring Disk Status Monitoring](#) on page 2127
- [Configuring Disk Usage Monitoring](#) on page 2127
- [Configuring FireAMP Status Monitoring](#) on page 2128
- [Configuring FireSIGHT Host Usage Monitoring](#) on page 2129
- [Configuring Hardware Alarm Monitoring](#) on page 2130
- [Configuring Health Status Monitoring](#) on page 2131
- [Configuring Intrusion Event Rate Monitoring](#) on page 2132
- [Understanding License Monitoring](#) on page 2133
- [Configuring Link State Propagation Monitoring](#) on page 2134
- [Configuring Memory Usage Monitoring](#) on page 2134
- [Configuring Power Supply Monitoring](#) on page 2136
- [Configuring Process Status Monitoring](#) on page 2137
- [Configuring RRD Server Process Monitoring](#) on page 2138
- [Configuring Security Intelligence Monitoring](#) on page 2139
- [Configuring Time Series Data Monitoring](#) on page 2140
- [Configuring Time Synchronization Monitoring](#) on page 2141
- [Configuring Traffic Status Monitoring](#) on page 2142

- [Configuring URL Filtering Monitoring](#) on page 2142
 - [Configuring User Agent Status Monitoring](#) on page 2143
 - [Configuring VPN Status Monitoring](#) on page 2144
4. You have three options:
 - To save your changes to this module and return to the Health Policy page, click **Save Policy and Exit**.
 - To return to the Health Policy page without saving any of your settings for this module, click **Cancel**.
 - To temporarily save your changes to this module and switch to another module's settings to modify, select the other module from the list at the left of the page. If you click **Save Policy and Exit** when you are done, all changes you made will be saved; if you click **Cancel**, you discard all changes.
 5. Reapply the policy to the appropriate appliances as described in [Applying Health Policies](#) on page 2145.

Comparing Health Policies

LICENSE: Any

To review policy changes for compliance with your organization's standards or to optimize health monitoring performance, you can examine the differences between two health policies. You can compare any two health policies or two revisions of the same health policy, for the health policies you can access. To quickly compare your active health policy to another, you can select the **Running Configuration** option. 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 health policies or health policy revisions:

- The comparison view displays only the differences between two health policies or health 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 web interface, with their differences highlighted.

- The comparison report creates a record of only the differences between two health policies or health policy revisions in a format similar to the health 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 health policy comparison tools, see:

- [Using the Health Policy Comparison View](#) on page 2151
- [Using the Health Policy Comparison Report](#) on page 2152

Using the Health Policy Comparison View

LICENSE: Any

The comparison view displays both health 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 Health Policy page displays the time a policy was last modified in local time, but the health policy report lists the time modified in UTC.

Custom Policy 1 (2011-09-27 12:01:58 by admin)	Sourcefire Default Health Policy (2011-09-24 18:13:36)
Policy Information	
Name: Custom Policy 1	Name
Modified: 2011-09-27 12:01:58 by admin	Description
Applied To: katsura	Modified
Modules	
Appliance Heartbeat: Disabled	Applied To
Automatic Application Bypass: Disabled	Appliance Heartbeat
Disk Usage: Disabled	Automatic Application Bypass
	Disk Usage
	Disk Usage
	Red
	Yellow

Differences between the two health 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 in the [Health Policy Comparison View Actions](#) table.

Health 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 health policy comparison view	click New Comparison . The Select Comparison window appears. See Using the Health Policy Comparison Report for more information.
generate a health policy comparison report	click Comparison Report . The health policy comparison report creates a PDF containing information identical to the comparison view.

Using the Health Policy Comparison Report

LICENSE: Any

A health policy comparison report is a record of all differences between two health policies or two revisions of the same health policy identified by the health policy comparison view, presented as a PDF. You can use this report to further examine the differences between two health policy configurations and to save and disseminate your findings.

You can generate a health policy comparison report from the comparison view for any health policies to which you have access. Remember to commit any potential changes before you generate a health policy report; only committed changes appear in the report.

Depending on your configuration, a health policy comparison report can contain one or more sections. The following sample graphic displays the Policy Information and Modules sections of a health policy comparison report, and lists the configuration for each rule for both health policy configurations. Each section uses the same format and provides the same level of detail. Note that the Value A

and Value B columns represent the policies or policy revisions you configured in the comparison view.

Policy Information

Field	Value A	Value B
Name	Example Health Policy	Test Health Policy
Description	Example Health Policy	Test Health Policy
Modified	2012-05-24 17:01:50 by admin	2012-05-24 17:29:05 by admin

Policy Run Time Interval

Field	Value A	Value B
Run Interval (mins)	5	20

Modules

Field	Value A	Value B
Automatic Application Bypass	Enabled	Disabled
Hardware Alarms	Enabled	Disabled
Process Monitor	Enabled	Disabled
Short Event Rate > Red	50	40
Short Event Rate > Yellow	30	20
Time Synchronization	Enabled	Disabled

You use a similar procedure to compare other types of policies on the Sourcefire 3D System. For more information, see:

- [Comparing Two Access Control Policies](#) on page 462
- [Comparing Two Intrusion Policies](#) on page 685
- [Comparing System Policies](#) on page 1965

To compare two health policies or two revisions of the same policy:

ACCESS: Admin/Maint

1. Select **Health > Health Policy**.

The Health Policy page appears.

2. Click **Compare Policies**.

The Select Comparison window appears.

Select Comparison ? x

Compare Against	<input type="text" value="Other Policy"/>
Policy A	<input type="text" value="Custom Policy 1"/>
Policy B	<input type="text" value="Sourcefire Default Health Policy"/>

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**.
 - To compare another policy to the currently active policy, select **Running Configuration**.

Remember to commit any changes before you generate a health policy report; only committed changes appear in the report.

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.
 - If you are comparing the running configuration to another policy, select the second policy from the **Policy B** drop-down list.
5. Click **OK** to display the health policy comparison view.
The comparison view appears.
6. Click **Comparison Report** to generate the health policy comparison report.
The health policy 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.

Deleting Health Policies

LICENSE: Any

You can delete health policies that you no longer need. If you delete a policy that is still applied to an appliance, the policy settings remain in effect until you apply a different policy. In addition, if you delete a health policy that is applied to a device, any health monitoring alerts in effect for the device remain active until you disable the underlying associated alert response; see [Enabling and Disabling Alert Responses](#) on page 534.

TIP! To stop health monitoring for an appliance, create a health policy with all modules disabled and apply it to the appliance. For more information on creating health policies, see [Creating Health Policies](#) on page 2118. For more information on applying health policies, see [Applying Health Policies](#) on page 2145.

To delete a health policy:

ACCESS: Admin/Maint

1. Select **Health > Health Policy**.
The Health Policy page appears.
2. Click the delete icon () next to the policy you want to delete.
A message appears, indicating if the deletion was successful.

Using the Health Monitor Blacklist

LICENSE: Any

In the course of normal network maintenance, you disable appliances or make them temporarily unavailable. Because those outages are deliberate, you do not want the health status from those appliances to affect the summary health status on your Defense Center.

You can use the health monitor blacklist feature to disable health monitoring status reporting on an appliance or module. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a managed device on that segment to prevent the health status on the Defense Center from displaying a warning or critical state because of the lapsed connection to the device.

When you disable health monitoring status, health events are still generated, but they have a disabled status and do not affect the health status for the health monitor. If you remove the appliance or module from the blacklist, the events that were generated during the blacklisting continue to show a status of disabled.

To temporarily disable health events from an appliance, go to the blacklist configuration page and add an appliance to the blacklist. After the setting takes effect, the system no longer includes the blacklisted appliance when calculating the overall health status. The Health Monitor Appliance Status Summary lists the appliance as disabled.

At times it may be more practical to just blacklist an individual health monitoring module on an appliance. For example, when you run out of FireSIGHT host licenses on an appliance, you can blacklist the FireSIGHT Host License Limit status messages.

Note that on the main Health Monitor page you can distinguish between appliances that are blacklisted if you expand to view the list of appliances with a particular status by clicking the arrow in that status row. For more information on expanding that view, see [Using the Health Monitor](#) on page 2163.

A blacklist icon () and a notation are visible after you expand the view for a blacklisted or partially blacklisted appliance.

IMPORTANT! On a Defense Center, Health Monitor blacklist settings are local configuration settings. Therefore, if you blacklist a device, then delete it and later re-register it with the Defense Center, the blacklist settings remain persistent. The newly re-registered device remains blacklisted.

For more information, see:

- [Blacklisting Health Policies or Appliances](#) on page 2156
- [Blacklisting a Health Policy Module](#) on page 2158

Blacklisting Health Policies or Appliances

LICENSE: Any

If you want to set health events to disabled for all appliances with a particular health policy, you can blacklist the policy. If you need to disable the results of a group of appliances' health monitoring, you can blacklist the group of appliances. After the blacklist settings take effect, the appliances report a disabled status in the Appliance Status Summary.

Note that if your Defense Center is in a high availability configuration, you can blacklist a managed device on one high availability peer and not the other. You can also blacklist the high availability peer to cause it to mark events generated by it and the devices from which it receives health events as disabled. Defense Centers in a high availability pair have the option to completely or partially blacklist their peer.

To blacklist an entire health policy or group of appliances:

ACCESS: Admin/Maint

1. Select **Health > Blacklist**.
The Blacklist page appears.

- Use the drop-down list on the right to sort the list by group, policy, or model. (Groups on a Defense Center are managed devices.)

Note that appliances with some, but not all, health modules blacklisted will appear as **(Partially Blacklisted)**. If you edit their blacklist status on the main blacklist page, you can either blacklist all modules on those appliances or remove all blacklisting. For information on blacklisting individual health modules on an appliance, see [Blacklisting a Health Policy Module](#) on page 2158.

▼ Ungrouped (4 total)			
<input type="checkbox"/>	katsura 10.10.10.3 - Defense Center 3500 v5.1.0	Health Policy katsura health policy	System Policy katsura system policy
<input type="checkbox"/>	linden 10.10.10.4 - 3D8250 v5.0.2.1	Health Policy linden health policy	System Policy katsura system policy
<input type="checkbox"/>	tamarix 10.10.10.5 - 3D8140 v5.1.0	Health Policy katsura health policy	System Policy katsura system policy
<input type="checkbox"/>	xiramat 10.10.10.6 - 3D8140 v5.1.0	Health Policy katsura health policy	System Policy katsura system policy

TIP! The status icon next to the Health Policy column () indicates the current health status for the appliance. The status icon next to the System Policy column () indicates the communication status between the Defense Center and the device.

- You have two options:
 - To blacklist all appliances in a group, model, or policy category, select the category, then click **Blacklist Selected Devices**.
 - To clear blacklisting from all appliances in a group, model, or policy category, select the category, then click **Clear Blacklist on Selected Devices**.

The page refreshes, now indicating the new blacklist state of the appliances.

Blacklisting an Appliance

LICENSE: Any

If you need to set the events and health status for an individual appliance to disabled, you can blacklist the appliance. After the blacklist settings take effect, the appliance shows as disabled in the Health Monitor Appliance Module Summary and health events for the appliance have a status of disabled.

To blacklist an individual appliance:

ACCESS: Admin/Maint

1. Select **Health > Blacklist**.

The Blacklist page appears.

▼ Ungrouped (4 total)			
<input type="checkbox"/>	katsura 10.10.10.3 - Defense Center 3500 v5.1.0	Health Policy katsura health policy !	System Policy katsura system policy !
<input type="checkbox"/>	linden 10.10.10.4 - 3D8250 v5.0.2.1	Health Policy linden health policy ✓	System Policy katsura system policy ✓
<input type="checkbox"/>	tamarix 10.10.10.5 - 3D8140 v5.1.0	Health Policy katsura health policy !	System Policy katsura system policy ✓
<input type="checkbox"/>	xiramat 10.10.10.6 - 3D8140 v5.1.0	Health Policy katsura health policy !	System Policy katsura system policy ✓

Buttons:

2. Use the drop-down list on the right to sort the list by appliance group, model, or by policy.

3. You have two options:

- To blacklist all appliances in a group, model, or policy category, select the category, then click **Blacklist Selected Devices**.
- To clear blacklisting from all appliances in a group, model, or policy category, select the category, then click **Clear Blacklist on Selected Devices**.

The page refreshes and indicates the new blacklist state of the appliances. Click **Edit** and see [Blacklisting a Health Policy Module](#) on page 2158 to blacklist individual health policy modules.

Blacklisting a Health Policy Module

LICENSE: Any

You can blacklist individual health policy modules on appliances. You may want to do this to prevent events from the module from changing the status for the appliance to warning or critical.

When any part of a module is blacklisted, the line for that module appears in boldface type in the Defense Center web interface.

TIP! After the blacklist settings take effect, the appliance shows as **Partially Blacklisted** or **All Modules Blacklisted** on the Blacklist page and in the Appliance Health Monitor Module Status Summary, but only in expanded views on the main Appliance Status Summary page. Make sure that you keep track of individually blacklisted modules so you can reactivate them when you need them. You may miss necessary warning or critical messages if you accidentally leave a module disabled.

To blacklist an individual health policy module:

ACCESS: Admin/Maint

1. Select **Health > Blacklist**.
The Blacklist page appears.
2. Sort by Group, Policy, or Model, then click **Edit** to display the list of health policy modules for an appliance.
The health policy modules appear.

▼ Ungrouped (4 total)			
<input type="checkbox"/>	katsura 10.10.10.3 - Defense Center 3500 v5.1.0	Health Policy katsura health policy	System Policy katsura system policy
<input type="checkbox"/>	linden 10.10.10.4 - 3D8250 v5.0.2.1	Health Policy linden health policy	System Policy katsura system policy
<input type="checkbox"/>	tamarix 10.10.10.5 - 3D8140 v5.1.0	Health Policy katsura health policy	System Policy katsura system policy
<input type="checkbox"/>	xirammat 10.10.10.6 - 3D8140 v5.1.0	Health Policy katsura health policy	System Policy katsura system policy

3. Select each module that you want to blacklist.
4. Click **Save**.

Configuring Health Monitor Alerts

LICENSE: Any

You can set up alerts to notify you through email, through SNMP, or through the system log when the status changes for the modules in a health policy. You can associate an existing alert response with health event levels to trigger and alert when health events of a particular level occur.

For example, if you are concerned that your appliances may run out of hard disk space, you can automatically send an email to a system administrator when the remaining disk space reaches the warning level. If the hard drive continues to fill, you can send a second email when the hard drive reaches the critical level.

For more information, see the following topics:

- [Creating Health Monitor Alerts](#) on page 2159
- [Interpreting Health Monitor Alerts](#) on page 2161
- [Editing Health Monitor Alerts](#) on page 2162
- [Deleting Health Monitor Alerts](#) on page 2163

Creating Health Monitor Alerts

LICENSE: Any

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert response. You can use an existing

alert or configure a new one specifically to report on system health. When the severity level occurs for the selected module, the alert triggers.

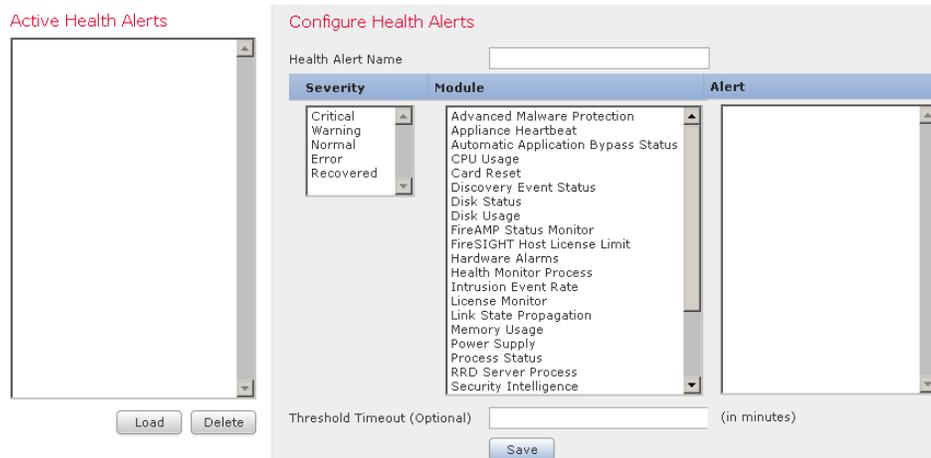
Note that if you create or update a threshold in a way that duplicates an existing threshold, you are notified of the conflict. When duplicate thresholds exist, the health monitor uses the threshold that generates the fewest alerts and ignores the others. The timeout value for the threshold must be between 5 and 4,294,967,295 minutes.

To create health monitor alerts:

ACCESS: Admin

1. Select **Health > Health Monitor Alerts**.

The Health Monitor Alerts page appears.



2. Type a name for the health alert in the **Health Alert Name** field.
3. From the **Severity** list, select the severity level you want to use to trigger the alert.
4. From the **Module** list, select the modules for which you want the alert to apply.

TIP! To select multiple modules, press Shift + Ctrl and click the module names.

5. From the **Alert** list, select the alert response that you want to trigger when the selected severity level is reached.

TIP! Click **Alerts** to open the Alerts page. For more information on creating alerts, see [Working with Alert Responses](#) on page 526.

6. Optionally, in the **Threshold Timeout** field, type the number of minutes that should elapse before each threshold period ends and the threshold count resets. The default value is 5 minutes.

Note that even if the policy run time interval value is less than the threshold timeout value, the interval between two reported health events from a given module is always greater, such that if the threshold timeout is 8 minutes and the policy run time interval is 5 minutes, there will be a 10-minute interval (5 x 2) between reported events.

7. Click **Save** to save the health alert.

A message appears, indicating if the alert configuration was successfully saved. The Active Health Alerts list now includes the alert you created.

Interpreting Health Monitor Alerts

LICENSE: Any

The alerts generated by the health monitor contain the following information:

- Severity, which indicates the severity level of the alert.
- Module, which specifies the health module whose test results triggered the alert.
- Description, which includes the health test results that triggered the alert.

For more information on health alert severity levels, see the [Alert Severities](#) table.

Alert Severities

SEVERITY	DESCRIPTION
Critical	The health test results met the criteria to trigger a Critical alert status.
Warning	The health test results met the criteria to trigger a Warning alert status.
Normal	The health test results met the criteria to trigger a Normal alert status.
Error	The health test did not run.
Recovered	The health test results met the criteria to return to a normal alert status, following a Critical or Warning alert status.

For more information on health modules, see [Understanding Health Modules](#) on page 2112.

Editing Health Monitor Alerts

LICENSE: Any

You can edit existing health monitor alerts to change the severity level, health module, or alert response associated with the health monitor alert.

To edit health monitor alerts:

ACCESS: Admin

1. Select **Health > Health Monitor Alerts**.
The Health Monitor Alerts page appears.
2. Select the alert you want to modify in the **Active Health Alerts** list.

The screenshot displays the 'Configure Health Alerts' interface. On the left, there is an 'Active Health Alerts' list which is currently empty, with 'Load' and 'Delete' buttons below it. The main configuration area includes a 'Health Alert Name' input field, a table with three columns: 'Severity', 'Module', and 'Alert'. The 'Severity' column has a dropdown menu with options: Critical, Warning, Normal, Error, and Recovered. The 'Module' column contains a list of system components such as 'Advanced Malware Protection', 'Appliance Heartbeat', 'Automatic Application Bypass Status', 'CPU Usage', 'Card Reset', 'Discovery Event Status', 'Disk Status', 'Disk Usage', 'FireAMP Status Monitor', 'FireSIGHT Host License Limit', 'Hardware Alarms', 'Health Monitor Process', 'Intrusion Event Rate', 'License Monitor', 'Link State Propagation', 'Memory Usage', 'Power Supply', 'Process Status', 'RRD Server Process', and 'Security Intelligence'. The 'Alert' column has a dropdown menu with 'linden (Syslog)' selected. Below the table is a 'Threshold Timeout (Optional)' input field with the unit '(in minutes)' and a 'Save' button.

3. Click **Load** to load the configured settings for the selected alert.
4. Modify settings as needed. For more information, see [Creating Health Monitor Alerts](#) on page 2159.
5. Click **Save** to save the modified health alert.
A message appears, indicating if the alert configuration was successfully saved.

Deleting Health Monitor Alerts

LICENSE: Any

You can delete existing health monitor alerts.

IMPORTANT! Deleting a health monitor alert does not delete the associated alert response. You must disable or delete the underlying alert response to ensure that alerting does not continue. For more information, see [Enabling and Disabling Alert Responses](#) on page 534 and [Deleting an Alert Response](#) on page 534.

To delete health monitor alerts:

ACCESS: Admin

1. Select **Health > Health Monitor Alerts**.

The Health Monitor Alerts page appears.

2. Select the alert you want to delete in the **Active Health Alerts** list.

3. Click **Delete**.

A message appears, indicating if the alert configuration was successfully deleted.

Using the Health Monitor

LICENSE: Any

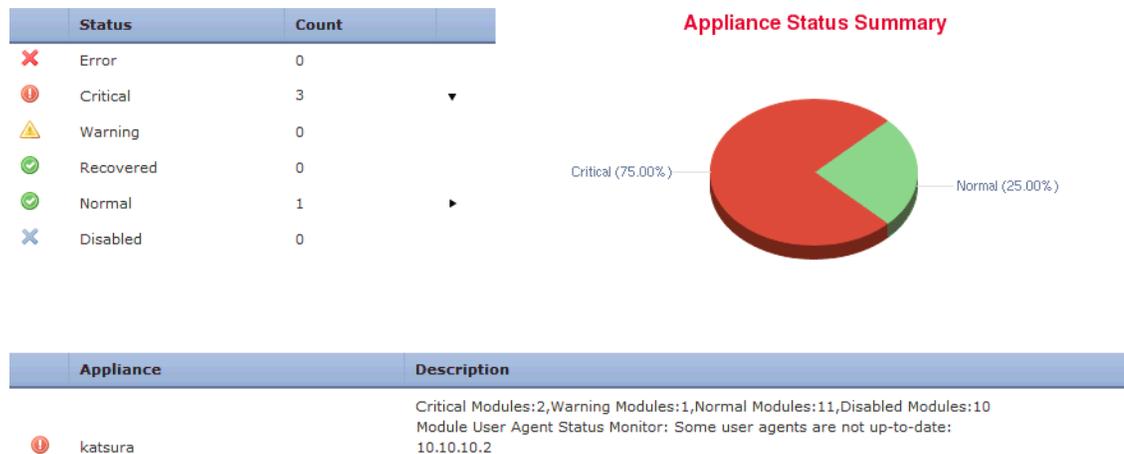
The Health Monitor page provides the compiled health status for all devices managed by the Defense Center, plus the Defense Center. The Status table provides a count of the managed appliances for this Defense Center by overall health status. The pie chart supplies another view of the health status breakdown, indicating the percentage of appliances currently in each health status category.

To use the health monitor:

ACCESS: Admin/Maint/Any Security Analyst

1. Click **Health > Health Monitor**.

The Health Monitor page appears.



2. Select the appropriate status in the Status column of the table or the appropriate portion of the pie chart to the list appliances with that status.

TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

The following topics provide details on the tasks you can perform from the Health Monitor page:

- [Interpreting Health Monitor Status](#) on page 2165
- [Using Appliance Health Monitors](#) on page 2166
- [Configuring Health Policies](#) on page 2116
- [Configuring Health Monitor Alerts](#) on page 2159

Interpreting Health Monitor Status

LICENSE: Any

Available status categories, by severity, include Error, Critical, Warning, Normal, Recovered, and Disabled, as described in the [Health Status Indicator](#) table.

Health Status Indicator

STATUS LEVEL	STATUS ICON	STATUS COLOR	DESCRIPTION
Error		White	Indicates that at least one health monitoring module has failed on the appliance and has not been successfully re-run since the failure occurred. Contact your technical support representative to obtain an update to the health monitoring module.
Critical		Red	Indicates that the critical limits have been exceeded for at least one health module on the appliance and the problem has not been corrected.
Warning		Yellow	Indicates that warning limits have been exceeded for at least one health module on the appliance and the problem has not been corrected.
Normal		Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance.
Recovered		Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance, including modules that were in a Critical or Warning state.
Disabled		Blue	Indicates that an appliance is disabled or blacklisted, that the appliance does not have a health policy applied to it, or that the appliance is currently unreachable.

Using Appliance Health Monitors

LICENSE: Any

The Appliance health monitor provides a detailed view of the health status of an appliance.

IMPORTANT! Your session normally logs you out after 1 hour of inactivity (or another configured interval). If you plan to passively monitor the health monitor for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings. For more information, see [Managing User Login Settings](#) on page 1903 and [Configuring User Interface Settings](#) on page 1994.

To view the status summary for a specific appliance:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Health > Health Monitor**.
The Health Monitor page appears.
2. To show the list of appliances with a particular status, click the arrow in that status row.

TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the **Appliance** column of the appliance list, click the name of the appliance for which you want to view details in the health monitor toolbar.

The Health Monitor Appliance page appears.

Health Monitor

Appliance	
❗	katsura Generate Troubleshooting Files

Module Status Summary

Disabled (75.00%)
Normal (20.00%)
Critical (5.00%)

Alert Detail (katsura)

Alert	Time	Description	▼ Display	Run All Modules
❗ Power Supply	2011-09-27 15:00:13	Power Supply 1 is No Power Power Supply 2 is Online	▶	Run Events
✅ Health Monitor Process	2011-09-27 14:59:13	1 are running Health Monitoring as scheduled	▶	Run Events

4. Optionally, in the **Module Status Summary** graph, click the color for the event status category you want to view. The Alert Detail list toggles the display to show or hide events.

For more information, see the following sections:

- [Understanding Health Modules](#) on page 2112
- [Interpreting Health Monitor Status](#) on page 2165
- [Viewing Alerts by Status](#) on page 2167
- [Running All Modules for an Appliance](#) on page 2167
- [Running a Specific Health Module](#) on page 2168
- [Generating Health Module Alert Graphs](#) on page 2170
- [Using the Health Monitor to Troubleshoot](#) on page 2171

Viewing Alerts by Status

LICENSE: Any

You can show or hide categories of alerts by status.

To show alerts by status:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ Click the status icon or the color segment in the pie chart that corresponds to the health status of the alerts you want to view. The alerts for that category appear in the Alert Detail list.

To hide alerts by status:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ Click the status icon or the color segment in the pie chart that corresponds to the health status of the alerts you want to view. The alerts in the Alert Detail list for that category disappear.

Running All Modules for an Appliance

LICENSE: Any

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run all health module tests on demand to collect up-to-date health information for the appliance.

To run all health modules for the appliance:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Health > Health Monitor**.
The Health Monitor page appears.

- To expand the appliance list to show appliances with a particular status, click the arrow in that status row.

TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

- In the **Appliance** column of the appliance list, click the name of the appliance for which you want to view details.

The Health Monitor Appliance page appears.

Health Monitor

Appliance	
❗	katsura Generate Troubleshooting Files

Alert Detail (katsura)

Alert	Time	Description	Display	Run All Modules		
❗	2011-09-27 15:00:13	Power Supply 1 is No Power Power Supply 2 is Online	▶	Run	Events	
✅	2011-09-27 14:59:13	1 are running Health Monitoring as scheduled	▶	Run	Events	Graph

Module Status Summary

Status	Percentage
Disabled	75.00%
Normal	20.00%
Critical	5.00%

- Click **Run All Modules**.

The status bar indicates the progress of the tests, then the Health Monitor Appliance page refreshes.

IMPORTANT! When you manually run health modules, the first refresh that automatically occurs may not reflect the data from the manually run tests. If the value has not changed for a module that you just ran manually, wait a few seconds, then refresh the page by clicking the device name. You can also wait for the page to refresh again automatically.

Running a Specific Health Module

LICENSE: Any

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run a health module test on demand to collect up-to-date health information for that module.

To run a specific health module:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Health > Health Monitor**.
The Health Monitor page appears.
2. To expand the appliance list to show appliances with a particular status, click the arrow in that status row.

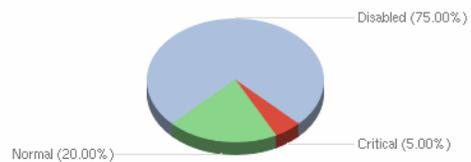
TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the **Appliance** column of the appliance list, click the name of the appliance for which you want to view details.
The Health Monitor Appliance page appears.
4. In the **Module Status Summary** graph of the Health Monitor Appliance page, click the color for the health alert status category you want to view.
The Alert Detail list expands to list the health alerts for the selected appliance for that status category.

Health Monitor

Appliance	
 katsura	<input type="button" value="Generate Troubleshooting Files"/>

Module Status Summary



Alert Detail (katsura)

Alert	Time	Description	Display	Run	Events	Graph
 Power Supply	2011-09-27 15:00:13	Power Supply 1 is No Power Power Supply 2 is Online	▶	Run	Events	
 Health Monitor Process	2011-09-27 14:59:13	1 are running Health Monitoring as scheduled	▶	Run	Events	Graph

5. In the **Alert Detail** row for the alert for which you want to view a list of events, click **Run**.
The status bar indicates the progress of the test, then the Health Monitor Appliance page refreshes.

IMPORTANT! When you manually run health modules, the first refresh that automatically occurs may not reflect the data from the manually run tests. If the value has not changed for a module that you just manually ran, wait a few seconds, then refresh the page by clicking the device name. You can also wait for the page to refresh automatically again.

Generating Health Module Alert Graphs

LICENSE: Any

You can graph the results over a period of time of a particular health test for a specific appliance.

To generate a health module alert graph:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Health > Health Monitor**.

The Health Monitor page appears.

2. To expand the appliance list to show appliances with a particular status, click the arrow in that status row.

TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the **Appliance** column of the appliance list, click the name of the appliance for which you want to view details.

The Health Monitor Appliance page appears.

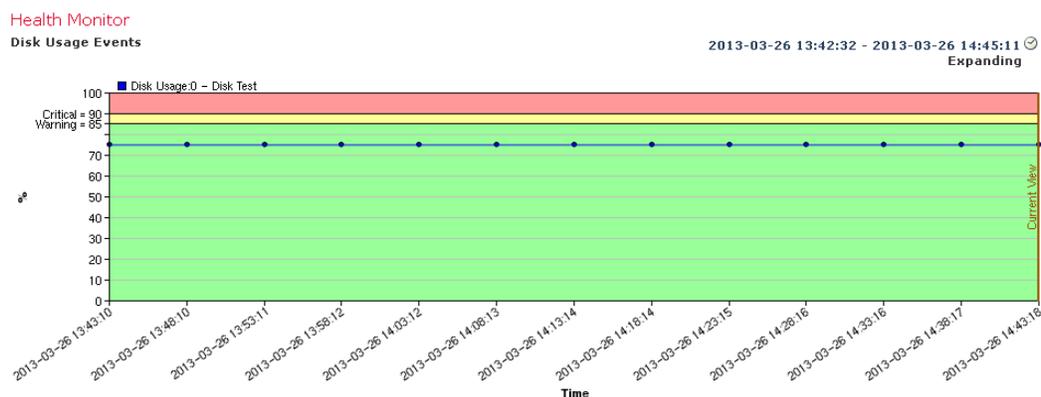
4. In the **Module Status Summary** graph of the Health Monitor Appliance page, click the color for the health alert status category you want to view.

The Alert Detail list expands to list the health alerts for the selected appliance for that status category.

- In the **Alert Detail** row for the alert for which you want to view a list of events, click **Graph**.

A graph appears, showing the status of the event over time. The Alert Detail section below the graph lists all health alerts for the selected appliance.

TIP! If no events appear, you may need to adjust the time range. See [Setting Event Time Constraints](#) on page 1820 for more information.



Using the Health Monitor to Troubleshoot

LICENSE: Any

In some cases, if you have a problem with your appliance, Sourcefire 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 health monitor reports.

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

Selectable Troubleshoot Options (Continued)

THIS OPTION...	REPORTS...
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](#) on page 2172
- [Downloading Troubleshooting Files](#) on page 2173

Generating Appliance Troubleshooting Files

LICENSE: Any

Use the following procedure to generate customized troubleshooting files that you can send to Sourcefire Support.

To generate troubleshooting files:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Health > Health Monitor**.
The Health Monitor page appears.
2. To expand the appliance list to show appliances with a particular status, click the arrow in that status row.

TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the Appliance column of the appliance list, click the name of the appliance for which you want to view details.

The Health Monitor Appliance page appears.

4. Click **Generate Troubleshooting Files**.

The Troubleshooting Options pop-up window appears.

Troubleshooting Options

Please select the data to include:

- All Data
 - Snort Performance and Configuration
 - Hardware Performance and Logs
 - System Configuration, Policy, and Logs
 - Detection Configuration, Policy, and Logs
 - Interface and Network Related Data
 - Discovery, Awareness, VDB Data, and Logs
 - Upgrade Data and Logs
 - All Database Data
 - All Log Data
 - Network Map Information

Note: This may take several minutes.

Generate Cancel

5. 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](#) on page 2171.

6. Click **OK**.

The Defense Center generates the troubleshooting files. You can monitor the file generation process in the task queue (**System > Monitoring > Task Status**).

7. 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:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **System > Monitoring > Task Status**.

The Task Status page appears.

2. Find the task that corresponds to the troubleshooting files you generated.



3. After the appliance generates the troubleshooting files and the task status changes to **Completed**, click **Click to retrieve generated files**.
4. Follow your browser's prompts to download the files.
The files are downloaded in a single `.tar.gz` file.
5. Follow the directions from Support to send the troubleshooting files to Sourcefire.

Working with Health Events

LICENSE: Any

The Defense Center provides fully customizable event views that allow you to quickly and easily analyze the health status events gathered by the health monitor. These event views allow you to search and view event data and to easily access other information that may be related to the events you are investigating.

Many functions that you can perform on the health event view pages are constant across all event view pages. See [Understanding Health Event Views](#) on page 2174 for more information about these common procedures.

From the **Health > Health Events** menu option, you can view health events and can search for specific events.

See the following sections for more information about viewing events:

- [Understanding Health Event Views](#) on page 2174 describes the types of events that FireSIGHT generates.
- [Viewing Health Events](#) on page 2175 describes how to access and use the Event View page.
- [Searching for Health Events](#) on page 2184 describes how to search for specific events using the Event Search page.

Understanding Health Event Views

LICENSE: Any

The Defense Center health monitor logs health events, which you can see on the Health Event View page. If you understand what conditions each health module tests for, you can more effectively configure alerting for health events. For more information on the different types of health modules that generate health events, see [Understanding Health Modules](#) on page 2112.

For more information about viewing and searching for health events, see the following sections:

- [Viewing Health Events](#) on page 2175
- [Understanding the Health Events Table](#) on page 2183
- [Searching for Health Events](#) on page 2184

Viewing Health Events

LICENSE: Any

You can view the appliance health data collected by your health monitor in several ways.

For more information, see the following topics:

- [Viewing All Health Events](#) on page 2175
- [Viewing Health Events by Module and Appliance](#) on page 2176
- [Working with the Health Events Table View](#) on page 2178
- [Interpreting Hardware Alert Details for 3D9900 Devices](#) on page 2179
- [Interpreting Hardware Alert Details for Series 3 Devices](#) on page 2181

Viewing All Health Events

LICENSE: Any

The Table View of Health Events page provides a list of all health events on the selected appliance. For a description of the health modules that generated the events that you may see on this page, see [Understanding Health Modules](#) on page 2112.

When you access health events from the Health Monitor page on your Defense Center, you retrieve all health events for all managed appliances.

To view all health events on all managed appliances:

ACCESS: Admin/Maint/Any Security Analyst

- ▶ Select **Health > Health Events**.

The Events page appears, containing all health events.

Health Events
Health Events ▶ **Table View of Health Events** 2011-09-27 13:40:21 - 2011-09-27 14:46:09 Expanding

No Search Constraints ([Edit Search](#))

<input type="checkbox"/>	Module Name ×	Test Name ×	Time ×	Description ×	Value ×	Units ×	Status ×	Device ×
↓ <input type="checkbox"/>	Power Supply	Power Supply	2011-09-27 14:45:00	Power Supply 1 is No Power Power Supply 2 is Online	0			katsura
↓ <input type="checkbox"/>	License Monitor	License Monitor	2011-09-27 14:42:00	Licenses are up to date	0	Licenses		katsura
↓ <input type="checkbox"/>	URL Filtering Monitor	URL Filtering Monitor	2011-09-27 14:42:00	Process is running correctly	0			katsura

Displaying rows 1-25 of 41 rows << Page 1 of 2 >>

IMPORTANT! If no events appear, you may need to adjust the time range. See [Setting Event Time Constraints](#) on page 1820 for more information.

TIP! You can bookmark this view to allow you to return to the page in the health events workflow containing the Health Events table of events. The bookmarked view retrieves events within the time range you are currently viewing, but you can then modify the time range to update the table with more recent information if needed. For more information, see [Setting Event Time Constraints](#) on page 1820.

Viewing Health Events by Module and Appliance

LICENSE: Any

You can query for events generated by a specific health module on a specific appliance.

To view the health events for a specific module:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Health > Health Monitor**.

The Health Monitor page appears.

2. To expand the appliance list to show appliances with a particular status, click the arrow in that status row.

TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the Appliance column of the appliance list, click the name of the appliance for which you want to view details.

The Health Monitor Appliance page appears.

4. In the Module Status Summary graph of the Health Monitor Appliance page, click the color for the health alert status category you want to view.

The Alert Detail list expands to list the health alerts for the selected appliance for that status category.

5. In the Alert Detail row for the alert for which you want to view a list of events, click **Events**.

The Health Events page appears, containing query results for a query with the name of the appliance and the name of the selected health alert module as constraints.

Health Events
Health Events ▶ **Table View of Health Events** 2011-10-13 11:17:37 - 2011-10-13 12:18:36

▶ Search Constraints (Edit Search Save Search) Expanded Columns

<input type="checkbox"/>	Test Name	Time	Description	Value	Units	Status	Device	Count
<input type="checkbox"/>	Disk Usage - Disk Test	2011-10-13 12:16:56	/ using 29%: 531M (1.4G Avail) of 2.0G	29	%		katsura.sfeng.sourcefire.com	1
<input type="checkbox"/>	Disk Usage - Disk Test	2011-10-13 12:11:56	/ using 29%: 531M (1.4G Avail) of 2.0G	29	%		katsura.sfeng.sourcefire.com	1
<input type="checkbox"/>	Disk Usage - Disk Test	2011-10-13 12:06:56	/ using 29%: 531M (1.4G Avail) of 2.0G	29	%		katsura.sfeng.sourcefire.com	1
<input type="checkbox"/>	Disk Usage - Disk Test	2011-10-13 11:21:56	/ using 29%: 531M (1.4G Avail) of 2.0G	29	%		katsura.sfeng.sourcefire.com	1

Displaying rows 1-12 of 12 rows << Page 1 of 1 >>

If no events appear, you may need to adjust the time range. See [Setting Event Time Constraints](#) on page 1820 for more information.

6. If you want to view all health events for the selected appliance, expand **Search Constraints** and click the **Module Name** constraint to remove it.

Working with the Health Events Table View

LICENSE: Any

The following table describes each action you can perform from the Event View page.

Health Event View Functions

To...	YOU CAN...
learn more about the contents of the columns that appear in the Health event view	find more information in Understanding the Health Events Table on page 2183.
modify the time and date range for events listed in the Health table view	find more information in Setting Event Time Constraints on page 1820. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
sort the events that appear, change what columns display in the table of events, or constrain the events that appear	find more information in Sorting Drill-Down Workflow Pages on page 1835.
delete health events	select the check box next to the events you want to delete and click Delete . To delete all the events in the current constrained view, click Delete All , then confirm you want to delete all the events.
navigate through event view pages	find more information in Navigating to Other Pages in the Workflow on page 1836.
navigate to other event tables to view associated events	find more information in Navigating Between Workflows on page 1836.
bookmark the current page so that you can quickly return to it	click Bookmark This Page , provide a name for the bookmark and click Save . See Using Bookmarks on page 1838 for more information.
navigate to the bookmark management page	click View Bookmarks from any event view. See Using Bookmarks on page 1838 for more information.
generate a report based on data in the table view	click Report Designer . See Creating a Report Template from an Event View on page 1725 for more information.
select another health events workflow	click (switch workflow) . See Selecting Workflows on page 1809 for more information.

Health Event View Functions (Continued)

To...	YOU CAN...
view the details associated with a single health event	click the down arrow link on the left side of the event.
view event details for multiple health events	select the check box next to the rows that correspond with the events you want to view details for and then click View .
view event details for all events in the view	click View All .
view all events of a particular status	click the status icon in the Status column for an event with that status.

Interpreting Hardware Alert Details for 3D9900 Devices

For 3D9900 device models, hardware alarms generate in response to the events described in the [Conditions Monitored for 3D9900 Devices](#) table. The triggering condition can be found in the message detail for the alert.

Conditions Monitored for 3D9900 Devices

CONDITION MONITORED	CAUSES OF YELLOW OR RED ERROR CONDITIONS
NFE card presence	If NFE hardware is detected that is not valid for the appliance, health status for the Hardware Alarms module changes to red and the message details include a reference to the NFE card presence.
NFE temperature	<ul style="list-style-type: none"> If NFE temperature exceeds 95 degrees Celsius, health status for the Hardware Alarms module changes to yellow and the message details include a reference to the NFE temperature. If NFE temperature exceeds 99 degrees Celsius, health status for the Hardware Alarms module changes to red and the message details include a reference to the NFE temperature.
NFE Platform daemon	If the NFE Platform daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.
NFE Message daemon	If the NFE Message daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.

Conditions Monitored for 3D9900 Devices (Continued)

CONDITION MONITORED	CAUSES OF YELLOW OR RED ERROR CONDITIONS
NFE TCAM daemon	If the NFE TCAM daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.
LBIM presence	If the Load Balancing Interface Module (LBIM) switch assembly is not present or not communicating, health status for the Hardware Alarms module changes to red and the message details include a reference to the LBIM presence.
Scmd daemon	If the Scmd daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.
Ps1s daemon	If the Ps1s daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.
Ftwo daemon	If the Ftwo daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.
Rulesd (host rules) daemon	If the Rulesd daemon goes down, health status for the Hardware Alarms module changes to yellow and the message details include a reference to the daemon.
nfm_ipfragd (host frag) daemon	If the nfm_ipfragd daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon.

Interpreting Hardware Alert Details for Series 3 Devices

For Series 3 devices, hardware alarms generate in response to the events described in the [Conditions Monitored for Series 3 Devices](#) table. The triggering condition appears in the message detail for the alert.

Conditions Monitored for Series 3 Devices

CONDITION MONITORED	CAUSES OF YELLOW OR RED ERROR CONDITIONS
Cluster status	If clustered devices are no longer communicating with each other (due, for example, to a cabling problem), the Hardware Alarms module changes to red.
ftwo daemon status	If the ftwo daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon.
NFE cards detected	Indicates the number of NFE cards detected on the system. If this value does not match the appliance's expected NFE count, the Hardware Alarms module changes to red.
NFE hardware status	If one or more NFE cards are not communicating, the Hardware Alarms module changes to red and the applicable card appears in the message details.
NFE heartbeat	If the system detects no NFE heartbeat, the Hardware Alarms module changes to red and message details include a reference to the relevant card(s).
NFE internal link status	If the link between the NMSB and NFE card(s) goes down, the Hardware Alarms module changes to red and message details include a reference to the relevant ports.
NFE Message daemon	If the NFE Message daemon goes down, health status for the Hardware Alarms module changes to red and the message details include a reference to the daemon (and, if applicable, the NFE card number).

Conditions Monitored for Series 3 Devices (Continued)

CONDITION MONITORED	CAUSES OF YELLOW OR RED ERROR CONDITIONS
NFE temperature	<ul style="list-style-type: none"> • If NFE temperature exceeds 97 degrees Celsius, health status for the Hardware Alarms module changes to yellow and message details include a reference to the NFE temperature (and, if applicable, the NFE card number). • If NFE temperature exceeds 102 degrees Celsius, health status for the Hardware Alarms module changes to red and message details include a reference to the NFE temperature. (and, if applicable, the NFE card number).
NFE temperature status	Indicates the current temperature status of the given NFE card. The Hardware Alarms module indicates green for OK, yellow for Warning, and red for Critical (and, if applicable, the NFE card number).
NFE TCAM daemon	If the NFE TCAM daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon (and, if applicable, the NFE card number).
nfe_ipfragd (host frag) daemon	If the nfm_ipfragd daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon (and, if applicable, the NFE card number).
NFE Platform daemon	If the NFE Platform daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon (and, if applicable, the NFE card number).
NMSB communications	If the Media assembly is not present or not communicating, health status for the Hardware Alarms module changes to red and message details include a reference to the NFE temperature (and, if applicable, the NFE card number).
ps1s daemon status	If the ps1s daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon.
Ru1esd (host rules) daemon	If the Ru1esd daemon goes down, health status for the Hardware Alarms module changes to yellow and message details include a reference to the daemon (and, if applicable, the NFE card number).
scmd daemon status	If the scmd daemon goes down, health status for the Hardware Alarms module changes to red and message details include a reference to the daemon.

Understanding the Health Events Table

LICENSE: Any

You can use the Defense Center's health monitor to determine the status of critical functionality within the Sourcefire 3D System. You create and apply health policies to your appliances, which monitor a variety of aspects, including hardware and software status. The Health Monitor modules you choose to enable in your health policy run various tests to determine appliance health status. When the health status meets criteria that you specify, a health event is generated. For more information on health monitoring, see [Monitoring the System](#) on page 2095.

The fields in the health events table are described in the [Health Event Fields](#) table.

Health Event Fields

FIELD	DESCRIPTION
Test Name	The name of the health module that generated the event. For a list of health modules, see the Health Modules table on page 2112.
Time	The timestamp for the health event.
Description	The description of the health module that generated the event. For example, health events generated when a process was unable to execute are labeled unable to Execute .
Value	The value (number of units) of the result obtained by the health test that generated the event. For example, if the Defense Center generates a health event whenever a device it is monitoring is using 80 percent or more of its CPU resources, the value could be a number from 80 to 100.
Units	The units descriptor for the result. You can use the asterisk (*) to create wildcard searches. For example, if the Defense Center generates a health event when a device it is monitoring is using 80 percent or more of its CPU resources, the units descriptor is a percentage sign (%).
Status	The status (Critical, Yellow, Green, or Disabled) reported for the appliance.
Device	The appliance where the health event was reported.

To display the table view of health events:

ACCESS: Admin/Maint/Any Security Analyst

► Select **Health > Health Events**.

The table view appears. For information on working with health events, see [Working with Health Events](#) on page 2174.

TIP! If you are using a custom workflow that does not include the table view of health events, click **(switch workflow)**. On the Select Workflow page, click **Health Events**.

Searching for Health Events

LICENSE: Any

You can search for specific health events. You may want to create searches customized for your network environment, then save them to reuse later. The following table describes the search criteria you can use.

Health Event Search Criteria

SEARCH FIELD	DESCRIPTION
Module Name	Specify the name of the module which generated the health events you want to view. For example, to view events that measure CPU performance, type CPU . The search should retrieve applicable CPU Usage and CPU temperature events.
Value	Specify the value (number of units) of the result obtained by the health test for the events you want to view. For example, if you specify a value of 15 and type CPU in the Units field, you retrieve events where the appliance CPU was running at 15% utilization at the time the test ran.
Description	Specify the description of the events you want to view. For example, you could enter unable to Execute to view any health events where a process was unable to execute. You can use an asterisk (*) in this field to create wildcard searches.

Health Event Search Criteria (Continued)

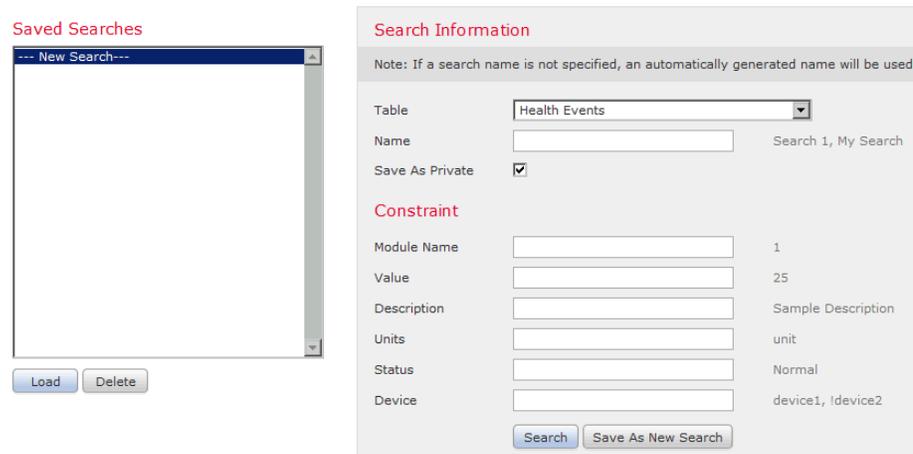
SEARCH FIELD	DESCRIPTION
Units	<p>Specify the units descriptor for the result obtained by the health test for the events you want to view. You can use an asterisk (*) in this field to create wildcard searches.</p> <p>For example, if you type % in the Units field, you retrieve all events for the Disk Usage modules, because the Disk Usage module has a “%” label in the Units field (and no additional text). However, if you type *% in the Units field, you retrieve all events for any modules that contain text followed by a “%” sign in the Units field.</p>
Status	<p>Specify the status for the health events that you want to view. Valid status levels are Critical, Warning, Normal, Error, and Disabled.</p> <p>For example, type <code>critical</code> to retrieve all health events that indicate a critical status.</p>
Device	Specify the name of the device.

For more information on searching, including information on special search syntax as well as saving and loading searches, see [Performing and Saving Searches](#) on page 1770.

To run and save health event searches:

ACCESS: Admin/Maint/Any Security Analyst

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Health Events**.
The Health Event Search page appears.



3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields, as described in the [Health Event Search Criteria](#) table.
If you enter multiple criteria, the search returns only the records that match all the criteria.
5. Optionally, if you want to save the search so that other users can access it, disable the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default health events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click (**switch workflow**). For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**), so that you can run it at a later time.

For more information about searching, see the following sections:

- [Loading a Saved Search](#) on page 1774
- [Deleting a Saved Search](#) on page 1774

CHAPTER 54

AUDITING THE SYSTEM

You can audit activity on your system in two ways. The appliances that are part of the Sourcefire 3D System generate an audit record for each user interaction with the web interface, and also record system status messages in the system log.

The following sections provide more information about the monitoring features that the system provides:

- [Managing Audit Records](#) on page 2187 describes how to view and manage system audit information.
- [Viewing the System Log](#) on page 2200 describes how to view the system log, which contains system status messages.

TIP! Defense Centers and managed devices with Protection licenses also provide full reporting features that allow you to generate reports for almost any type of data accessible in an event view, including auditing data. For more information, see [Working with Reports](#) on page 1724.

Managing Audit Records

LICENSE: Any

Defense Centers and managed devices log read-only auditing information for user activity. Audit logs are presented in a standard event view that allows you to view, sort, and filter audit log messages based on any item in the audit view. You can easily delete and report on audit information and can view detailed reports of the changes that users make.

The audit log stores a maximum of 100,000 entries. When the number of audit log entries exceeds 100,000, the appliance prunes the oldest records from the database to reduce the number to 100,000.

IMPORTANT! If you reboot a Series 3 appliance, then log into the CLI as soon as you are able, any commands you execute are not recorded in the audit log until the web interface is available.

For more information, see the following sections:

- [Viewing Audit Records](#) on page 2188
- [Suppressing Audit Records](#) on page 2191
- [Understanding the Audit Log Table](#) on page 2196
- [Using the Audit Log to Examine Changes](#) on page 2196
- [Searching Audit Records](#) on page 2197

Viewing Audit Records

LICENSE: Any

You can use the appliance to view a table of audit records. Then, you can manipulate the view depending on the information you are looking for. The predefined audit workflow includes a single table view of events. You can also create a custom workflow that displays only the information that matches your specific needs. For information on creating a custom workflow, see [Creating Custom Workflows](#) on page 1841.

The [Audit Log Actions](#) table below describes some of the specific actions you can perform on an audit log workflow page.

Audit Log Actions

To...	YOU CAN...
learn more about the contents of the columns in the table	find more information in Understanding the Audit Log Table on page 2196.
modify the time range used when viewing audit records	find more information at Setting Event Time Constraints on page 1820. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
sort and constrain events on the current workflow page	find more information in Sorting Table View Pages and Changing Their Layout on page 1834.

Audit Log Actions (Continued)

To...	YOU CAN...
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1836.
navigate between pages in the current workflow, keeping the current constraints	click the appropriate page link at the top left of the workflow page. For more information, see Using Workflow Pages on page 1812.
drill down to the next page in the workflow	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To drill down to the next workflow page constraining on a specific value, click a value within a row. Note that this only works on drill-down pages. Clicking a value within a row in a table view constrains the table view and does not drill down to the next page. • To drill down to the next workflow page constraining on some events, select the check boxes next to the events you want to view on the next workflow page, then click View. • To drill down to the next workflow page keeping the current constraints, click View All. <p>TIP! Table views always include “Table View” in the page name. For more information, see Constraining Events on page 1830.</p>
constraining on a specific value	<p>Click a value within a row.</p> <p>If you click a value on a drill-down page, you move to the next page and constrain on the value.</p> <p>Note that clicking a value within a row in a table view constrains the table view and does not drill down to the next page.</p> <p>TIP! Table views always include “Table View” in the page name. For more information, see Constraining Events on page 1830.</p>
delete audit records	<p>use one of the following methods:</p> <ul style="list-style-type: none"> • To delete some items, select the check boxes next to events you want to delete, then click Delete. • To delete all items in the current constrained view, click Delete All, then confirm you want to delete all the events.
temporarily use a different workflow	click (switch workflow) . For more information, see Selecting Workflows on page 1809.
bookmark the current page so you can quickly return to it	click Bookmark This Page . For more information, see Using Bookmarks on page 1838.

Audit Log Actions (Continued)

To...	You CAN...
navigate to the bookmark management page	click View Bookmarks . For more information, see Using Bookmarks on page 1838.
generate a report based on the data in the current view	click Report Designer . For more information, see Creating a Report Template from an Event View on page 1725.
view a summary of a change recorded in the audit log	click the compare icon () next to applicable events in the Message column. For more information, see Using the Audit Log to Examine Changes on page 2196.

To view audit records:

ACCESS: Admin

- ▶ Select **System > Monitoring > Audit**.

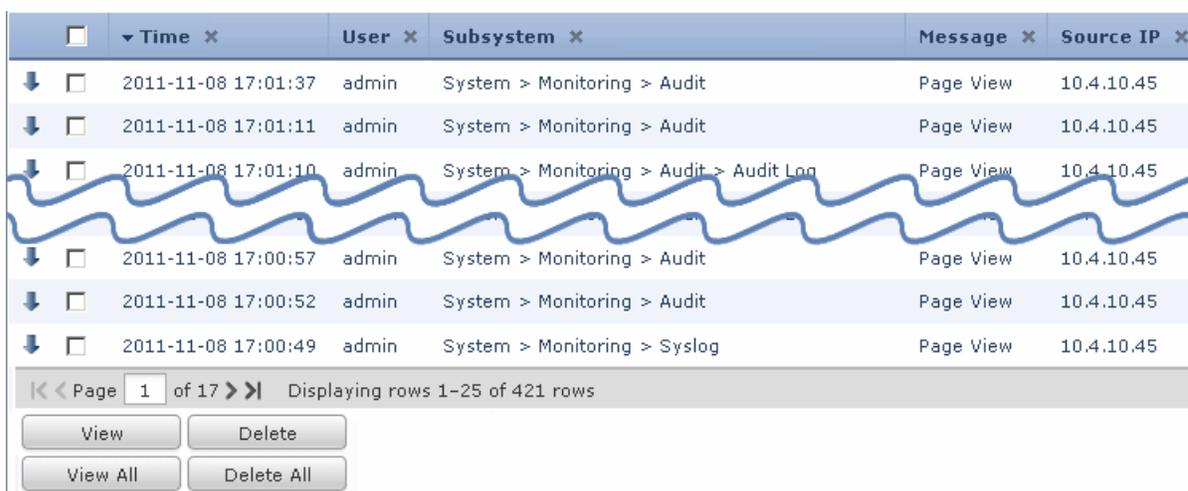
The first (and only) page of the default audit log workflow appears. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218. If no events appear, you may need to adjust the time range. For more information, see [Setting Event Time Constraints](#) on page 1820.

TIP! If you are using a custom workflow that does not include the table view of audit events, click **(switch workflow)**, then select **Audit Log**.

Working with Audit Events

LICENSE: Any

You can change the layout of the event view or constrain the events in the view by a field value.



<input type="checkbox"/>	Time ×	User ×	Subsystem ×	Message ×	Source IP ×
<input type="checkbox"/>	2011-11-08 17:01:37	admin	System > Monitoring > Audit	Page View	10.4.10.45
<input type="checkbox"/>	2011-11-08 17:01:11	admin	System > Monitoring > Audit	Page View	10.4.10.45
<input type="checkbox"/>	2011-11-08 17:01:10	admin	System > Monitoring > Audit > Audit Log	Page View	10.4.10.45
<input type="checkbox"/>	2011-11-08 17:00:57	admin	System > Monitoring > Audit	Page View	10.4.10.45
<input type="checkbox"/>	2011-11-08 17:00:52	admin	System > Monitoring > Audit	Page View	10.4.10.45
<input type="checkbox"/>	2011-11-08 17:00:49	admin	System > Monitoring > Syslog	Page View	10.4.10.45

« < Page 1 of 17 > » Displaying rows 1–25 of 421 rows

View Delete
View All Delete All

When disabling columns, after you click the close icon (×) in the column heading that you want to hide, in the pop-up window that appears, click **Apply**. When you disable a column, it is disabled for the duration of your session (unless you add it back later). Note that when you disable the first column, the Count column is added.

To hide or show other columns, or to add a disabled column back to the view, select or clear the appropriate check boxes before you click **Apply**.

Clicking a value within a row in a table view constrains the table view and does not drill down to the next page.

TIP! Table views always include “Table View” in the page name.

For more information, see the following topics:

- [Constraining Events](#) on page 1830.
- [Using Compound Constraints](#) on page 1832
- [Sorting Drill-Down Workflow Pages](#) on page 1835
- [Understanding the Audit Log Table](#) on page 2196

Suppressing Audit Records

LICENSE: Any

If your auditing policy does not require that you audit specific types of user interactions with the Sourcefire 3D System, you can prevent those interactions

from generating audit records. For example, by default, each time a user views the online help, the Sourcefire 3D System generates an audit record. If you do not need to keep a record of these interactions, you can automatically suppress them.

To configure audit event suppression, you must have access to an appliance's **admin** user account, and you must be able to either access the appliance's console or open a secure shell.

WARNING! Make sure that only authorized personnel have access to the appliance and to its **admin** account.

To suppress audit records, you must create one or more files in the `/etc/sf` directory in the following form:

`AuditBlock.type`

where *type* is `address`, `message`, `subsystem`, or `user`.

IMPORTANT! If you create an `AuditBlock.type` file for a specific type of audit message, but later decide that you no longer want to suppress them, you must delete the contents of the `AuditBlock.type` file but leave the file itself on the Sourcefire 3D System.

The contents for each audit block type must be in a specific format, as described in the [Audit Block Types](#) table. Make sure you use the correct capitalization for the file names. Note also that the contents of the files are case sensitive.

Audit Block Types

TYPE	DESCRIPTION
Address	Create a file named <code>AuditBlock.address</code> and include, one per line, each IP address that you want to suppress from the audit log. You can use partial IP addresses provided that they map from the beginning of the address. For example, the partial address <code>10.1.1</code> matches addresses from <code>10.1.1.0</code> through <code>10.1.1.255</code> .
Message	Create a file named <code>AuditBlock.message</code> and include, one per line, the message substrings that you want to suppress. Note that substrings are matched so that if you include <code>backup</code> in your file, all messages that include the word <code>backup</code> are suppressed.

Audit Block Types (Continued)

TYPE	DESCRIPTION
Subsystem	<p>Create a file named <code>AuditBlock.subsystem</code> and include, one per line, each subsystem that you want to suppress.</p> <p>Note that substrings are not matched. You must use exact strings. See the Subsystem Names table for a list of subsystems that are audited.</p>
User	<p>Create a file named <code>AuditBlock.user</code> and include, one per line, each user account that you want to suppress. You can use partial string matching provided that they map from the beginning of the username. For example, the partial username <code>IPSAlyst</code> matches the user names <code>IPSAlyst1</code> and <code>IPSAlyst2</code>.</p>

Note that when you add an `AuditBlock` file, an audit record with a subsystem of `Audit` and a message of `Audit Filter type Changed` is added to the audit events. For security reasons, this audit record **cannot** be suppressed.

The following table lists audited subsystems.

Subsystem Names

NAME	INCLUDES USER INTERACTIONS WITH...
Admin	Administrative features such as system and access configuration, time synchronization, backup and restore, device management, user account management, and scheduling
Alerting	Alerting functions such as email, SNMP, and syslog alerting
Audit Log	Audit event views
Audit Log Search	Audit event searches
Command Line	Command line interface
Configuration	Email alerting
COOP	Continuity of operations feature
Date	Date and time range for event views
Default Subsystem	Options that do not have assigned subsystems
Detection & Prevention Policy	Menu options for intrusion policies

Subsystem Names (Continued)

NAME	INCLUDES USER INTERACTIONS WITH...
Error	System-level errors
eStreamer	eStreamer configuration
EULA	Reviewing the end user license agreement
Events	Intrusion and discovery event views
Events Clipboard	Intrusion event clipboard
Events Reviewed	Reviewed intrusion events
Events Search	Any event search
Failed to install rule update <i>rule_update_id</i>	Installing rule updates
Header	Initial presentation of the user interface after a user logs in
Health	Health monitoring
Health Events	Health monitoring event views
Help	Online help
High Availability	High availability feature
IDS Impact Flag	Impact flag configuration
IDS Policy	Intrusion policies
IDSPolicy > <i>policy_name</i> > Appliance > <i>det_engine_name</i>	Applying intrusion policies
IDSRule sid: <i>sig_id</i> rev: <i>rev_num</i>	Intrusion rules by SID
Incidents	Intrusion incidents
Insert Policy Apply Job	Applying policies
Install	Installing updates

Subsystem Names (Continued)

NAME	INCLUDES USER INTERACTIONS WITH...
Intrusion Events	Intrusion events
Login	Web interface login and logout functions
Menu	Any menu option
Object export > <i>obj_type</i> > <i>obj_name</i>	Importing objects of a specific type and name
Permission Escalation	User role escalation
Preferences	User preferences, such as the time zone for a user account and individual event preferences
Policy	Any policy, including intrusion policies
Register	Registering devices on a Defense Center
RemoteStorageDevice	Configuring remote storage devices
Reports	Report listing and report designer features
Rules	Intrusion rules, including the rule editor and the rule importation process
Rule Update Import Log	Viewing the rule update import log
Rule Update Install	Installing rule updates
Status	Syslog, as well as host and performance statistics
System	Various system-wide settings
System Policy > <i>policy_name</i> Appliance > <i>appliance_name</i>	Applying system policies
Task Queue	Viewing the task queue
Users	Creating and modifying user accounts and roles

Understanding the Audit Log Table

LICENSE: Any

Each appliance generates an audit event for each user interaction with the web interface. Each event includes a time stamp, the user name of the user whose action generated the event, a source IP, and text describing the event. The fields in the audit log table are described in the [Audit Log Fields](#) table.

Audit Log Fields

FIELD	DESCRIPTION
Time	Time and date that the appliance generated the audit record.
User	User name of the user that triggered the audit event.
Subsystem	Menu path the user followed to generate the audit record. For example, System > Monitoring > Audit is the menu path to view the audit log. In a few cases where a menu path is not relevant, the Subsystem field displays only the event type. For example, Login classifies user login attempts.
Message	Action the user performed. For example, Page View signifies that the user simply viewed the page indicated in the Subsystem, while Save means that the user clicked the Save button on the page. Changes made to the Sourcefire 3D System appear with a compare icon () that you can click to see a summary of the changes. For more information, see Using the Audit Log to Examine Changes on page 2196.
Source IP	IP address associated with the host used by the user.
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows.

Using the Audit Log to Examine Changes

LICENSE: Any

You can use the audit log to view detailed reports of changes to your system. These reports compare the current configuration of your system to its most recent configuration before a particular change.

A compare icon () appears next to audit log events that reflect changes to the system. You can click the compare icon to access the Compare Configurations page and view a detailed report of a change.

The Compare Configurations page displays the differences between the system configuration before changes and the running configuration in a side-by-side format. The audit event type, time of last modification, and name of the user who made the change are displayed in the title bar above each configuration.

Differences between the two configurations are highlighted:

- Blue indicates that the highlighted setting is different in the two configurations, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one configuration but not the other.

To examine a change in the audit log:

ACCESS: Admin

1. Select **System > Monitoring > Audit**.

The first page of the default audit log workflow appears.

If you are using a custom workflow that does not include the table view of audit events, click **(switch workflow)**, then select **Audit Log**.

2. Click the compare icon (🔍) next to an applicable audit log event in the **Message** column.

The Compare Configurations page appears.

System Setting (2011-11-22 18:08:37 by admin from 127.0.0.1)	
System Setting	
Change Reconciliation	
Enable Change Reconciliation	off

System Setting (2011-11-28 18:40:18 by admin from 10.10.0.10)	
System Setting	
Change Reconciliation	
Enable Change Reconciliation	on
Time to Run (Hour)	0
Time to Run (Minute)	0

Note that you can navigate through changes individually by clicking **Previous** or **Next** above the title bar. If the change summary is more than one page long, you can also use the scroll bar on the right to view additional changes.

Searching Audit Records

LICENSE: Any

You can search audit records to find information specific to a user, a specific subsystem, or an audit record message.

You may want to create searches customized for your network environment, then save them to reuse later. The search criteria you can use are described in the [Audit Record Search Criteria](#) table. Note that audit searches are not case

sensitive. For example, searching for Analyst01 or analyst01 yields the same results.

Audit Record Search Criteria

SEARCH FIELD	DESCRIPTION	EXAMPLE
User	Enter the user name of the user who triggered the audit events you want to see. You can use an asterisk (*) as a wildcard character in this field.	jsmith returns all audit records involving the user jsmith.
Subsystem	Enter the full menu path a user would follow to generate the audit records you want to see. You can use an asterisk (*) as a wildcard character in this field.	System > Monitoring > Audit and *Audit both return audit records that involve using the audit log. *Audit* returns all of the above records, plus records that involve searching for audit records.
Message	The action the user performed or the button the user clicked on the page. You can use an asterisk (*) as a wildcard character in this field.	Apply returns audit records where the user applied an intrusion policy. Save Rule returns audit records where the user saved a correlation rule. Page view returns audit records where the user viewed the page.
Time	Specify the date and time the audit record was generated. See Specifying Time Constraints in Searches on page 1775 for the syntax for entering time.	> 2006-01-15 13:30:00 returns all audit records generated after January 15, 2006 at 1:30 PM.
Source IP	Enter the IP address of the host that you want to view audit records for. IMPORTANT! You must type a specific IP address. You cannot use IP ranges when searching audit logs.	172.16.1.37 returns all audit records generated by a user from the 172.16.1.37 IP address.
Configuration Change	Specify whether or not you want to view audit records of configuration changes.	yes returns audit records of configuration changes.

For more information on searching, including how to load and delete saved searches, see [Searching for Events](#) on page 1770.

To search for audit records:

ACCESS: Admin

1. Select **Analysis > Search**.
The Search page appears.
2. From the **Table** drop-down list, select **Audit Log Events**.
The Audit Log search page appears.

Search Information

Note: If a search name is not specified, an automatically generated name will be used.

Table: Audit Log Events

Name: Search 1, My Search

Save As Private:

Constraint:

User: username

Subsystem: subsystem

Message: message

Time: > 2009-07-16 13:00:31, < today at 4:30pm

Source IP: 192.168.1.3, 2001:db8:85a3::1370

Configuration Change: yes, no

Buttons: Search, Save As New Search

TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

3. Optionally, if you want to save the search, enter a name for the search in the **Name** field.
If you do not enter a name, one is created automatically when you save the search.
4. Enter your search criteria in the appropriate fields, as described in the [Audit Record Search Criteria](#) table.
If you enter multiple criteria, the search returns only the records that match all the criteria.
5. If you want to save the search so that other users can access it, clear the **Save As Private** check box. Otherwise, leave the check box selected to save the search as private.

TIP! If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

6. You have the following options:
 - Click **Search** to start the search.
Your search results appear in the default audit log workflow, constrained by the current time range. To use a different workflow, including a custom workflow, click **(switch workflow)**. For information on specifying a different default workflow, see [Configuring Event View Settings](#) on page 2218.
 - Click **Save** if you are modifying an existing search and want to save your changes.
 - Click **Save as New Search** to save the search criteria. The search is saved (and associated with your user account if you selected **Save As Private**) so that you can run it at a later time.

Viewing the System Log

LICENSE: Any

The System Log (syslog) page provides you with system log information for the appliance. The system log displays each message generated by the system. The following items are listed in order:

- the date that the message was generated
- the time that the message was generated
- the host that generated the message
- the message itself

IMPORTANT! System log information is local. For example, you **cannot** use the Defense Center to view system status messages in the system logs on your managed devices.

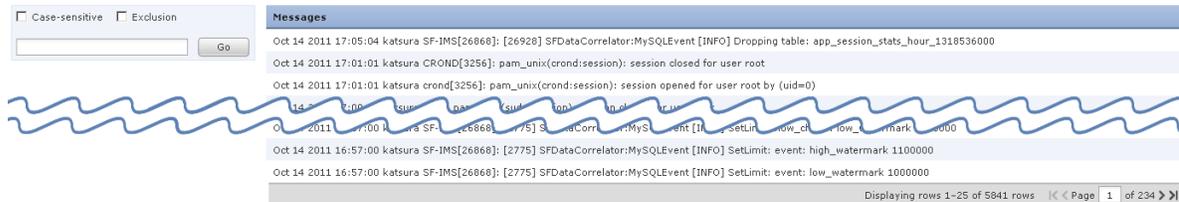
You can view system log messages for specific components by using the filter feature. For more information, see [Filtering System Log Messages](#) on page 2201.

To view the syslog:

ACCESS: Admin/Maint

- ▶ Select **System > Monitoring > Syslog**.

The System Log page appears. The Defense Center version of the page is shown below.



TIP! On the 3D9900, the Load Balancing Interface Module (LBIM) forwards messages to the device's syslog. You can find these messages by filtering on `lbim`.

Filtering System Log Messages

LICENSE: Any

You can view system log messages for specific components by using the filter feature. Filtering allows you to search for specific messages based on content.

The filter functionality uses the UNIX file search utility Grep, and as such, you can use most syntax accepted by Grep. This includes using Grep-compatible regular expressions for pattern matching. You can use a single word as a filter, or you can use Grep-supported regular expressions to search for content.

WARNING! The System Log page does not allow the use of pipe characters for OR expressions. For example, if you use `[word_1|word_2]`, you will receive an invalid filter error.

The [System Log Filter Syntax](#) table shows the regular expression syntax you can use in System Log filters:

System Log Filter Syntax

SYNTAX COMPONENT	DESCRIPTION	EXAMPLE
.	Matches any character or white space	Admi. matches Admin, Admin, Admi1, and Admi&
[[[:alpha:]]]	Matches any alphabetic character	[[[:alpha:]]]dmin matches Admin, badmin, Cadmin, and so on.
[[[:upper:]]]	Matches any uppercase alphabetic character	[[[:upper:]]]dmin matches Admin, Badmin, Cadmin, and so on.
[[[:lower:]]]	Matches any lowercase alphabetic character	[[[:lower:]]]dmin matches admin, badmin, cadmin, and so on.
[[[:digit:]]]	Matches any numeric character	[[[:digit:]]]dmin matches 0admin, 1admin, 2admin, and so on.
[[[:alnum:]]]	Matches any alphanumeric character	[[[:alnum:]]]dmin matches 1admin, admin, 2admin, badmin, and so on.
[[[:space:]]]	Matches any white space, including tabs	Feb[[[:space:]]]29 matches logs from February 29th.
*	Matches one or more instances of the pattern it follows	ab* matches ab, abb, abbb, abbbb, and so on. [ab]* matches ab, abab, ababab, and so on.
?	Matches zero or one instances	ab? matches a or ab.
\	Allows you to search for a character typically interpreted as regular expression syntax	alert\? matches alert?.

The [System Log Filter Examples](#) table shows some example filters you can use on the System Log page.

System Log Filter Examples

TO SEARCH FOR ALL LOG ENTRIES THAT...	USE...
Are generated on November 5	Nov[[[:space:]]]*5

System Log Filter Examples (Continued)

To SEARCH FOR ALL LOG ENTRIES THAT...	USE...
Contain the user name "Admin"	Admin
Contain authorization debugging information on November 5	Nov[[:space:]]*5.*AUTH.*DEBUG

To search for specific message content in the system log:

ACCESS: Admin/Maint

1. Select **System > Monitoring > Syslog**.

The System Log page appears.

2. Enter a word or query in the **Filter** field.

See the [System Log Filter Syntax table](#) on page 2202 and the [System Log Filter Examples table](#) on page 2202 for more information about the filter syntax you can use.

IMPORTANT! Only Grep-compatible search syntax is supported. For example, you could search for all NTP-related system log messages by using `ntp` as a filter, or search for all messages generated in November by using `Nov` as a filter. You could view messages from November 27th by using `Nov[[:space:]]*27` or `Nov.*27`, but you could not, however, use `Nov 27` or `Nov*27` to view these messages.

3. Optionally, to make your search case-sensitive, check **Case-sensitive**. (By default, filters are not case-sensitive.)
4. Optionally, check **Exclusion** to search for all system log messages that do **not** meet the criteria you entered.
5. Click **Go**.
The messages that match the filter appear.

CHAPTER 55

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 Sourcefire 3D System provides a mechanism for archiving data so that the Defense Center or managed device can be restored in case of disaster. Note that backups are valid only for the product version on which you created them, and you cannot create or restore backup files for virtual managed devices.

You can restore a backup onto a replacement appliance if the two appliances are the same model and are running the same version of the Sourcefire 3D System software.

WARNING! Do not use the backup and restore process to copy the configuration files between managed devices. The configuration files include information that uniquely identifies a device and cannot be shared.

By default, system configuration files are saved in the backup file. You can also choose to back up event data.

WARNING! If you applied any Sourcefire 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. Additionally, if you are using a Defense Center, you can use remote storage as detailed in

[Managing Remote Storage](#) on page 2017.

WARNING! Never insert a USB drive into any USB port on a 3D9900. Additionally, remove any device with external storage (for example, a KVM switch with external storage) from a 3D9900 before upgrading or restoring the device.

See the following sections for more information:

- See [Creating Backup Files](#) on page 2205 for information about backing up files from the appliance.
- See [Creating Backup Profiles](#) on page 2208 for information about creating backup profiles that you can use later as templates for creating backups.
- See [Backing up Your Managed Devices with a Defense Center](#) on page 2209 for information about backing up managed devices with the Defense Center.
- See [Uploading Backups from a Local Host](#) on page 2210 for information about uploading backup files from a local host.
- See [Restoring the Appliance from a Backup File](#) on page 2211 for information about how to restore a backup file to the appliance.

Creating Backup Files

LICENSE: Any

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 and packet 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.

WARNING! 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 85% of available disk space. If necessary, delete old backup files, transfer old backup files off the appliance, or use remote storage.

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 because web browsers do not support uploading files that large. On Defense Centers, the backup file can be saved to a remote

location; see [Managing Remote Storage](#) on page 2017 for more information.

IMPORTANT! While your backup task is collecting discovery events, data correlation is temporarily suspended.

If you perform a backup, then delete reviewed intrusion events, your backup restores the deleted intrusion events but does not restore their reviewed status. You view those restored intrusion events under Intrusion Events, not under Reviewed Events. See [Reviewing Intrusion Events](#) on page 611.

WARNING! 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](#) on page 187.

When you back up your managed device from the device itself, you back up the configuration only. Use the Defense Center that manages the physical device to perform a complete backup.

The Defense Center version of the page is shown below.

The screenshot shows a 'Create Backup' form with the following fields and options:

- Name: [Text Input Field]
- Storage Location: /var/sf/backup/
- Back Up Configuration:
- Back Up Events:
- Email: [Not available. You must set up your mail relay host.](#)
- Copy when complete:

Buttons: Start Backup, Save As New, Cancel

The physical managed device version of the page is shown below.

The screenshot shows a 'Create Backup' form with the following fields and options:

- Name: [Text Input Field]
- Storage Location: /var/sf/backup/
- Email when complete:
- Email Address: [Text Input Field]
- Copy when complete:

Buttons: Start Backup, Save As New, Cancel

To create a backup file:

ACCESS: Admin/Maint

1. Select **System > Tools > Backup/Restore**.
The Backup Management page appears.

2. Click **Managed Device Backup** or **Defense Center Backup**.
The Create Backup page appears.
3. In the **Name** field, type a name for the backup file. You can use alphanumeric characters, punctuation, and spaces.
4. On Defense Centers, you have two further options:
 - To archive the configuration, select **Back Up Configuration**.
 - To archive the entire event database, select **Back Up Events**.
5. Optionally, to be notified when the backup is complete, select the **Email when complete** check box and type your email address in the accompanying text box.

IMPORTANT! To receive email notifications, you must configure a relay host as described in [Configuring a Mail Relay Host and Notification Address](#) on page 1983.

6. Optionally, on Defense Centers, 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.

TIP! Sourcefire recommends that you periodically save backups to a remote location so the appliance can be restored in case of system failure.

7. 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. On Defense Centers, you can direct the backup file to a remote location; see [Managing Remote Storage](#) on page 2017.
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](#) on page 2211.
 - 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 **System > Tools > Backup/Restore**, then clicking **Backup Profiles**. See [Creating Backup Profiles](#) on page 2208 for more information.

Creating Backup Profiles

LICENSE: Any

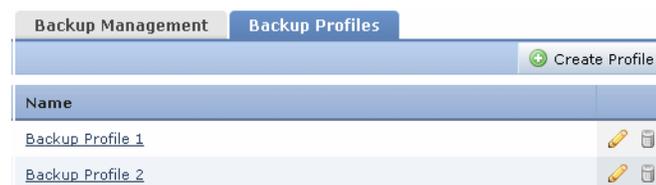
You can use the Backup Profiles page to create backup profiles that contain the settings that you want to use for different types of backups. You can later select one of these profiles when you back up the files on your appliance.

TIP! When you create a backup file as described in [Creating Backup Files](#) on page 2205, a backup profile is automatically created.

To create a backup profile:

ACCESS: Admin/Maint

1. Select **System > Tools > Backup/Restore**.
The Backup Management page appears.
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.

3. Click **Create Profile**.
The Create Backup page appears.
4. Type a name for the backup profile. You can use alphanumeric characters, punctuation, and spaces.
5. Configure the backup profile according to your needs.
See [Creating Backup Files](#) on page 2205 for more information about the options on this page.
6. Click **Save As New** to save the backup profile.
The Backup Profiles page appears and your new profile appears in the list.

Backing up Your Managed Devices with a Defense Center

LICENSE: Any

You can use the Defense Center to back up data on managed devices. The default name for the backup file uses the name of the managed device.

TIP! If you use a backup file name containing spaces or punctuation characters, they change to underscores.

To back up a managed device from a Defense Center:

ACCESS: Admin/Maint

1. Select **System > Tools > Backup/Restore**.
The Backup Management page appears.
2. Click **Managed Device Backup**.
The Managed Device Backup page appears.



The screenshot shows the 'Managed Device Backup' configuration page. At the top, the title 'Managed Device Backup' is displayed in red. Below the title, there is a 'Managed Devices' section with a text input field containing the name 'linden'. Underneath this, there are two checkboxes: 'Include All Unified Files' and 'Retrieve to Defense Center', both of which are currently unchecked. At the bottom of the form, there is a 'Start Backup' button.

3. In the **Managed Devices** field, select the managed devices you want to back up.
4. To include event data in addition to configuration data, select the **Include All Unified Files** check box. Note that unified files are binary files that the Sourcefire 3D System uses to log event data.

5. To save the backup file on the Defense Center, select the **Retrieve to Defense Center** check box. To save each device's backup file on the device itself, leave this check box unselected.

IMPORTANT! If you select **Retrieve to Defense Center** and your Defense Center is configured for remote storage of backups, the device backup file will be saved to the configured remote location, not the Defense Center itself.

6. Click **Start Backup**.

A success message appears and the backup task is created.

It may take several minutes to complete the backup. You can monitor its progress on the Task Status page (**System > Monitoring > Task Status**). When the backup is complete, the backup file appears on the Backup Management page.

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](#) on page 2211, you can upload it to a Defense Center.

TIP! You cannot upload a backup larger than 4GB from your local host because web browsers do not support uploading files that large. As an alternative, copy the backup via SCP to a remote host and retrieve it from there. On Defense Centers, the backup file can be saved to and retrieved from a remote location; see [Managing Remote Storage](#) on page 2017.

To upload a backup from your local host:

ACCESS: Admin/Maint

1. Select **System > Tools > Backup/Restore**.
The Backup Management page appears.
2. Click **Upload Backup**.
The Upload Backup page appears.



The screenshot shows a web interface titled "Upload Backup". Below the title, there is a text prompt "Please Select a File to Upload" followed by a text input field and a "Browse..." button. Below the input field is an "Upload Backup" button.

3. Click **Browse** and navigate to the backup file you want to upload.
After you select the file to upload, click **Upload Backup**.

4. Click **Backup Management** to return to the Backup Management page.
The backup file is uploaded and appears in the backup list. After the Defense Center appliance verifies 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 Sourcefire Rule Update.

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. On Defense Centers, select **Remote Storage** at the top of the Backup Management page to configure remote storage options; then, to enable remote storage, select the **Enable Remote Storage for Backups** check box on the Backup Management page. If you use remote storage, the protocol, backup system, and backup directory are listed at the bottom of the page. The [Backup Management](#) table describes each column and icon on the Backup Management page.

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
Location	The location of the backup file
Size (MB)	The size of the backup file, in megabytes
Events?	“Yes” indicates the backup includes event data
View	Click the name of the backup file to view a list of the files included in the compressed backup file.

Backup Management (Continued)

FUNCTIONALITY	DESCRIPTION
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	On a Defense Center, 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:

ACCESS: Admin

1. Select **System > Tools > Backup/Restore**.

The Backup Management page appears. A Defense Center version of the page is shown.

Defense Center Backups

<input type="checkbox"/>	System Information	Date Created	File Name	Location	Size (MB)	Events?
<input checked="" type="checkbox"/>	katsura Defense Center 3500 v5.0.0	2011-10-26 14:28:29	katsura2-2011-10-26T14-28-05.tgz	Local	31	
<input type="checkbox"/>	katsura Defense Center 3500 v5.0.0	2011-10-26 14:26:55	katsura.example.com2-2011-10-26T14-26-29.tgz	Local	31	Yes

Storage Location: /var/sf/backup/ (Disk Usage: 1%)

- 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. The Defense Center version of the page is truncated to show a sample of the files that are backed up.

katsura2-2011-10-26T14-28-05.tgz

Perm	User	Size	Date	Time	Name
drwx-----	www/www	0	2011-10-26	14:28	var/tmp/backupbj7d/
-rw-rw-rw-	www/www	106	2011-10-26	14:28	var/tmp/backupbj7d/files_not_to_tar
-rw-rw-rw-	www/www	653	2011-10-26	14:28	var/tmp/backupbj7d/files_to_tar
-rw-r--r--	root/root	170416654	2011-10-26	14:28	var/tmp/backupbj7d/db.dump
drwxr-xr-x	root/root	0	2011-10-22	17:09	etc/cron.d/
-rw-r--r--	root/root	570	2011-10-24	09:56	etc/group
-rw-r--r--	root/root	0	2011-10-10	20:40	etc/gshadow
-rw-r--r--	root/root	254	2011-10-24	10:11	etc/hosts
-rw-r-----	root/www	6024	2011-10-10	20:40	etc/httpd/httpsd.conf
-rw-r-----	root/www	103	2011-10-10	20:40	etc/httpd/ssl_certificates.conf
-rw-r--r--	root/root	1	2011-10-24	10:11	etc/issue
-rw-r--r--	root/root	8678	2011-10-10	22:08	etc/ldap.conf

- Click **Backup Management** to return to the Backup Management page.
- 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.

Restore Backup

Backup Name: katsura2-2011-10-26T14-28-05.tgz

Replace Configuration Data:

WARNING! This procedure overwrites all configuration files and, on the managed device, all event data.

- To restore files, select either or both:
 - Replace Configuration Data**
 - Restore Event Data**

IMPORTANT! Note that, when you restore the configuration of a managed device from a backup file, any device configuration changes you made from the device's managing Defense Center will also be restored, even changes you made after you created that backup file.

6. Click **Restore** to begin the restoration.
The appliance is restored using the backup file you specified.
7. Reboot the appliance.
8. Apply the latest Sourcefire Rule Update to reapply rule updates.
9. Reapply any access control, intrusion, network discovery, health, and system policies to the restored system.

CHAPTER 56

SPECIFYING USER PREFERENCES

You can configure the preferences that are tied to a single user account, such as the home page, account password, time zone, dashboard, and event viewing preferences.

Depending on your user role, you can specify certain preferences for your user account, including passwords, event viewing preferences, time zone settings, and home page preferences. See the following sections for more information:

- [Changing Your Password](#) on page 2216 explains how to change the password for your user account.
- [Specifying Your Home Page](#) on page 2217 explains how to use one of the existing pages as your default home page. After setting this value, this becomes the first page you see upon logging into the appliance.
- [Configuring Event View Settings](#) on page 2218 describes how the event preferences affect what you see as you view events.
- [Setting Your Default Time Zone](#) on page 2222 explains how to set the time zone for your user account and describes how that affects the time stamp on the events that you view.
- [Specifying Your Default Dashboard](#) on page 2223 explains how to choose which of the dashboards you want to use as your default dashboard.

Changing Your Password

LICENSE: Any
SUPPORTED DEVICES: Series 2, Series 3
SUPPORTED DEFENSE CENTERS: Any

All user accounts are protected with a password. You can change your password at any time, and depending on the settings for your user account, you may have to change your password periodically; see [Changing an Expired Password](#) on page 2216.

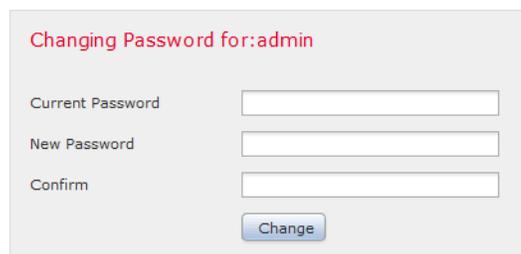
Note that if password strength checking is enabled, passwords must be at least eight alphanumeric characters of mixed case and must include at least one numeric character. Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters.

IMPORTANT! If you are an LDAP or a RADIUS user, you cannot change your password through the web interface.

To change your password:

ACCESS: Any

1. From the drop-down list under your user name, select **User Preferences**. The Change Password page appears.



Changing Password for: admin

Current Password

New Password

Confirm

2. In the **Current Password** field, type your current password and click **Change**.
3. In the **New Password** and **Confirm** fields, type your new password.
4. Click **Change**.

A success message appears on the page when your new password is accepted by the system.

Changing an Expired Password

LICENSE: Any
SUPPORTED DEVICES: Series 2, Series 3
SUPPORTED DEFENSE CENTERS: Any

Depending on the settings for your user account, your password may expire. Note that the password expiration time period is set when your account is created and

cannot be changed. If your password has expired, the Password Expiration Warning page appears.

To respond to the password expiration warning:

ACCESS: Any

► You have two choices:

- Click **Change Password** to change your password now.

If you have zero warning days left, you **must** change your password. Also, if password strength checking is enabled, passwords must be at least eight alphanumeric characters of mixed case and must include at least one numeric character. Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters.

- Click **Skip** to change your password later.

Specifying Your Home Page

LICENSE: Any

You can specify a page within the web interface as your home page for the appliance. The default home page is the Summary Dashboard (**Overview > Dashboards**), except for user accounts with no dashboard access, which use the Welcome page.

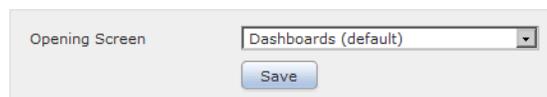
To specify your home page:

ACCESS: Any

1. From the drop-down list under your user name, select **User Preferences**. The Change Password page appears.

2. Click **Home Page**.

The Home Page page appears.



3. Select the page you want to use as your home page from the drop-down list. The options in the drop-down list are based on the access privileges for your user account. For more information, see [User Account Privileges](#) on page 1914.

4. Click **Save**.

Your home page preference is saved.

Configuring Event View Settings

LICENSE: Any

Use the Event View Settings page to configure characteristics of event views in the Sourcefire 3D System.

To configure event preferences:

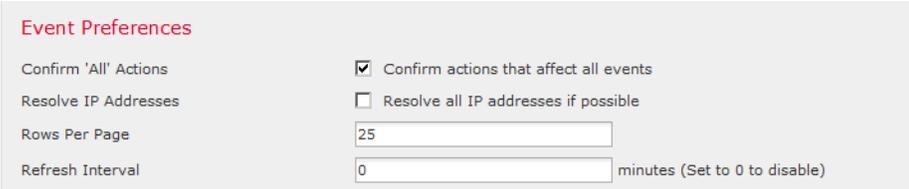
ACCESS: Any except Database

1. From the drop-down list under your user name, select **User Preferences**.
The User Preferences page appears.
2. Click **Event View Settings**.
The Event View Settings page appears.
3. Configure the basic characteristics of event views.
For more information, see [Event Preferences](#) on page 2218.
4. Configure the default time window or windows.
For more information, see [Default Time Windows](#) on page 2219.
5. Configure default workflows.
For more information, see [Default Workflows](#) on page 2221.
6. Click **Save**.
Your changes are implemented.

Event Preferences

LICENSE: Any

Use the Event Preferences section of the Event View Settings page to configure basic characteristics of event views in the Sourcefire 3D System.



The screenshot shows the 'Event Preferences' section of the Sourcefire 3D System interface. It contains four settings:

- Confirm 'All' Actions:** A checkbox labeled 'Confirm actions that affect all events' is checked.
- Resolve IP Addresses:** A checkbox labeled 'Resolve all IP addresses if possible' is unchecked.
- Rows Per Page:** A text input field containing the value '25'.
- Refresh Interval:** A text input field containing the value '0', with the text 'minutes (Set to 0 to disable)' to its right.

The following fields appear in the Event Preferences section:

- The **Confirm "All" Actions** field controls whether the appliance forces you to confirm actions that affect all events in an event view.
For example, if this setting is enabled and you click **Delete All** on an event view, you must confirm that you want to delete all the events that meet the current constraints (including events not displayed on the current page) before the appliance will delete them from the database.

- The **Resolve IP Addresses** field will allow the appliance, whenever possible, to display host names instead of IP addresses in event views.
Note that an event view may be slow to display if it contains a large number of IP addresses and you have enabled this option. Note also that for this setting to take effect, you must have a DNS server configured in the system settings; see [Configuring Network Settings](#) on page 2009.
- The **Rows Per Page** field controls how many rows of events per page you want to appear in drill-down pages and table views.
- The **Refresh Interval** field sets the refresh interval for event views in minutes. Entering “0” disables the refresh option. Note that this interval does not apply to dashboards.

Default Time Windows

LICENSE: Any

The time window, sometimes called the time range, imposes a time constraint on the events in any event view. Use the Default Time Windows section of the Event View Settings page to control the default behavior of the time window.

Default Time Windows

Number of Time Windows: Single Multiple

Events Time Window: Use End Time

Audit Log Time Window: Use End Time

Health Monitoring Time Window: Use End Time

Note that, regardless of the default time window setting, you can always manually change the time window for individual event views during your event analysis. Also, keep in mind that time window settings are valid for only the current session. When you log out and then log back in, time windows are reset to the defaults you configured on this page. For more information, see [Setting Event Time Constraints](#) on page 1820.

There are three types of events for which you can set the default time window:

- The **Events Time Window** sets a single default time window for most events that can be constrained by time.
- The **Audit Log Time Window** sets the default time window for the audit log.
- The **Health Monitoring Time Window** sets the default time window for health events.

You can only set time windows for event types your user account can access. All user types can set event time windows. Administrators, Maintenance Users, and

Security Analysts can set health monitoring time windows. Administrators and Maintenance Users can set audit log time windows.

Note that because not all event views can be constrained by time, time window settings have no effect on event views that display hosts, host attributes, applications, clients, vulnerabilities, user identity, or white list violations.

You can either use **Multiple** time windows, one for each of these types of events, or you can use a **Single** time window that applies to all events. If you use a single time window, the settings for the three types of time window disappear and a new **Global Time Window** setting appears.

The screenshot shows a configuration panel titled "Default Time Windows". It contains the following elements:

- "Number of Time Windows": Radio buttons for "Single" (selected) and "Multiple".
- "Global Time Window": A dropdown menu showing "Show the Last - Static/Expanding".
- A text input field containing the number "1".
- A dropdown menu for units, currently set to "hour(s)".
- A checkbox labeled "Use End Time" which is currently unchecked.

There are three types of time window:

- *static*, which displays all the events generated from a specific start time to a specific end time
- *expanding*, which displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view
- *sliding*, which displays all the events generated from a specific start time (for example, one day ago) to the present; as time moves forward, the time window "slides" so that you see only the events for the range you configured (in this example, for the last day)

The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC).

The following options appear in the **Time Window Settings** drop-down list:

- The **Show the Last - Sliding** option allows you configure a sliding default time window of the length you specify.

The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window "slides" so that you always see events from the last hour.

- The **Show the Last - Static/Expanding** option allows you to configure either a static or expanding default time window of the length you specify.
For **static** time windows, enable the **Use End Time** check box. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.
For **expanding** time windows, disable the **Use End Time** check box. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window expands to the present time.
- The **Current Day - Static/Expanding** option allows you to configure either a static or expanding default time window for the current day. The current day begins at midnight, based on the time zone setting for your current session.
For **static** time windows, enable the **Use End Time** check box. The appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.
For **expanding** time windows, disable the **Use End Time** check box. The appliance displays all the events generated from midnight to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 24 hours before you log out, this time window can be more than 24 hours.
- The **Current Week - Static/Expanding** option allows you to configure either a static or expanding default time window for the current week. The current week begins at midnight on the previous Sunday, based on the time zone setting for your current session.
For **static** time windows, enable the **Use End Time** check box. The appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window.
For **expanding** time windows, disable the **Use End Time** check box. The appliance displays all the events generated from midnight Sunday to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 1 week before you log out, this time window can be more than 1 week.

Default Workflows

LICENSE: Any

A workflow is a series of pages displaying data that analysts use to evaluate events. For each event type, the appliance ships with at least one predefined workflow. For example, as a Security Analyst, depending on the type of analysis

you are performing, you can choose among ten different intrusion event workflows, each of which presents intrusion event data in a different way.

The appliance is configured with a default workflow for each event type. For example, the Events by Priority and Classification workflow is the default for intrusion events. This means whenever you view intrusion events (including reviewed intrusion events), the appliance displays the Events by Priority and Classification workflow.

You can, however, change the default workflow for each event type using the Default Workflows sections of the Event View Settings page.

Keep in mind that the default workflows you are able to configure depend on your user role. For example, intrusion event analysts cannot set default discovery event workflows. For general information on workflows, see [Understanding and Using Workflows](#) on page 1791.

Setting Your Default Time Zone

LICENSE: Any

You can change the time zone used to display events from the standard UTC time that the appliance uses. When you configure a time zone, it applies only to your user account and is in effect until you make further changes to the time zone.

WARNING! The Time Zone function assumes that the default system clock is set to UTC time. If you have changed the system clock on the appliance to use a local time zone, you must change it back to UTC time in order to view accurate local time on the appliance. For more information about time synchronization between the Defense Center and the managed devices, see [Synchronizing Time](#) on page 1991.

To change your time zone:

ACCESS: Any

1. From the drop-down list under your user name, select **User Preferences**.
The Change Password page appears.

2. Click **Time Zone Settings**.

The Time Zone Preference page appears.



3. From the left list box, select the continent or area that contains the time zone you want to use.

For example, if you want to use a time zone standard to North America, South America, or Canada, select **America**.

4. From the right list box, select the zone (city name) that corresponds with the time zone you want to use.

For example, if you want to use Eastern Standard Time, you would select **New York** after selecting **America** in the first time zone box.

5. Click **Save**.

The time zone is set.

Specifying Your Default Dashboard

LICENSE: Any

You can specify one of the dashboards on the appliance as the default dashboard. The default dashboard appears when you select **Overview > Dashboards**. If you do not have a default dashboard defined, the Dashboard List page appears. For general information on dashboards, see [Using Dashboards](#) on page 66.

To specify your default dashboard:

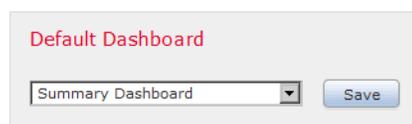
ACCESS: Admin/Maint/Any Security Analyst

1. From the drop-down list under your user name, select **User Preferences**.

The Change Password page appears.

2. Click **Dashboard Settings**.

The Dashboard Settings page appears.



3. Select the dashboard you want to use as your default from the drop-down list. If you select **None**, when you select **Overview > Dashboards**, the Dashboard List page appears. You can then select a dashboard to view.
4. Click **Save**.
Your default dashboard preference is saved.

APPENDIX A

IMPORTING AND EXPORTING OBJECTS

You can use the Import/Export feature to copy several types of objects, including policies, from one appliance to another appliance of the same type. Object import and export is not intended as a backup tool, but can be used to simplify the process of adding new appliances to your Sourcefire 3D System.

You can import and export the following objects:

- access control policies
- alert responses
- application detectors
- custom tables
- custom user roles
- custom workflows
- dashboards
- health policies
- intrusion policies
- report templates
- saved searches
- system policies
- third-party product mappings
- third-party vulnerability mappings

To import an exported object, both appliances must be running the same version of the Sourcefire 3D System. To import an exported intrusion policy (or an access

control policy that incorporates an intrusion policy), the rule update versions on both appliances must also match.

For more information, see the following sections:

- [Exporting Objects](#) on page 2226
- [Importing Objects](#) on page 2231

Exporting Objects

LICENSE: Any

You can export a single object, or you can export several objects (of the same type or of different types) at once. When you later import the package onto another appliance, you can choose which objects in the package to import.

When you export an object, the appliance also exports revision information for that object. The Sourcefire 3D System uses that information to determine whether you can import that object onto another appliance; you cannot import an object revision that already exists on an appliance.

In addition, when you export an object, the appliance also exports system objects that the object depends on, such as authentication objects. For example, if you set up authentication to an LDAP server on your Defense Center, then export a Defense Center system policy with authentication enabled, the authentication object is exported as well.

TIP! Many list pages in the Sourcefire 3D System 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 objects:

- *Alert responses* — An alert response is a set of configurations that allows the Sourcefire 3D System to interact with the external system where you plan to send the alert.
- *Custom tables* — A custom table is a table you can construct that combines fields from two or more of the predefined tables delivered with the Sourcefire 3D System.
- *Custom user roles* — A custom user role is a user role that you create with a specialized set of access privileges. Exporting a custom user role that requires saved searches also exports all of the necessary saved searches.

- *Custom workflows* — A custom workflow is a workflow that you create to meet the unique needs of your organization. On the Defense Center, you can export custom workflows that you create as well as the predefined custom workflows delivered with the appliance.

Note that if a Defense Center does not allow you to view the table on which an exported custom workflow is based, you can import the workflow but will not be able to view it.

- *Dashboards* — A dashboard is a customizable tabbed view that provides you with an at-a-glance display of your current system status. Dashboards use various widgets to present data about the events collected and generated by the Sourcefire 3D System, as well as information about the status and overall health of the appliances in your deployment.

Note that the dashboard widgets that you can view depend on the type of appliance you are using and on your user role. For more information, see [Understanding Widget Availability](#) on page 71.

- *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 as well as any objects the rules use, and may also include referenced intrusion and file policies. 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.

If an access control policy that you export references an intrusion policy, the rule update version on the exporting and importing appliances must match.

If an access control policy that you export references an unsupported DC500 or Series 2 device policy feature or rule condition, you cannot use a DC500 to apply the policy and you cannot apply the policy to a Series 2 device. Neither the DC500 nor Series 2 devices support user or URL rule conditions, Security Intelligence, or file policies that include rules that use the Block Malware or Malware Cloud Lookup action. Additionally, Series 2 devices do not support application rule conditions.

- *Health policies* — A health policy comprises the criteria used when checking the health of appliances in your deployment, that is, whether your Sourcefire hardware and software are working correctly.

- *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 include preprocessors; intrusion rules that inspect the protocol header values, payload content, and certain packet size characteristics; adaptive profile configurations; FireSIGHT recommended rules configurations; and tools that allow you to control how often events are logged and displayed.

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 SMTP 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 Defense Center to another, the imported policy may behave differently if the second Defense Center has differently configured system default variables.

IMPORTANT! You cannot use the Import/Export feature to update rules created by Sourcefire's Vulnerability Research Team (VRT). Instead, download and apply the latest rule update version; see [Importing Rule Updates and Local Rule Files](#) on page 2072.

- *Report templates* — Reports are document files formatted in PDF, HTML, or CSV that collate specific Sourcefire 3D System data. A report template specifies the data searches and formats for the report and its sections. When you export a report template, all saved searches, images, network objects, objects created in the object manager, and custom tables that are necessary for the report are exported also.
- *Saved searches* — A saved search provides access to predefined Sourcefire 3D System data for users with limited permissions. When you export a custom user role that requires saved searches, the necessary saved searches are exported also. You can also export individual user-defined saved searches.

- *System policies* — A system policy controls the aspects of an appliance that are likely to be similar to other Sourcefire 3D System appliances in your deployment, including database event limits, time settings, login banners, and so on.

If external authentication is enabled in the system policy you are exporting, the associated authentication objects are exported as well.

Note that system policies on Defense Centers contain database settings that do not apply to managed devices. If you export a system policy from a managed device and then import it onto a Defense Center, the database limits that you could not configure on the device are set to the default values on the Defense Center.

- *Third-party product mappings* — If you import data from a third-party application, you must map the product to the third-party name to assign vulnerabilities and perform impact correlation using that data. Mapping the product associates Sourcefire vulnerability information with the third-party product name, which allows the Sourcefire 3D System to perform impact correlation using that data. For information on creating a third-party product mapping, see [Mapping Third-Party Products](#) on page 1682.
- *Third-party vulnerability mappings* — To add vulnerability information from a third-party application to the vulnerability database, you must map the third-party identification string for each imported vulnerability to any existing Sourcefire, Bugtraq, or Snort ID. After you create a mapping for the vulnerability, the mapping works for all vulnerabilities imported to hosts in your network map and allows impact correlation for those vulnerabilities. For information on creating a third-party vulnerability mapping, see [Mapping Third-Party Vulnerabilities](#) on page 1687.
- *Application detectors* — When the system analyzes IP traffic, it uses detectors to collect information about and then identify the commonly used applications running on hosts on your network. You can export two kinds of detectors: user-defined detectors and individual add-on detectors provided by Sourcefire Professional Services. For more information on detectors, see [Working with Application Detectors](#) on page 1663.

IMPORTANT! Depending on the number of objects being exported and the number of objects those objects reference, the export process may take several minutes.

To export one or more objects:

ACCESS: Admin

1. Make sure that the appliance where you are exporting the objects and the appliance where you plan to import the objects are running the same version of the Sourcefire 3D System. If you are exporting an intrusion policy (or an access control policy that incorporates an intrusion policy), you must also make sure that the rule update versions match.

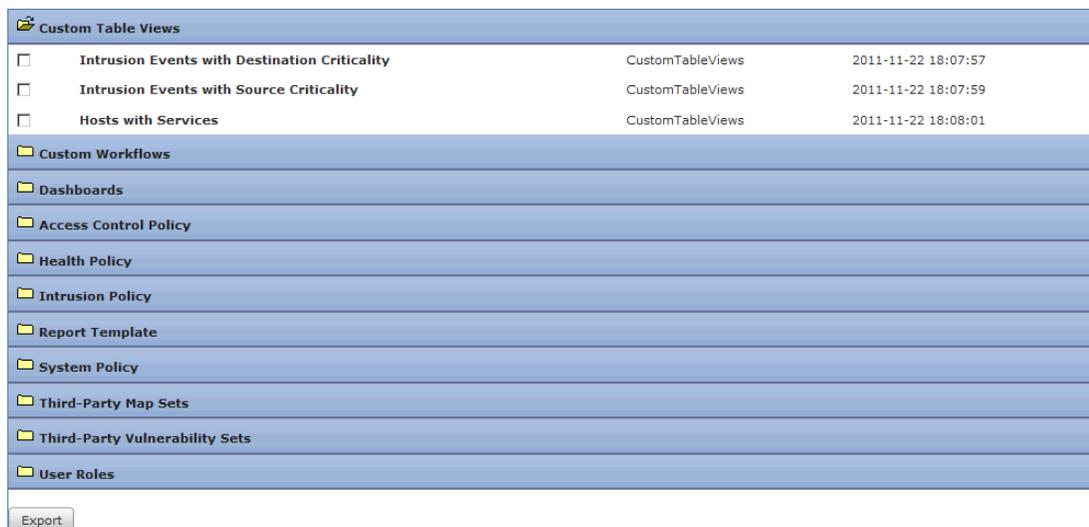
If the versions of the Sourcefire 3D System (and, if applicable, the rule update versions) do not match, the import will fail.

2. Select **Systems > Tools > Import/Export**.

The Import/Export page appears, including a list of the objects on the appliance. Note that object categories with no objects to export do not appear in this list.

TIP! You can click the collapse icon (🔼) next to an object type to collapse the list of objects. Click the expand folder icon (📁) next to an object type to reveal objects.

The Defense Center version of the page is shown below with some object types collapsed.



3. Select the check boxes next to the objects you want to export and click **Export**.
4. Follow your web browser's prompts to save the exported package to your computer.

Importing Objects

LICENSE: Any

After you export an object from an appliance, you can import it onto a different appliance as long as that appliance supports it. Note, however, that some imported objects may not be useful depending on the type of appliance you are using and on your user role.

Depending on the type of object you are importing, you should keep the following points in mind:

- You must make sure that the appliance where you import an object is running the same version of the Sourcefire 3D System as the appliance you used to export the object. If you are importing an intrusion policy (or an access control policy that incorporates an intrusion policy), the rule update versions on both appliances must also match. If the versions do not match, the import will fail.
- When you import a custom user role that requires saved searches, the necessary saved searches are imported also.
- The dashboard widgets that you can view depend on the type of appliance you are using and on your user role. For example, a dashboard created on the Defense Center and imported onto a managed device may display some invalid, disabled widgets.
- 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 devices managed by the importing Defense Center. When you map zones, their types must match. Therefore, you must create any zone types you need on the importing Defense Center before you begin the import. For more information about security zones, see [Working with Security Zones](#) on page 187.
- If you import an access control policy or saved search 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 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.

IMPORTANT! You cannot use the Import/Export feature to update rules created by Sourcefire's Vulnerability Research Team (VRT). Instead, download and apply the latest rule update version; see [Importing Rule Updates and Local Rule Files](#) on page 2072.

- When you import a system policy that was exported from a Defense Center where external authentication is enabled, you also import the authentication objects on which the system policy depends.

Because you can export several different objects in a single package, when you import the package you must choose which objects in the package to import. You can only import objects that are supported on the destination appliance.

When you attempt to import an object, your appliance determines whether that object already exists on the appliance. If a conflict exists, you can keep the existing object, replace the existing object with a new object, keep the newest object, or import the object as a new object. If you import an object and then later make a modification to the object on the destination system, and then re-import the object, you must choose which version of the object to keep.

Depending on the number of objects being imported and the number of objects those objects reference, the import process may take several minutes.

For information on using imported objects, see the following sections:

- [Working with Alert Responses](#) on page 526
- [Using Custom Tables](#) on page 1779
- [Managing Custom User Roles](#) on page 1908
- [Using Custom Workflows](#) on page 1840
- [Working with Dashboards](#) on page 106
- [Applying an Access Control Policy](#) on page 466
- [Applying Health Policies](#) on page 2145
- [Managing Intrusion Policies](#) on page 670
- [Exporting and Importing Report Templates](#) on page 1766
- [Loading a Saved Search](#) on page 1774
- [Applying a System Policy](#) on page 1965
- [Mapping Third-Party Products](#) on page 1682
- [Mapping Third-Party Vulnerabilities](#) on page 1687
- [Activating and Deactivating Detectors](#) on page 1678

To import one or more objects:

ACCESS: Admin

1. Make sure that the appliance where you are exporting the objects and the appliance where you plan to import the objects are running the same version of the Sourcefire 3D System. If you want to import an intrusion policy (or an access control policy that incorporates an intrusion policy), you must also make sure that the rule update versions match.

If the versions of the Sourcefire 3D System (and, if applicable, the rule update versions) do not match, the import will fail.

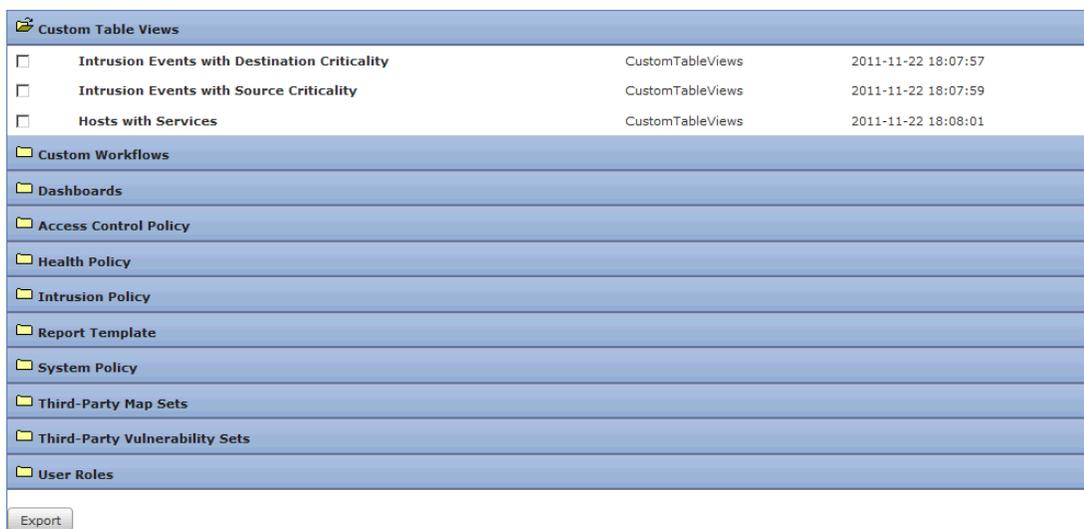
2. Export the objects you want to import; see [Exporting Objects](#) on page 2226.

3. On the appliance where you want to import the objects, select **System > Tools > Import/Export**.

The Import/Export page appears.

TIP! Click the collapse icon (🔼) next to an object type to collapse the list of objects. Click the expand folder icon (📁) next to an object type to reveal objects.

The Defense Center version of the page is shown below with some object types collapsed.



4. Click **Upload Package**.

The Upload Package page appears.

Package Name Browse...

Upload Cancel

5. You have two options:
 - Type the path to the package you want to upload.
 - Click **Browse** to browse to locate the package.

6. Click **Upload.**

The result of the upload depends on the contents of the package:

- If the objects 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 objects, so you do not need to import them.
- If there is a Sourcefire 3D System 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 Sourcefire 3D System or the rule update version and attempt the process again.
- If the package contains any object or rule versions that do not exist on your appliance, the Package Import page appears. Continue with the next step.

7. Select the objects you want to import and click **Import.**

The import process resolves, with the following results:

- If the objects you import do not have previous revisions on your appliance, 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 objects you import do have previous revisions on your appliance, the Import Resolution page appears. Continue with step 9.

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.

9. Expand each object and select the appropriate option:

- To keep the object on your appliance, select **Keep existing**.
- To replace the object on your appliance with the imported object, select **Replace existing**.
- To keep the newest object, select **Keep newest**.

- To save the imported object as a new object, select **Import as new** and, optionally, edit the object name.

If you are importing an access control policy that includes a file policy with Global Malware Whitelist 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.

10. Click **Import**.

The objects are imported.

APPENDIX B

PURGING DISCOVERY DATA FROM THE DATABASE

You can use the Discovery Data Purge page to purge files from the network discovery and user discovery event databases. Note that when you purge a database, the appropriate process is restarted.

WARNING! Purging a database removes the data you specify from the Defense Center. After the data is deleted, it **cannot** be recovered.

To purge the network and user discovery database:

ACCESS: Admin/Any Security Analyst

1. Select **System > Tools > Discovery Data Purge**.
The Discovery Data Purge page appears.
2. Under **Network Discovery Database**, perform any or all of the following:
 - Select **Network Discovery Events** to remove all network discovery events from the database.
 - Select **Hosts** to remove all hosts from the database.
 - Select **Connection Events** to remove all connection data from the database.
 - Select **Connection Summary Events** to remove all connection summary data from the database.

3. Under **User Discovery Database**, perform any or all of the following:
 - Select **User Activity** to remove all user activity events from the events database.
 - Select **User Identity** to remove all users from the database.
4. Click **Save & Restart Network/User Discovery**.
The items are purged and the appropriate processes are restarted.

APPENDIX C

VIEWING THE STATUS OF LONG-RUNNING TASKS

Some tasks that you can perform on the Sourcefire 3D System, 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](#) on page 2238
- [Managing the Task Queue](#) on page 2240

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. You can always see the status of tasks that you initiated. If your user account has the Administrator user role, or a user role with the **View Other Users' Tasks** permission enabled, you can see the status of every task,

regardless of who initiated it. For more information on configuring user roles, see [Configuring User Roles](#) on page 1905.

Job Summary Remove Completed Jobs Remove Failed Jobs

Running	0
Waiting	0
Completed	14
Retrying	0
Failed	0

Jobs

Task Description	Message	Creation Time	Last Change	Status	
Health Policy Apply 0 Running 0 Waiting 1 Completed 0 Retrying 0 Failed					
Health policy apply to appliance katsura Health Policy Apply	Health Policy Applied! Initial_Health_Policy 2011-11-03 14:24:10 -> katsura	2011-11-04 17:49:32	2011-11-04 17:50:32	Completed	
Network Discovery Policy Apply 0 Running 0 Waiting 11 Completed 0 Retrying 0 Failed					
Network Discovery policy apply to linden Network Discovery Policy	Network Discovery policy successfully apply	2011-11-06 17:51:17	2011-11-06 17:53:14	Completed	
Apply TestPolicy to linden Access Control Policy	Access Control Policy applied successfully	2011-11-06 17:51:17	2011-11-06 17:52:18	Completed	

The Job Summary section displays the state of the tasks listed on the page, as described in the following [Task Queue Task Types](#) table.

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, such as Network Discovery Policy Apply, appear together in a task group.

To make sure that the Task Status page loads quickly, once per week, the Sourcefire 3D System 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:

ACCESS: Admin/Maint/Network Admin/Security Approver/Security Analyst

► 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 **System > 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

If your user account is assigned the Administrator, Maintenance User, Network Admin, Security Approver, or Security Analyst user role, there are several actions

you can perform while viewing the task queue (see [Viewing the Task Queue](#) on page 2238), as described in the following table.

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 Sourcefire 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.

APPENDIX D

COMMAND LINE REFERENCE

This reference explains the command line interface (CLI) for Sourcefire systems. You can use the CLI to view, configure, and troubleshoot your Sourcefire systems.

IMPORTANT! The command line interface is not supported on Defense Centers and Series 2 appliances.

There are numerous CLI modes, such as **show** and **configuration**, that contain sets of commands beginning with the mode name. You may enter a mode and then enter valid commands within that mode, or you may enter an entire full command from any mode. For example, to display information about a user account called Analyst1, you can enter the following at the CLI prompt:

```
show user Analyst1
```

If you have previously entered **show** mode, enter the following at the CLI prompt:

```
user Analyst1
```

Within each mode, the commands available to a user depend on the user's CLI access. When you create a user account, you can assign it one of the following CLI access levels:

- Basic

The user has read-only access and cannot run commands that impact system performance.

- Configuration
The user has read-write access and can run commands that impact system performance.
- None
The user is unable to log in to the shell.

On Series 3 devices, you can assign command line permissions on the User Management page in the web interface; see [Managing Users](#) on page 1849 for more information. On virtual devices, you assign command line permissions through the CLI itself.

IMPORTANT! If you reboot a Series 3 device and then log in to the CLI as soon as you are able, any commands you execute are not recorded in the audit log until the web interface is available.

Note that CLI commands are case-insensitive with the exception of parameters whose text is not part of the CLI framework, such as user names and search filters.

For information about logging into the command line, see [Logging into the Appliance](#) on page 58.

The following sections describe the CLI commands:

- [Basic CLI Commands](#) on page 2243
- [Show Commands](#) on page 2246
- [Configuration Commands](#) on page 2268
- [System Commands](#) on page 2279

Basic CLI Commands

The basic CLI commands provide the ability to interact with the CLI. These commands do not affect the operation of the device. Basic commands are available to all CLI users.

The following sections describe the basic commands:

- [configure password](#) on page 2244
- [end](#) on page 2244
- [exit](#) on page 2244
- [help](#) on page 2245
- [history](#) on page 2245
- [logout](#) on page 2245
- [? \(question mark\)](#) on page 2246
- [?? \(double question marks\)](#) on page 2246

configure password

Allows the current user to change their password. After issuing the command, the CLI prompts the user for their current (or old) password, then prompts the user to enter the new password twice.

Access

Basic

Syntax

```
configure password
```

Example

```
> configure password
Enter current password:
Enter new password:
Confirm new password:
```

end

Returns the user to the default mode. (Moves the user up to the default mode from any lower-level CLI context.)

Access

Basic

Syntax

```
end
```

Example

```
configure network ipv4> end
>
```

exit

Moves the CLI context up to the next highest CLI context level. Issuing this command from the default mode logs the user out of the current CLI session, and is equivalent to issuing the `logout` CLI command.

Access

Basic

Syntax

```
exit
```

Example

```
configure network ipv4> exit
configure network>
```

help

Displays an overview of the CLI syntax.

Access

Basic

Syntax

```
help
```

Example

```
> help
```

history

Displays the command line history for the current session.

Access

Basic

Syntax

```
history limit
```

where *limit* sets the size of the history list. To set the size to unlimited, enter zero.

Example

```
history 25
```

logout

Logs the current user out of the current CLI console session.

Access

Basic

Syntax

```
logout
```

Example

```
> logout
```

? (question mark)

Displays context-sensitive help for CLI commands and parameters. Use the question mark (?) command as follows:

- To display help for the commands that are available within the current CLI context, enter a question mark (?) at the command prompt.
- To display a list of the available commands that start with a particular character set, enter the abbreviated command immediately followed by a question mark (?).
- To display help for a command's legal arguments, enter a question mark (?) in place of an argument at the command prompt.

Note that the question mark (?) is not echoed back to the console.

Access

Basic

Syntax

```
?  
abbreviated_command ?  
command [arguments] ?
```

Example

```
> ?
```

?? (double question marks)

Displays detailed context-sensitive help for CLI commands and parameters.

Access

Basic

Syntax

```
??  
abbreviated_command end??  
command [arguments] ??
```

Example

```
> configure manager add ??
```

Show Commands

Show commands provide information about the state of the device. These commands do not change the operational mode of the device and running them has minimal impact on system operation. Most show commands are available to

all CLI users; however, only users with configuration CLI access can issue the `show user` command.

The following sections describe the show commands:

- [access-control-config](#) on page 2248
- [alarms](#) on page 2248
- [arp-tables](#) on page 2249
- [audit-log](#) on page 2249
- [clustering](#) on page 2249
- [cpu](#) on page 2250
- [database](#) on page 2251
- [device-settings](#) on page 2252
- [disk](#) on page 2252
- [dns](#) on page 2252
- [expert](#) on page 2253
- [bypass](#) on page 2253
- [fan-status](#) on page 2253
- [fastpath-rules](#) on page 2254
- [gui](#) on page 2254
- [hostname](#) on page 2254
- [hyperthreading](#) on page 2255
- [inline-sets](#) on page 2255
- [interfaces](#) on page 2255
- [lcd](#) on page 2256
- [link-state](#) on page 2256
- [log-ips-connection](#) on page 2256
- [managers](#) on page 2257
- [memory](#) on page 2257
- [model](#) on page 2257
- [mpls-depth](#) on page 2258
- [NAT](#) on page 2258
- [network](#) on page 2260
- [network-modules](#) on page 2260
- [ntp](#) on page 2260
- [perfstats](#) on page 2261
- [portstats](#) on page 2261
- [power-supply-status](#) on page 2261

- [process-tree](#) on page 2262
- [processes](#) on page 2262
- [routing-table](#) on page 2262
- [serial-number](#) on page 2263
- [stacking](#) on page 2263
- [summary](#) on page 2263
- [time](#) on page 2264
- [traffic-statistics](#) on page 2264
- [user](#) on page 2264
- [users](#) on page 2265
- [version](#) on page 2266
- [virtual-routers](#) on page 2266
- [virtual-switches](#) on page 2267

access-control-config

Displays the currently applied access control configurations, including logging settings, Security Intelligence settings, file and malware detection limits (by file size), HTTP Response page content, all enabled access control rules, source ports for access control rules, destination port data (including type and code for ICMP entries) and the number of connections that matched each access control rule (hit counts).

Access

Basic

Syntax

```
show access-control-config
```

Example

```
> show access-control-config
```

alarms

Displays currently active (failed/down) hardware alarms on the device. This command is not available on virtual devices.

Access

Basic

Syntax

```
show alarms
```

Example

```
> show alarms
```

arp-tables

Displays the Address Resolution Protocol tables applicable to your network. This command is not available on virtual devices.

Access

Basic

Syntax

```
show arp-tables
```

Example

```
> show arp-tables
```

audit-log

Displays the audit log in reverse chronological order; the most recent audit log events are listed first.

Access

Basic

Syntax

```
show audit-log
```

Example

```
> show audit-log
```

clustering

Displays information about device clustering configuration, status, and member stacks. This command is not available on virtual devices.

Access

Basic

config

Displays the clustering configuration on the device.

Syntax

```
show clustering config
```

Example

```
> show clustering config
```

clustering ha-statistics

Displays state sharing statistics for a device in a cluster.

Access

Basic

Syntax

```
show clustering ha-statistics
```

Example

```
> show clustering ha-statistics
```

cpu

Displays the current CPU usage statistics appropriate for the platform for all CPUs on the device. For managed devices, the following values are displayed:

- CPU
Processor number.
- Load
The CPU utilization, represented as a number from 0 to 100. 0 is not loaded and 100 is completely loaded.

For virtual devices, the following values are displayed:

- CPU
Processor number.
- %user
Percentage of CPU utilization that occurred while executing at the user level (application).
- %nice
Percentage of CPU utilization that occurred while executing at the user level with nice priority.
- %sys
Percentage of CPU utilization that occurred while executing at the system level (kernel). This does not include time spent servicing interrupts or softirqs. A softirq (software interrupt) is one of up to 32 enumerated software interrupts that can run on multiple CPUs at once.
- %iowait
Percentage of time that the CPUs were idle when the system had an outstanding disk I/O request.
- %irq
Percentage of time spent by the CPUs to service interrupts.

- %soft
Percentage of time spent by the CPUs to service softirqs.
- %steal
Percentage of time spent in involuntary wait by the virtual CPUs while the hypervisor was servicing another virtual processor.
- %guest
Percentage of time spent by the CPUs to run a virtual processor.
- %idle
Percentage of time that the CPUs were idle and the system did not have an outstanding disk I/O request.

Access

Basic

Syntax

```
show cpu [procnum]
```

where *procnum* is the number of the processor for which you want the utilization information displayed. Valid values are 0 to one less than the total number of processors on the system. If *procnum* is used for a managed device, it is ignored because for that platform, utilization information can only be displayed for all processors.

Example

```
> show cpu
```

database

The show database commands configure the device's management interface.

Access

Basic

processes

Displays a list of running database queries.

Access

Basic

Syntax

```
show database processes
```

Example

```
> show database processes
```

slow-query-log

Displays the slow query log of the database.

Access

Basic

Syntax

```
show database slow-query-log
```

Example

```
> show database slow-query-log
```

device-settings

Displays information about application bypass settings specific to the current device.

Access

Basic

Syntax

```
show device-settings
```

Example

```
> show device-settings
```

disk

Displays the current disk usage.

Access

Basic

Syntax

```
show disk
```

Example

```
> show disk
```

dns

Displays the current DNS server addresses and search domains.

Access

Basic

Syntax

```
show dns
```

Example

```
> show dns
```

expert

Invokes the shell.

Access

Basic

Syntax

```
expert
```

Example

```
> expert
```

bypass

Lists the inline sets in use and shows the bypass mode status of those sets, either normal or bypass. This command is not available on virtual devices.

Access

Basic

Syntax

```
show bypass
```

Example

```
> show bypass
```

fan-status

Displays the current status of hardware fans. This command is not available on virtual devices.

Access

Basic

Syntax

```
show fan-status
```

Example

```
> show fan-status
```

fastpath-rules

Displays the currently configured fastpath rules. This command is not available on virtual devices.

Access

Basic

Syntax

```
show fastpath-rules
```

Example

```
> show fastpath-rules
```

gui

Displays the current state of the web interface. This command is not available on virtual devices.

Access

Basic

Syntax

```
show gui
```

Example

```
> show gui
```

hostname

Displays the device's hostname and appliance UUID.

Access

Basic

Syntax

```
show hostname
```

Example

```
> show hostname
```

hyperthreading

Displays whether hyperthreading is enabled or disabled.

Access

Basic

Syntax

```
show hyperthreading
```

Example

```
> show hyperthreading
```

inline-sets

Displays configuration data for all inline security zones and associated interfaces.

Access

Basic

Syntax

```
show inline-sets
```

Example

```
> show inline-sets
```

interfaces

If no parameters are specified, displays a list of all configured interfaces. If a parameter is specified, displays detailed information about the specified interface.

Access

Basic

Syntax

```
show interfaces [interface]
```

where *interface* is the specific interface for which you want the detailed information.

Example

```
> show interfaces
```

lcd

Displays whether the LCD hardware display is enabled or disabled. This command is not available on virtual devices.

Access

Basic

Syntax

```
show lcd
```

Example

```
> show lcd
```

link-state

Displays type, link, speed, duplex state, and bypass mode of the ports on the device.

Access

Basic

Syntax

```
show link-state
```

Example

```
> show link-state
```

log-ips-connection

Displays whether the logging of connection events that are associated with logged intrusion events is enabled or disabled.

Access

Basic

Syntax

```
show log-ips-connection
```

Example

```
> show log-ips-connection
```

managers

Displays the configuration and communication status of the Defense Center. Registration key and NAT ID are only displayed if registration is pending. If a device is registered to a high availability pair, information about both managing Defense Centers is displayed. If a device is configured as a secondary device in a stacked configuration, information about both the managing Defense Center and the primary device is displayed.

Access

Basic

Syntax

```
show managers
```

Example

```
> show managers
```

memory

Displays the total memory, the memory in use, and the available memory for the device.

Access

Basic

Syntax

```
show memory
```

Example

```
> show memory
```

model

Displays model information for the device.

Access

Basic

Syntax

```
show model
```

Example

```
> show model
```

mpls-depth

Displays the number of MPLS layers configured on the management interface, from 0 to 6. This command is not available on virtual devices.

Access

Basic

Syntax

```
show mpls-depth
```

Example

```
> show mpls-depth
```

NAT

The `show nat` commands display NAT data and configuration information for the management interface. This command is not available on virtual devices.

Access

Basic

active-dynamic

Displays NAT flows translated according to dynamic rules. These entries are displayed when a flow matches a rule, and persist until the rule has timed out. Therefore, the list can be inaccurate.

Syntax

```
show nat active-dynamic
```

Example

```
> show nat active-dynamic
```

active-static

Displays NAT flows translated according to static rules. These entries are displayed as soon as you apply the rule to the device, and the list does not indicate active flows that match a static NAT rule.

Syntax

```
show nat active-static
```

Example

```
> show nat active-static
```

allocators

Displays information for all NAT allocators, the pool of translated addresses used by dynamic rules

Syntax

```
show nat allocators
```

Example

```
> show nat allocators
```

config

Displays the current NAT policy configuration for the management interface.

Syntax

```
show nat config
```

Example

```
> show nat config
```

dynamic-rules

Displays dynamic NAT rules that use the specified allocator ID.

Syntax

```
show nat dynamic-rules allocator_id
```

Example

```
> show nat dynamic-rules 9
```

where *allocator_id* is a valid allocator ID number.

flows

Displays the number of flows for rules that use the specified allocator ID.

Syntax

```
show nat flows allocator-id
```

Example

```
> show nat flows 81
```

where *allocator_id* is a valid allocator ID number.

static-rules

Displays all static NAT rules.

Syntax

```
show nat static-rules
```

Example

```
> show nat static-rules
```

network

Displays the IPv4 and IPv6 configuration of the management interface, as well as its MAC address.

Access

Basic

Syntax

```
show network
```

Example

```
> show network
```

network-modules

Displays all installed modules and information about them, including serial numbers. This command is not available on virtual devices.

Access

Basic

Syntax

```
show network-modules
```

Example

```
> show network-modules
```

ntp

Displays the ntp configuration.

Access

Basic

Syntax

```
show ntp
```

Example

```
> show ntp
```

perfstats

Displays performance statistics for the device.

Access

Basic

Syntax

```
show perfstats
```

Example

```
> show perfstats
```

portstats

Displays port statistics for all installed ports on the device. This command is not available on virtual devices.

Access

Basic

Syntax

```
show portstats [copper | fiber | internal | external | all]
```

where **copper** specifies for all copper ports, **fiber** specifies for all fiber ports, **internal** specifies for all internal ports, **external** specifies for all external (copper and fiber) ports, and **all** specifies for all ports (external and internal).

Example

```
> show portstats fiber
```

power-supply-status

Displays the current state of hardware power supplies. This command is not available on virtual devices.

Access

Basic

Syntax

```
show power-supply-status
```

Example

```
> show power-supply-status
```

process-tree

Displays processes currently running on the device, sorted in tree format by type.

Access

Basic

Syntax

```
show process-tree
```

Example

```
> show process-tree
```

processes

Displays processes currently running on the device, sorted by descending CPU usage.

Access

Basic

Syntax

```
show processes [sort-flag] [filter]
```

where *sort-flag* can be *-m* to sort by memory (descending order), *-u* to sort by username rather than the process name, or *verbose* to display the full name and path of the command. The *filter* parameter specifies the search term in the command or username by which results are filtered. The header row is still displayed.

Example

```
> show processes -u user1
```

routing-table

If no parameters are specified, displays routing information for all virtual routers. If parameters are specified, displays routing information for the specified router and, as applicable, its specified routing protocol type. All parameters are optional. This command is not available on virtual devices.

Access

Basic

Syntax

```
show routing-table [name] [ospf | rip | static ]
```

where *name* is the name of the specific router for which you want information, and *ospf*, *rip*, and *static* specify the routing protocol type.

Example

```
> show routing-table vrouter1 static
```

serial-number

Displays the chassis serial number. This command is not available on virtual devices.

Access

Basic

Syntax

```
show serial-number
```

Example

```
> show serial-number
```

stacking

Shows the stacking configuration and position on managed devices; on devices configured as primary, also lists data for all secondary devices. For clustered stacks, this command also indicates that the stack is a member of a cluster. The user must use the web interface to enable or (in most cases) disable stacking; if stacking is not enabled, the command will return **stacking not currently configured**. This command is not available on virtual devices.

Access

Basic

Syntax

```
show stacking
```

Example

```
> show stacking
```

summary

Displays a summary of the most commonly used information (version, type, UUID, and so on) about the device. For more detailed information, see the following `show` commands: [version](#) on page 2266, [interfaces](#) on page 2255, [device-settings](#) on page 2252, and [access-control-config](#) on page 2248.

Access

Basic

Syntax

```
show summary
```

Example

```
> show summary
```

time

Displays the current date and time in UTC and in the local time zone configured for the current user.

Access

Basic

Syntax

```
show time
```

Example

```
> show time
```

traffic-statistics

If no parameters are specified, displays details about bytes transmitted and received from all ports. If a port is specified, displays that information only for the specified port.

Access

Basic

Syntax

```
show traffic-statistics [port]
```

where `port` is the specific port for which you want information.

Example

```
> show traffic-statistics s1p1
```

user

Applicable to virtual devices only. Displays detailed configuration information for the specified user(s). The following values are displayed:

- Login — the login name
- UID — the numeric user ID
- Auth (**Local** or **Remote**) — how the user is authenticated
- Access (**Basic** or **Config**) — the user's privilege level
- Enabled (**Enabled** or **Disabled**) — whether the user is active
- Reset (**Yes** or **No**) — whether the user must change password at next login
- Exp (**Never** or a number) — the number of days until the user's password must be changed
- Warn (**N/A** or a number) — the number of days a user is given to change their password before it expires

- Str (**Yes** or **No**) — whether the user's password must meet strength checking criteria
- Lock (**Yes** or **No**) — whether the user's account has been locked due to too many login failures
- Max (**N/A** or a number) — the maximum number of failed logins before the user's account is locked

Access

Configuration

Syntax

```
show user username username username ...
```

where *username* specifies the name of the user and the usernames are space-separated.

Example

```
> show user jdoe
```

users

Applicable to virtual devices only. Displays detailed configuration information for all local users. The following values are displayed:

- Login — the login name
- UID — the numeric user ID
- Auth (**Local** or **Remote**) — how the user is authenticated
- Access (**Basic** or **Config**) — the user's privilege level
- Enabled (**Enabled** or **Disabled**) — whether the user is active
- Reset (**Yes** or **No**) — whether the user must change password at next login
- Exp (**Never** or a number) — the number of days until the user's password must be changed
- Warn (**N/A** or a number) — the number of days a user is given to change their password before it expires
- Str (**Yes** or **No**) — whether the user's password must meet strength checking criteria

- Lock (Yes or No) — whether the user's account is locked due to too many login failures
- Max (N/A or a number) — the maximum number of failed logins before the user's account is locked

Access

Configuration

Syntax

```
show users
```

Example

```
> show users
```

version

Displays the product version and build. If the *detail* parameter is specified, displays the versions of additional components.

Access

Basic

Syntax

```
show version [detail]
```

Example

```
> show version
```

virtual-routers

If no parameters are specified, displays a list of all currently configured virtual routers with DHCP relay, OSPF, and RIP information. If parameters are specified, displays information for the specified router, limited by the specified route type. All parameters are optional. This command is not available on virtual devices.

Access

Basic

Syntax

```
show virtual-routers [ dhcprelay | ospf | rip ] [name]
```

where *dhcprelay*, *ospf*, and *rip* specify for route types, and *name* is the name of the specific router for which you want information. If you specify *ospf*, you can then further specify *neighbors*, *topology*, or *lsadb* between the route type and (if present) the router name.

Example

```
> show virtual-routers ospf VRouter2
```

virtual-switches

If no parameters are specified, displays a list of all currently configured virtual switches. If parameters are specified, displays information for the specified switch. This command is not available on virtual devices.

Access

Basic

Syntax

```
show virtual-switches [name]
```

Example

```
> show virtual-switches vswitch1
```

VPN

The `show VPN` commands display VPN status and configuration information for VPN connections. This command is not available on virtual devices.

Access

Basic

config

Displays the configuration of all VPN connections.

Syntax

```
show vpn config
```

Example

```
> show vpn config
```

config by virtual router

Displays the configuration of all VPN connections for a virtual router.

Syntax

```
show vpn config [virtual router]
```

Example

```
> show vpn config VRouter1
```

status

Displays the status of all VPN connections.

Syntax

```
show vpn status
```

Example

```
> show vpn status
```

status by virtual router

Displays the status of all VPN connections for a virtual router.

Syntax

```
show vpn status [virtual router]
```

Example

```
> show vpn status VRouter1
```

counters

Displays the counters for all VPN connections.

Syntax

```
show vpn counters
```

Example

```
> show vpn counters
```

counters by virtual router

Displays the counters of all VPN connections for a virtual router.

Syntax

```
show vpn counters [virtual router]
```

Example

```
> show vpn counters VRouter1
```

Configuration Commands

The configuration commands enable the user to configure and manage the system. These commands affect system operation; therefore, with the exception of Basic-level **configure password**, only users with configuration CLI access can issue these commands.

The following sections describe the configuration commands:

- [clustering](#) on page 2269
- [bypass](#) on page 2269
- [gui](#) on page 2269
- [lcd](#) on page 2270
- [log-ips-connections](#) on page 2270
- [manager](#) on page 2270
- [mpls-depth](#) on page 2271

- [network](#) on page 2271
- [password](#) on page 2275
- [stacking disable](#) on page 2275
- [user](#) on page 2275

clustering

Disables or configures bypass for clustering on the device. This command is not available on virtual devices or on devices configured as secondary stack members.

Access

Configuration

Syntax

```
configure clustering {disable | bypass}
```

Example

```
> configure clustering disable
```

bypass

Opens or closes the bypass mode of an inline pair. This command is not available on virtual devices.

Access

Configuration

Syntax

```
configure bypass {open | close} {interface}
```

where *interface* is the name of either hardware port in the inline pair.

Example

```
> configure bypass open s1p1
```

gui

Enables or disables the device web interface, including the streamlined upgrade web interface that appears during major updates to the system. This command is not available on virtual devices.

Access

Configuration

Syntax

```
configure gui {enable | disable}
```

Example

```
> configure gui disable
```

lcd

Enables or disables the LCD display on the front of the device. This command is not available on virtual devices.

Access

Configuration

Syntax

```
configure lcd {enable | disable}
```

Example

```
> configure lcd disable
```

log-ips-connections

Enables or disables logging of connection events that are associated with logged intrusion events.

Access

Configuration

Syntax

```
configure log-ips-connections {enable | disable}
```

Example

```
> configure log-ips-connections disable
```

manager

The `configure manager` commands configure the device's connection to its managing Defense Center.

Access

Configuration

add

Configures the device to accept a connection from a managing Defense Center. This command works only if the device is not actively managed.

A unique alphanumeric registration key is always required to register a device to a Defense Center. In most cases, you must provide the hostname or the IP address along with the registration key. However, if the device and the Defense Center are separated by a NAT device, you must enter a unique NAT ID, along with the registration key, and specify `DONTRESOLVE` instead of the hostname.

Syntax

```
configure manager add {hostname | IPv4_address |  
IPv6_address | DONTRESOLVE} regkey [nat_id]
```

where `{hostname | IPv4_address | IPv6_address | DONTRESOLVE}` specifies the DNS host name or IP address (IPv4 or IPv6) of the Defense Center that manages this device. If the Defense Center is not directly addressable, use `DONTRESOLVE`. If you use `DONTRESOLVE`, `nat_id` is required. `regkey` is the unique alphanumeric registration key required to register a device to the Defense Center. `nat_id` is an optional alphanumeric string used during the registration process between the Defense Center and the device. It is required if the hostname is set to `DONTRESOLVE`.

Example

```
> configure manager add DONTRESOLVE abc123 efg456
```

delete

Removes the Defense Center's connection information from the device. This command only works if the device is not actively managed.

Syntax

```
configure manager delete
```

Example

```
> configure manager delete
```

mpls-depth

Configures the number of MPLS layers on the management interface. This command is not available on virtual devices.

Access

Configuration

Syntax

```
configure mpls-depth {depth}
```

where `depth` is a number between 0 and 6.

Example

```
> configure mpls-depth 3
```

network

The `configure network` commands configure the device's management interface.

Access

Configuration

dns searchdomains

Replaces the current list of DNS search domains with the list specified in the command.

Syntax

```
configure network dns searchdomains {searchlist}
```

where *searchlist* is a comma-separated list of domains.

Example

```
> configure network dns searchdomains foo.bar.com,bar.com
```

dns servers

Replaces the current list of DNS servers with the list specified in the command.

Syntax

```
configure network dns servers {dnslist}
```

where *dnslist* is a comma-separated list of DNS servers.

Example

```
> configure network dns servers 10.123.1.10,10.124.1.10
```

hostname

Sets the hostname for the device.

Syntax

```
configure network hostname {name}
```

where *name* is the new hostname.

Example

```
> configure network hostname sfrocks
```

ipv4 delete

Disables the IPv4 configuration of the device's management interface.

Syntax

```
configure network ipv4 delete
```

Example

```
> configure network ipv4 delete
```

ipv4 dhcp

Sets the IPv4 configuration of the device's management interface to DHCP. The management interface communicates with the DHCP server to obtain its configuration information.

Syntax

```
configure network ipv4 dhcp
```

Example

```
> configure network ipv4 dhcp
```

ipv4 manual

Manually configures the IPv4 configuration of the device's management interface.

Syntax

```
configure network ipv4 manual ipaddr netmask gw
```

where *ipaddr* is the IP address, *netmask* is the subnet mask, and *gw* is the IPv4 address of the default gateway.

Example

```
> configure network ipv4 manual 10.123.1.10 255.255.0.0  
10.123.1.1
```

ipv6 delete

Disables the IPv6 configuration of the device's management interface.

Syntax

```
configure network ipv6 delete
```

Example

```
> configure network ipv6 delete
```

ipv6 dhcp

Sets the IPv6 configuration of the device's management interface to DHCP. The management interface communicates with the DHCP server to obtain its configuration information.

Syntax

```
configure network ipv6 dhcp
```

Example

```
> configure network ipv6 dhcp
```

ipv6 router

Sets the IPv6 configuration of the device's management interface to Router. The management interface communicates with the IPv6 router to obtain its configuration information.

Syntax

```
configure network ipv6 router
```

Example

```
> configure network ipv6 router
```

ipv6 manual

Manually configures the IPv6 configuration of the device's management interface.

Syntax

```
configure network ipv6 manual ip6addr/ip6prefix [ip6gw]
```

where *ip6addr/ip6prefix* is the IP address and prefix length and *ip6gw* is the IPv6 address of the default gateway.

Example

```
> configure network ipv6 manual  
2001:DB8:3ffe:1900:4545:3:200:f8ff:fe21:67cf 64
```

management-port

Sets the value of the device's TCP management port.

Syntax

```
configure network management-port number
```

where *number* is the management port value you want to configure.

Example

```
> configure network management-port 8500
```

password

Allows the current user to change their password. After issuing the command, the CLI prompts the user for their current (or old) password, then prompts the user to enter the new password twice.

Access

Basic

Syntax

```
configure password
```

Example

```
> configure password
Enter current password:
Enter new password:
Confirm new password:
```

stacking disable

On managed devices, removes any stacking configuration present on that device: on devices configured as primary, the stack is removed entirely; on devices configured as secondary, that device is removed from the stack. This command is not available on virtual devices, and you cannot use it to break a clustered stack.

Use this command when you cannot establish communication with appliances higher in the stacking hierarchy. If the Defense Center is available for communication, a message appears instructing you to use the Defense Center web interface instead; likewise, if you enter **stacking disable** on a device configured as secondary when the primary device is available, a message appears instructing you to enter the command from the primary device.

Access

Configuration

Syntax

```
configure stacking disable
```

Example

```
> configure stacking disable
```

user

Applicable only to virtual devices, the **configure user** commands manage the device's local user database.

Access

Configuration

access

Modifies the access level of the specified user. This command takes effect the next time the specified user logs in.

Syntax

```
configure user access username [basic | config]
```

Example

```
> configure user access jdoe basic
```

where *username* specifies the name of the user for which you want to modify access, **basic** indicates basic access, and **config** indicates configuration access.

add

Creates a new user with the specified name and access level. This command prompts for the user's password.

Syntax

```
configure user add username [basic | config]
```

where *username* specifies the name of the new user, **basic** indicates basic access, and **config** indicates configuration access.

Example

```
> configure user add jdoe basic  
Enter new password for user jdoe:  
Confirm new password for user jdoe:
```

aging

Forces the expiration of the user's password.

Syntax

```
configure user aging username max_days warn_days
```

where *username* specifies the name of the user, *max_days* indicates the maximum number of days that the password is valid, and *warn_days* indicates the number of days that the user is given to change the password before it expires.

Example

```
> configure user aging jdoe 100 3
```

delete

Deletes the user and the user's home directory.

Syntax

```
configure user delete username
```

where *username* specifies the name of the user.

Example

```
> configure user delete jdoe
```

disable

Disables the user. Disabled users cannot login.

Syntax

```
configure user disable username
```

where *username* specifies the name of the user.

Example

```
> configure user disable jdoe
```

enable

Enables the user.

Syntax

```
configure user enable username
```

where *username* specifies the name of the user.

Example

```
> configure user enable jdoe
```

forcereset

Forces the user to change their password the next time they login. When the user logs in and changes the password, strength checking is automatically enabled.

Syntax

```
configure user forcereset username
```

where *username* specifies the name of the user.

Example

```
> configure user forcereset jdoe
```

maxfailedlogins

Sets the maximum number of failed logins for the specified user.

Syntax

```
configure user maxfailedlogins username number
```

where *username* specifies the name of the user and *number* specifies the maximum number of failed logins.

Example

```
> configure user maxfailedlogins jdoe 3
```

password

Sets the user's password. This command prompts for the user's password.

Syntax

```
configure user password username
```

where *username* specifies the name of the user.

Example

```
> configure user password jdoe
Enter new password for user jdoe:
Confirm new password for user jdoe:
```

strengthcheck

Enables or disables the strength requirement for a user's password. When a user's password expires or if the configure user forcereset command is used, this requirement is automatically enabled the next time the user logs in.

Syntax

```
configure user strengthcheck username {enable | disable}
```

where *username* specifies the name of the user, **enable** sets the requirement for the specified user's password, and **disable** removes the requirement for the specified user's password.

Example

```
> configure user strengthcheck jdoe enable
```

unlock

Unlocks a user that has exceeded the maximum number of failed logins.

Syntax

```
configure user unlock username
```

where *username* specifies the name of the user.

Example

```
> configure user unlock jdoe
```

System Commands

The system commands enable the user to manage system-wide files and access control settings. Only users with configuration CLI access can issue commands in system mode.

The following sections describe the system commands:

- [access-control](#) on page 2279
- [disable-http-user-cert](#) on page 2280
- [file](#) on page 2280
- [generate-troubleshoot](#) on page 2282
- [ldapsearch](#) on page 2282
- [lockdown-sensor](#) on page 2283
- [nat rollback](#) on page 2283
- [reboot](#) on page 2283
- [restart](#) on page 2284
- [shutdown](#) on page 2284

access-control

The system `access-control` commands enable the user to manage the access control configuration on the device.

Access

Configuration

archive

Saves the currently applied access control policy as a text file on `/var/common`.

Syntax

```
system access-control archive
```

Example

```
> system access-control archive
```

clear-rule-counts

Resets the access control rule hit count to 0.

Syntax

```
system access-control clear-rule-counts
```

Example

```
> system access-control clear-rule-counts
```

rollback

Reverts the system to the previously applied access control configuration. You cannot use this command with clustered or stacked devices.

Syntax

```
system access-control rollback
```

Example

```
> system access-control rollback
```

disable-http-user-cert

Removes all HTTP user certification present on the system.

Access

Configuration

Syntax

```
system disable-http-user-cert
```

Example

```
> system disable-http-user-cert
```

file

The `system file` commands enable the user to manage the files in the common directory on the device.

Access

Configuration

copy

Uses FTP to transfer files to a remote location on the host using the login username. The local files must be located in the common directory.

Syntax

```
system file copy hostname username path filenames filenames  
...
```

where *hostname* specifies the name or ip address of the target remote host, *username* specifies the name of the user on the remote host, *path* specifies the destination path on the remote host, and *filenames* specifies the local files to transfer; the file names are space-separated.

Example

```
> system file copy sfrocks jdoe /pub *
```

delete

Removes the specified files from the common directory.

Syntax

```
system file delete filenames filenames ...
```

where *filenames* specifies the files to delete; the file names are space-separated.

Example

```
> system file delete *
```

list

If no file names are specified, displays the modification time, size, and file name for all the files in the common directory. If file names are specified, displays the modification time, size, and file name for files that match the specified file names.

Syntax

```
system file list {filenames filenames ...}
```

where *filenames* specifies the files to display; the file names are space-separated.

Example

```
> system file list
```

secure-copy

Uses SCP to transfer files to a remote location on the host using the login username. The local files must be located in the `/var/common` directory.

Syntax

```
system file secure-copy hostname username path filenames  
filenames ...
```

where *hostname* specifies the name or ip address of the target remote host, *username* specifies the name of the user on the remote host, *path* specifies the destination path on the remote host, and *filenames* specifies the local files to transfer; the file names are space-separated.

Example

```
> system file secure-copy 10.123.31.1 jdoe /tmp *
```

generate-troubleshoot

Generates troubleshooting data for analysis by Sourcefire.

Access

Configuration

Syntax

```
system generate-troubleshoot
```

This syntax displays a list of optional parameters to specify what troubleshooting data should be displayed.

Example

```
> system generate-troubleshoot
```

ldapsearch

Enables the user to perform a query of the specified LDAP server. Note that all parameters are required.

Access

Configuration

Syntax

```
system ldapsearch host port baseDN userDN basefilter
```

where *host* specifies the LDAP server domain, *port* specifies the LDAP server port, *baseDN* specifies the DN (distinguished name) that you want to search under, *userDN* specifies the DN of the user who binds to the LDAP directory, and *basefilter* specifies the record or records you want to search for.

Example

```
> system ldapsearch ldap.example.com 389 cn=users,  
dc=example,dc=com cn=user1,cn=users,dc=example,dc=com,  
cn=user2
```

lockdown-sensor

Removes the `expert` command and access to the bash shell on the device.

WARNING! This command is irreversible without a hotfix from Sourcefire Support. Use with care.

Access

Configuration

Syntax

```
system lockdown-sensor
```

Example

```
> system lockdown-sensor
```

nat rollback

Reverts the system to the previously applied NAT configuration.

Access

Configuration

Syntax

```
system nat rollback
```

Example

```
> system nat rollback
```

reboot

Reboots the device.

Access

Configuration

Syntax

```
system reboot
```

Example

```
> system reboot
```

restart

Restarts the device application.

Access

Configuration

Syntax

```
system restart
```

Example

```
> system restart
```

shutdown

Shuts down the device.

Access

Configuration

Syntax

```
system shutdown
```

Example

```
> system shutdown
```

APPENDIX E

THIRD-PARTY PRODUCTS

Sourcefire products contain certain third-party open source code products that are distributed for use in combination with the Sourcefire products. These products are free and distributed “as-is” under the terms set forth in their respective license agreements. The following table lists the major open source code products and applicable license agreements that are distributed by Sourcefire for use in combination with the Sourcefire products.

Third-Party Software

SOFTWARE	LICENSE AGREEMENT
Apache HTTPD Web Server 2.4.3	Apache License
Linux Kernel 2.6.32.24 (Series 2)	GNU General Public License Version 2 (CPLv2)
Linux Kernel 2.6.35.14 (Series 3)	GNU General Public License Version 2 (GPLv2)
Perl 5.10.1 and related modules	Perl Artistic License
Snort 2.9.5	GNU General Public License Version 2 (GPLv2)

A complete list of all third-party open source code products and the full text of all applicable license agreements that are distributed with the Sourcefire products can be obtained by logging into the product command line and viewing the following:

```
/usr/share/license-files
```

If you would like to receive the source code to any of the third-party open source code products used in combination with the Sourcefire products, submit a request to the Sourcefire Support Site.

APPENDIX F

END USER LICENSE AGREEMENT

SOURCEFIRE NETWORK SECURITY PRODUCTS

IMPORTANT: PLEASE READ THIS END USER LICENSE AGREEMENT CAREFULLY.

THIS END USER LICENSE AGREEMENT IS A LEGAL AGREEMENT BETWEEN YOU AND SOURCEFIRE, INC. OR ITS DESIGNATED SUBSIDIARY LICENSING THE LICENSED MATERIALS TO YOU HEREUNDER INSTEAD OF SOURCEFIRE, INC. ("SOURCEFIRE"). THE TERMS AND CONDITIONS UNDER WHICH YOU MAY USE THE LICENSED MATERIALS ARE SET FORTH IN THIS END USER LICENSE AGREEMENT ("EULA"), IN ADDITION TO ANY OTHER TERMS AS MAY BE SET FORTH IN ANY SUPPLEMENTAL LICENSE AGREEMENT(S) WHICH MAY ACCOMPANY ANY SOURCEFIRE PRODUCTS (TOGETHER WITH THE EULA, THE "AGREEMENT"). BY DOWNLOADING, INSTALLING AND USING ANY OF THE SOURCEFIRE PRODUCTS, YOU ARE BINDING YOURSELF AND THE BUSINESS ENTITY THAT YOU REPRESENT (COLLECTIVELY, "YOU") TO THIS AGREEMENT AND AGREEING THAT THIS AGREEMENT WITH SOURCEFIRE IS ENFORCEABLE LIKE ANY WRITTEN CONTRACT SIGNED BY YOU.

IF YOU DO NOT AGREE TO ALL OF THE TERMS AND CONDITIONS CONTAINED IN THIS AGREEMENT, THEN SOURCEFIRE IS UNWILLING TO LICENSE THE LICENSED MATERIALS TO YOU, IN WHICH CASE YOU MAY NOT DOWNLOAD, INSTALL OR USE ANY OF THE LICENSED MATERIALS.

IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT DO NOT INITIATE USE OF THE PRODUCT. BY SELECTING "I ACCEPT," "OK," "CONTINUE," "YES," "NEXT" OR BY INSTALLING OR USING THE LICENSED MATERIALS IN ANY WAY, YOU ARE INDICATING YOUR COMPLETE UNDERSTANDING AND ACCEPTANCE OF THE TERMS AND

CONDITIONS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL OR USE THE PRODUCT.

If You are located outside of the United States, then Sourcefire International GmbH, a subsidiary located in Switzerland, shall be a party to this Agreement with You and the party licensing the Licensed Materials to You hereunder. This Agreement governs Your access and use of the Sourcefire Products, except to the extent there is a separate written agreement signed by both You and Sourcefire that expressly states that it governs Your use of the Sourcefire Products. In the event of a conflict between the provisions of such a written agreement and this Agreement, the order of precedence shall be (1) the separate signed agreement, and (2) this Agreement.

1. DEFINITIONS

The following capitalized terms shall have the following meanings in this EULA:

- 1.1. "Appliance" means any Sourcefire-branded network security appliance made available to You, consisting of Hardware and pre-installed Sourcefire Software and/or other Licensed Materials.
- 1.2. "Documentation" means written information contained in user manuals and technical specifications pertaining to the use of the Sourcefire Products made available by Sourcefire with the Sourcefire Products in any manner (including on CD-ROM, or on-line).
- 1.3. "Hardware" means the hardware components of any Appliance on which Sourcefire Software is installed and runs.
- 1.4. "Laws" means, collectively, all international and national laws, treaties, statutes, ordinances, regulations and other types of government authority.
- 1.5. "Licensed Materials" means any Sourcefire Software, Documentation and Subscription Data licensed by Sourcefire to You hereunder.
- 1.6. "Party" or "Parties" means, individually each party hereto, and collectively all the parties to this Agreement.
- 1.7. "Products" means the Sourcefire Products and/or the Third Party Products.
- 1.8. "Reseller" means an authorized reseller or distributor of Sourcefire.
- 1.9 "Sourcefire Products" means the Appliance(s) and/or Licensed Materials.
- 1.10. "Sourcefire Software" means the machine-readable computer software programs licensed by Sourcefire to You hereunder including any software provided to You for use on a subscription, term or software-as-a-services (SaaS) basis, and all Updates to any of the foregoing.
- 1.11. "Subscription Data" means that data made available to You by Sourcefire for use with the Sourcefire Products including, but not limited to, URL data and IP address blacklists.
- 1.12 "Third Party Products" means any products or other materials made available to You for use with Sourcefire Products and which are not Sourcefire Products.

1.13. "Updates" means with respect to Licensed Materials any Sourcefire-approved periodic patches, bug-fixes, work-arounds, error corrections, enhancements, rules updates, vulnerability database updates, security enhancement updates and additions and other modifications thereto, or revised versions thereof, which may be made available from time to time.

Unless otherwise defined herein, the capitalized terms used in this EULA shall be defined in the context in which they are used.

2. YOUR PAYMENT OBLIGATIONS

In consideration for Your purchase of an Appliance and Your license to use the Licensed Materials, You agree to pay all amounts due or incurred by You, including all applicable taxes, as are specified in an invoice provided by Sourcefire or a Reseller, as applicable.

3. LICENSE GRANT

Subject to the terms and conditions of this Agreement, Sourcefire grants to You a limited, non-exclusive and non-transferable license to download, install and use for Your internal operations the Licensed Materials for which You have paid the required fees to Sourcefire or a Reseller, as applicable. Such Licensed Materials may be delivered to You pre-installed on an Appliance, made available to You separately via download by Sourcefire or otherwise made available on a subscription, term or software-as-a-service (SaaS) basis. In order to use the Products, You may be required to input a registration number, product authorization key or otherwise register Your Products online at Sourcefire's designated website to obtain the necessary license key or license file. You shall own the Appliance that You purchase and the magnetic or other physical media upon which the Licensed Materials are originally or subsequently recorded or fixed, but Sourcefire and Sourcefire's licensors, as applicable, retain all title, copyright and other intellectual proprietary rights in, and ownership of, the Licensed Materials regardless of the media upon which the original or any copy may be recorded or fixed. You may make one (1) copy of the Licensed Materials solely for internal backup purposes. Sourcefire and its licensors expressly reserve any rights in Licensed Materials not granted herein.

4. SCOPE OF USE

If You purchased an Appliance, You may only use the Licensed Materials included on that Appliance for use on such Appliance. If Sourcefire Products are made available to You for use without an Appliance on a "virtual" basis, Your use of such Sourcefire Products may not exceed the applicable number of licenses purchased and other limitations associated with the fees paid or payable by You for such use. If You purchased a license to use the Licensed Materials on a subscription or term basis, You may not deploy or use such Licensed Materials in a manner that exceeds the term of subscription, the permitted number of users, hosts or endpoints, or other subscription or term limitations associated with the applicable fees paid or payable by You.

5. LICENSE RESTRICTIONS

You agree not to directly or indirectly: (i) sell, lease, rent, distribute, sublicense, assign or transfer any of the Licensed Materials; (ii) reverse engineer, decompile, disassemble, decrypt or otherwise attempt to determine the source code of any of the Licensed Materials, except to the limited extent permitted by law; (iii) modify, make error corrections to or create derivative works based on the Licensed Materials; (iv) use any Licensed Materials for the benefit of any third parties (e.g., in an ASP, SaaS, outsourcing or service bureau relationship) or in any way other than in its intended manner, except as otherwise permitted by Sourcefire; (v) remove, alter or obscure any proprietary or copyright notice, labels, or marks on the Hardware or within the Licensed Materials; (vi) disable or circumvent any access control or related security measure, process or procedure established with respect to the Appliance or any Licensed Materials or any other part thereof; (vii) create Internet “links” to the Subscription Data or “frame” or “mirror” the Subscription Data on any other server or wireless or Internet-based device; or (viii) utilize the Subscription Data in order to: (1) build a competitive product or service; (2) build a product using similar ideas, features, functions or graphics; (3) copy any ideas, features, functions or graphics; or (4) aggregate subscriptions to the Subscription Data, either by sublicensing or by rebranding of the Subscription.

You may use the Sourcefire Products to conduct internal performance and benchmark testing, the results of which only You may publish or publicly disseminate, provided that (i) Sourcefire has reviewed and approved the methodology, assumptions and parameters of Your testing, (ii) You publish a full description of the test environment and methods, assumptions and parameters used in the testing, and (iii) You do not publish false, deceptive or misleading statements relating to the test or Sourcefire Products. Please contact a Sourcefire technical support representative regarding approved testing methodology, assumptions and parameters.

You are responsible for all use of the Products obtained by You and for compliance with this Agreement; any breach of this Agreement by You or other user in connection with the use of those Products obtained by You shall be deemed to have been made by You.

6. INTELLECTUAL PROPERTY; TITLE

This Agreement does not transfer to You any title or any ownership right or interest in any Licensed Materials or in any other intellectual property rights of Sourcefire or Sourcefire’s licensors. You acknowledge that the Appliance(s) and the Licensed Materials contain, embody and are based upon patented or patentable inventions, trade secrets, copyrights and other intellectual property rights owned by Sourcefire and its licensors. If You purchased an Appliance, title and risk of loss to each Appliance transfers to You when the Appliance is delivered to Sourcefire’s designated carrier for shipment; Products are shipped FOB Sourcefire’s designated shipping facility. If you purchased an Appliance from a Reseller, the terms of such purchase are between You and such Reseller. If You purchased an Appliance directly from Sourcefire, the terms of such purchase are

as set forth in the Sourcefire sales order submitted by You and accepted by Sourcefire. In all instances, Licensed Materials are licensed to You pursuant to this Agreement and not sold to You.

7. TECHNICAL SUPPORT

You may purchase technical support for Sourcefire Products by separately enrolling in Sourcefire's customer support plan (the "Support Plan") and paying Sourcefire or a Reseller the then-applicable customer support fee. Your rights and Sourcefire's obligations under the Support Plan are set forth in the Support Plan terms and conditions, a current copy of which is located at www.sourcefire.com/customer-support. All Updates received by You pursuant to the Support Plan shall be governed by, and licensed to You under, this Agreement.

8. CONFIDENTIALITY

As used herein, "Confidential Information" means any non-public technical or business information of either Party, including without limitation, the terms and conditions of this Agreement, any information relating to Sourcefire's techniques, algorithms, software, know-how and current or future product designs, financial information, procurement requirements, and manufacturing or business forecasts. Confidential Information does not include information that (i) is or becomes generally known to the public through no fault or breach of this Agreement by the receiving party; (ii) the receiving party can demonstrate by written evidence was rightfully in the receiving party's possession at the time of disclosure, without an obligation of confidentiality; (iii) is independently developed by the receiving party without use of or access to the disclosing party's Confidential Information or otherwise in breach of this Agreement; (iv) the receiving party rightfully obtains from a third party not under a duty of confidentiality and without restriction on use or disclosure; or (v) is required to be disclosed pursuant to, or by, any Laws, court order or other legal process to do so, provided that the receiving party shall, promptly upon learning that such disclosure is required, give written notice of such disclosure to the disclosing party. The party receiving Confidential Information will employ all reasonable measures to maintain the confidentiality of such Confidential Information, but in no event shall such measures be less than the measures the receiving party employs to protect its own confidential information. The party receiving the Confidential Information will limit the disclosure of the other party's Confidential Information to its employees and contractors with a bona fide need to access such Confidential Information in order to exercise its rights and obligations under this Agreement; provided that, all such employees and contractors are bound by a written non-disclosure agreement that contains restrictions at least as protective as those set forth herein. The Parties agree that the party disclosing Confidential Information will suffer irreparable harm in the event that the receiving party breaches any obligation under this Section 8 and that monetary damages will be inadequate to compensate the non-breaching party for such breach. In the event of a breach, or threatened breach, of any of the provisions of this Section 8, the non-breaching party, in addition to and not in limitation of any other rights,

remedies or damages available to it at law or in equity, shall be entitled to seek a temporary restraining order, preliminary injunction and/or permanent injunction in order to prevent or to restrain any such breach.

9. INSTALLATION

You represent, warrant and covenant that You are solely responsible for the proper installation, configuration and management of the Appliance on which the Licensed Materials will be installed, as well as the installation of any separately provided Licensed Materials. You further understand and hereby acknowledge that the failure to properly configure and manage an Appliance, and the failure to properly install any separately provided Licensed Materials, may adversely affect the performance of the Appliance and the Licensed Materials. You represent, warrant and covenant that You will adhere to the recommended minimum requirements specified in the Documentation. Sourcefire shall have no obligation under this Agreement to the extent an Appliance or any separately provided Licensed Materials fails to substantially perform the functions described in the Documentation, in whole or in part, because (i) You fail to meet specified minimum requirements; (ii) Your separate hardware fails to perform properly; (iii) You improperly configured an Appliance; or (iv) the Licensed Materials had been improperly installed.

10. WARRANTY AND DISCLAIMER

Sourcefire warrants that, for a period of ninety (90) days from the date of initial shipment of the Appliance or, in the case of Sourcefire Software separately provided to You, the date the Sourcefire Software is made available to You for download or delivered on a fixed media (as the case may be, the "Software Warranty Period"), the unmodified Sourcefire Software will, under normal use, substantially perform the functions described in its Documentation. Sourcefire also warrants that for a period of one (1) year from the date of initial shipment of a new Appliance (the "Hardware Warranty Period") that the unmodified Hardware comprising such Appliance will, under normal use, be free of substantial defects in materials and workmanship. Neither of the aforementioned warranties apply if the Sourcefire Software or Appliance (i) has been altered, except by Sourcefire or its authorized representative; (ii) has not been installed, operated, repaired or maintained in accordance with the Documentation and/or instructions supplied by Sourcefire, (iii) has been subjected to abnormal physical or electrical stress, abnormal environmental conditions, misuse, negligence or accident by You; or (iv) is licensed for beta, evaluation, testing or demonstration purposes. If a court of competent jurisdiction determines that the statutory warranty periods of such jurisdiction apply rather than the Software Warranty Period and Hardware Warranty Periods referenced above, then such statutory warranty periods will control only in the event of a conflict with the terms of this Section 10.

EXCEPT AS EXPRESSLY WARRANTED IN THIS SECTION 10, THE SOURCEFIRE PRODUCTS (INCLUDING, ANY EVALUATION AND BETA PRODUCTS), AND ANY OTHER DOCUMENTATION, MATERIALS AND/OR DATA PROVIDED BY SOURCEFIRE ARE PROVIDED "AS IS" AND "WITH ALL FAULTS," AND

SOURCEFIRE EXPRESSLY DISCLAIMS ALL OTHER WARRANTIES OF ANY KIND OR NATURE, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF OPERABILITY, CONDITION, TITLE, NON-INFRINGEMENT, NON-INTERFERENCE, QUIET ENJOYMENT, VALUE, ACCURACY OF DATA, OR QUALITY, AS WELL AS ANY WARRANTIES OF MERCHANTABILITY, SYSTEM INTEGRATION, WORKMANSHIP, SUITABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR THE ABSENCE OF ANY DEFECTS THEREIN, WHETHER LATENT OR PATENT.

THE SOURCEFIRE PRODUCTS ARE NOT FAULT-TOLERANT AND ARE NOT DESIGNED OR INTENDED FOR USE IN ANY HAZARDOUS ENVIRONMENT REQUIRING FAIL-SAFE PERFORMANCE OR OPERATION. SOURCEFIRE PRODUCTS ARE NOT FOR USE IN THE OPERATION OF AIRCRAFT NAVIGATION, NUCLEAR FACILITIES, WEAPONS SYSTEMS, DIRECT OR INDIRECT LIFE-SUPPORT SYSTEMS, AIR TRAFFIC CONTROL SYSTEMS, OR ANY APPLICATION OR INSTALLATION WHERE FAILURE COULD RESULT IN DEATH, PHYSICAL INJURY OR PROPERTY DAMAGE.

NO WARRANTY IS MADE BY SOURCEFIRE ON THE BASIS OF TRADE USAGE, COURSE OF DEALING OR COURSE OF TRADE. SOURCEFIRE DOES NOT WARRANT THAT THE APPLIANCE, THE LICENSED MATERIALS OR ANY OTHER INFORMATION, MATERIALS, DOCUMENTATION OR TECHNOLOGY PROVIDED UNDER THIS AGREEMENT WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION THEREOF WILL BE UNINTERRUPTED OR ERROR-FREE, OR THAT ALL ERRORS WILL BE CORRECTED. YOU ACKNOWLEDGE THAT SOURCEFIRE'S OBLIGATIONS UNDER THIS AGREEMENT ARE FOR YOUR BENEFIT ONLY. NOTWITHSTANDING ANYTHING TO THE CONTRARY IN THIS AGREEMENT, ANY THIRD PARTY PRODUCTS ARE PROVIDED "AS IS" WITHOUT ANY WARRANTY WHATSOEVER.

Sourcefire's sole obligation and liability, and Your sole and exclusive remedy under the warranties set forth in Section 10 shall be for Sourcefire to use commercially reasonable efforts to remedy the problem, or to replace the defective Hardware and/or the Sourcefire Software, if Sourcefire is notified in writing of all warranty problems during the applicable warranty period.

11. LIMITATION OF LIABILITY

IN NO EVENT WILL SOURCEFIRE'S OR ANY OF ITS SUBSIDIARIES' OR AFFILIATES' AGGREGATE LIABILITY (INCLUDING, BUT NOT LIMITED TO, LIABILITY FOR NEGLIGENCE, STRICT LIABILITY, BREACH OF CONTRACT, MISREPRESENTATION AND OTHER CONTRACT OR TORT CLAIMS) ARISING FROM OR RELATED TO THIS AGREEMENT, OR THE USE OF THE PRODUCTS, EXCEED THE AMOUNT OF FEES YOU PAID TO SOURCEFIRE OR ITS RESELLER, AS APPLICABLE, FOR THE PRODUCTS THAT GAVE RISE TO SUCH LIABILITY. UNDER NO CIRCUMSTANCES SHALL SOURCEFIRE OR ANY OF ITS SUBSIDIARIES, AFFILIATES, SUPPLIERS OR LICENSORS BE LIABLE FOR ANY OF THE FOLLOWING: (I) THIRD PARTY CLAIMS, EXCEPT AS SET FORTH IN SECTION 13; (II) LOSS OR DAMAGE TO ANY SYSTEMS, RECORDS OR DATA; (III) INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, PUNITIVE,

RELIANCE, OR COVER DAMAGES (INCLUDING LOST PROFITS AND LOST SAVINGS); OR (IV) DAMAGES ARISING OUT OF ANY THIRD PARTY PRODUCTS, IN EACH CASE EVEN IF SOURCEFIRE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. YOU ARE SOLELY RESPONSIBLE AND LIABLE FOR VERIFYING THE SECURITY, ACCURACY AND ADEQUACY OF ANY OUTPUT FROM ANY PRODUCTS, AND FOR ANY RELIANCE THEREON. THE LIMITATIONS OF LIABILITY IN THIS SECTION 11 ARE INTENDED TO APPLY WITHOUT REGARD TO WHETHER OTHER PROVISIONS OF THIS AGREEMENT HAVE BEEN BREACHED OR HAVE PROVEN INEFFECTIVE.

12. ESSENTIAL BASIS

The disclaimers, exclusions and limitations of liability set forth in this Agreement form an essential basis of the bargain between the Parties, and, absent any of such disclaimers, exclusions or limitations of liability, the provisions of this Agreement, including, without limitation, the economic terms, would be substantially different.

13. INFRINGEMENT OBLIGATIONS

13.1. Sourcefire will defend You from any unaffiliated third party claim that Your use of the Sourcefire Software as provided by Sourcefire to You under this Agreement, when used within the scope of this Agreement, infringes any unaffiliated third party's U.S. copyright ("Claim"). Sourcefire's obligations to You under this Section 13 are limited solely to paying (i) counsel hired by Sourcefire to defend the Claim; (ii) the reasonable and verifiable out-of-pocket costs incurred directly by You in connection with defending the Claim and/or assisting Sourcefire in the defense thereof; and (iii) subject to Section 11 herein, any direct damages finally awarded to such third party by a court of competent jurisdiction (after any appeals) or any settlement of the Claim to which Sourcefire consents in writing. Sourcefire's obligations under this Section 13 are expressly contingent upon: (x) You giving prompt written notice to Sourcefire of any such Claim; (y) You allowing Sourcefire exclusive control of the defense and any related settlement of any such Claim; and (z) You furnishing Sourcefire with reasonable assistance in connection with the Claim without prejudicing Sourcefire in any manner. Subject to the foregoing conditions, nothing in this Agreement shall prohibit You from hiring separate counsel, at Your own expense.

13.2. If Your use of the Products hereunder is, or in Sourcefire's opinion is likely to be, enjoined due to the type of Claim specified in Section 13.1, then Sourcefire may, at its sole option and expense but without obligation to do so: (i) procure for You the right to continue to use the Products under the terms of this Agreement; (ii) replace the Products with a functional equivalent; (iii) modify the Products so that they become non-infringing (including disabling the challenged functionality), provided the modified Products remain substantially equivalent in function to the enjoined Products; or (iv) repurchase the affected Products less depreciation at the rate of twenty-five percent (25%) per year, or pro rata for the part of the year, from the date of payment to the date of removal of the Products, and terminate the Agreement with respect to those Products. Further, if as a result of a Claim a

court of competent jurisdiction issues a final injunction (which has not been appealed) against Your use of any part of the Products, then Sourcefire will, at its sole option, perform one of the remedy options listed in this Section 13.2. In either case, if Sourcefire selects option (ii), (iii) or (iv) listed in this Section 13.2, You shall immediately refrain from use of the allegedly infringing Products.

13.3. Sourcefire shall have no indemnification obligation or liability for any Claim to the extent that it arises out of or relates to: (i) Your use of the Products after Sourcefire notifies You to discontinue use due to a Claim; (ii) the combination of the Sourcefire Products with a non-Sourcefire application, product, data or business process; (iii) damages attributable to a non-Sourcefire application, product, data or business process; (iv) modifications to the Products made other than by Sourcefire; (v) changes made by Sourcefire on behalf of You; (vi) continued use of the Products for which Sourcefire has provided You with modifications or substitute Products if use of such modifications or substitute Products would have prevented the Claim; or (vii) use of the Products in a manner prohibited under this Agreement.

13.4. THE PROVISIONS OF THIS SECTION 13 SET FORTH SOURCEFIRE'S SOLE AND EXCLUSIVE OBLIGATIONS, AND YOUR SOLE AND EXCLUSIVE REMEDIES, WITH RESPECT TO INFRINGEMENT, VIOLATION OR MISAPPROPRIATION OF INTELLECTUAL PROPERTY RIGHTS OF ANY KIND. IN NO EVENT SHALL SOURCEFIRE'S LIABILITY TO YOU UNDER SECTION 13 EXCEED THE AMOUNT OF THE FEES PAID BY YOU FOR THE SOURCEFIRE PRODUCT THAT IS THE SUBJECT OF SUCH CLAIM.

14. VERIFICATION

You agree that Sourcefire or its designee shall have the right to periodically conduct on-site audits of Your use of the Products for the purpose of verifying that You are in compliance with Your obligations under this Agreement and have paid all applicable fees. These audits will be conducted during regular business hours, and Sourcefire will make reasonable efforts to minimize interference with Your regular business activities. Alternatively, Sourcefire may request that You complete a self-audit questionnaire in a form provided by Sourcefire. If an audit or such questionnaire reveals unlicensed use of the Products, You agree to promptly order and pay for sufficient licenses to permit all usage disclosed.

15. EXPORT; RE-EXPORT

The Products are subject to export controls under the Laws of the United States and other countries. You shall comply with all such Laws governing export, re-export, transfer and use of the Products and will obtain all required U.S. and local authorizations, permits and licenses. Sourcefire assumes no responsibility or liability for Your failure to obtain such necessary authorizations, permits and licenses. Information regarding U.S. export laws can be found at www.bis.doc.gov. You agree not to use or transfer the Products for any use relating to the operation of nuclear facilities, chemical or biological weapons, or

missile technology, unless authorized by the U.S. Government by regulation or specific written license.

16. U.S. GOVERNMENT END USERS

The Licensed Materials, information and data provided under this Agreement are prepared entirely at private expense and are "Commercial Items" as that term is defined in 48 C.F.R. 2.101. If you are an agency, department, or other entity of the United States Government, or funded in whole or in part by the United States Government, then your use, duplication, reproduction, release, modification, disclosure or transfer of this commercial product and data, is restricted in accordance with 48 C.F.R. §12.211, 48 C.F.R. §12.212, 48 C.F.R. §227.7102-2, and 48 C.F.R. §227.7202, as applicable. Consistent with 48 C.F.R. §12.211, 48 C.F.R. §12.212, 48 C.F.R. §227.7102-1 through 48 C.F.R. §227.7102-3, and 48 C.F.R. §§227.7202-1 through 227.7202-4, as applicable, this commercial product and data are licensed to U.S. Government end users (i) only as Commercial Items, and (ii) with only those rights as are granted to all other users pursuant to the Sourcefire's standard end user license agreement. In case of conflict between any of the FAR and DFARS provisions listed herein and this Agreement, the construction that provides greater limitations on the U.S. Government's rights shall control. For purpose of any public disclosure provision under any federal, state or local law, it is agreed that this commercial product and data are a trade secret and proprietary commercial products and not subject to disclosure.

17. FREE SOFTWARE

You acknowledge and agree that while certain open source code Third Party Products are made available to You hereunder for free for use in combination with the Sourcefire Products, the terms and conditions under which such Third Party Products are being made available to You are as set forth in their respective third party agreements (the "Third Party Agreements"), and that this Agreement in no way supplements or detracts from any term or condition of such Third Party Agreements. Sourcefire is not giving any warranties for these Third Party Products and Your use of these Third Party Products will be subject solely to such Third Party Agreements. A listing of these Third Party Products, including the applicable Third Party Agreements and other applicable disclosures, is available in the Documentation. You may obtain the source code to such open source code software in accordance with the directions set forth in the Documentation.

18. EVALUATION PRODUCTS

If You have been provided Products on an evaluation-only basis or beta-release basis (each, "Evaluation Products") to evaluate their suitability for purchase and/or licensing on a for-fee basis (as the case may be, for "Evaluation"), You acknowledge and agree that the evaluation license key(s) for these Evaluation Products will be set with a set expiration date (the "Expiration Date"), pursuant to which upon activation of the Evaluation Products, You may use the Evaluation Products through the Expiration Date (the "Evaluation Period") solely for their

Evaluation. All Evaluation Products are provided to You “AS IS” without warranty or any kind, whether express, implied, statutory, or otherwise, and the limited warranties referenced in Section 10 and the indemnification obligations referenced in Section 13 above will not be applicable to Your use of the Evaluation Products. Sourcefire bears no liability for any damages resulting from use (or attempted use) of the Evaluation Products.

19. COLLECTION OF DATA

Sourcefire hereby informs You that the Products use data collection technology to collect and analyze certain information about Your network and endpoints including, but not limited to, the IP addresses of Your endpoints and the metadata of certain executable files in order to (i) identify malware on Your network and endpoints; (ii) provide support and related services to You; and (iii) improve Sourcefire's products. You do have the ability to configure the Products to limit some of the data that can be collected. You grant Sourcefire a perpetual right and license to use the information and data made available by You via the Products in order to attempt to prevent malware from running on Your network and endpoints, to conduct related analysis and provide support and for product improvement purposes. By accepting this Agreement, You (x) acknowledge and agree that the technology included in the Products can collect traffic and data from Your network and endpoints to detect malware and conduct related analysis; (y) agree to upload from Your network and endpoints certain metadata and other required information for the purpose of being scanned by the remote cloud-based servers operated by Sourcefire; and (z) covenant that You have the right to provide Sourcefire all such information and data.

Sourcefire may engage other companies and individuals to perform functions on its behalf, such as payment processing, order fulfillment, marketing programs and customer service so Sourcefire may share such information with such subcontractors in order to perform these functions, but such subcontractors may not use Your personal information for other purposes, unless You agree.

20. GOVERNING LAW

If You are located in the United States, then Sourcefire, Inc. is the Sourcefire entity that is the party hereto and this Agreement shall be governed in all respects by the laws of the State of New York, USA, without regard to choice-of-law rules or principles. If You are located outside of the United States, then Sourcefire International GmbH is the Sourcefire entity that is the party hereto and this Agreement shall be governed in all respects by the laws of Switzerland, without regard to choice-of-law rules or principles. In either case, You expressly agree with Sourcefire that the U.N. Convention on Contracts for the International Sale of Goods shall not govern this Agreement. You represent that You understand, and You hereby agree to comply with, all requirements of the U.S. Foreign Corrupt Practices Act and all other applicable Laws. Except for instances where equitable relief is permitted under this Agreement, any and all claims, disputes, or controversies arising under, out of, or in connection with this Agreement or the breach thereof (each, a “Dispute”) shall be submitted to the

chief operating officer (or equivalent) of each Party (or their designee) for a good faith attempt to resolve the Dispute. The position of each Party shall be submitted, and the individuals promptly thereafter shall meet at a neutral site in an attempt to resolve such Dispute.

21. ASSIGNMENT

You may not assign or otherwise transfer this Agreement without Sourcefire's prior written consent. Notwithstanding the foregoing, You may assign this Agreement if a majority of Your outstanding voting capital stock is sold to a third party, or if You sell all or substantially all of Your assets or if You otherwise undergo a change of control, provided, that, in such instance such assignment will not become effective until You provide Sourcefire written notice of such event. Sourcefire may assign or transfer this Agreement, in whole or in part, at any time in its sole discretion without Your consent. This Agreement shall be binding upon and inure to the benefit of the Parties' successors and permitted assigns.

22. TERM; TERMINATION

This Agreement will continue in effect indefinitely, subject to the right of either Party to terminate as provided below. Either Party may terminate this Agreement if the other does not comply with any of its terms, if the one who is not complying is given written notice and reasonable time to comply. Sourcefire may terminate Your licenses and other rights herein, immediately by providing notice, if You breach or fail to comply with any of the terms and conditions of this Agreement. You agree that, upon such termination, You will cease using the Licensed Materials and either destroy or return all copies thereof.

23. GENERAL

This Agreement and the terms and conditions referenced herein are the Parties' complete agreement regarding its subject matter, superseding any prior oral or written communications. Under no circumstances will the terms of any purchase order issued by You control or otherwise negate the terms set forth in this Agreement. If any provision of this Agreement is held to be illegal or unenforceable for any reason, then such provision shall be deemed to be restated to be enforceable to the maximum extent permissible under law, and the remainder of this Agreement shall remain in full force and effect. Neither Party shall be liable for any delay or failure due to a force majeure event and other causes beyond its reasonable control, provided, however, this provision shall not apply to Your payment obligations. Any notices under this Agreement to Sourcefire will be personally delivered or sent by certified or registered mail, return receipt requested, or by nationally recognized overnight express courier, to 9770 Patuxent Woods Drive, Columbia, Maryland U.S.A. 21046 if to Sourcefire, Inc., and to Rheinweg 1c, CH-8200 Schaffhausen, Switzerland if to Sourcefire International GmbH, or such other address as Sourcefire may specify in writing. Such notices will be effective upon receipt, which may be shown by confirmation of delivery. All notices to Sourcefire shall be sent to the attention of General

Counsel (unless otherwise specified by Sourcefire). Amendments or changes to this Agreement must be in mutually executed writings to be effective. Sections 1-2, 5-6, 8-12 and 14-23, including all warranty disclaimers and use restrictions, shall survive the termination or expiration of this Agreement. The Parties are independent contractors for all purposes under this Agreement.

-- End of Agreement --

Glossary

7000 Series	A group of Series 3 Sourcefire managed devices . The devices in this series include the 70xx Family (the 3D7010/7020/7030 models) and the 71xx Family (3D7110/7120/3D7115/3D7125 models).
8000 Series	A group of Series 3 Sourcefire managed devices . The devices in this series include the 81xx Family (the 3D8120/8130/8140 models) and the 82xx Family (the 3D8250/8260/8270/8290 models). 8000 Series devices are generally more powerful than 7000 Series devices.
access control	A feature of the Sourcefire 3D System that allows you to specify, inspect, and log the traffic that traverses your network. Access control includes the intrusion detection and prevention , file control , and advanced malware protection features, and also determines the traffic you can inspect with the discovery feature.
access control policy	A policy that you apply to managed devices to perform access control on the network traffic monitored by those devices. An access control policy may include multiple access control rules ; it also specifies a default action , which determines the handling and logging of traffic that does not meet the criteria of any of those rules. An access control policy can also specify HTTP response page , Security Intelligence , and other advanced settings.
access control rule	A set of conditions the Sourcefire 3D System uses to examine your monitored network traffic and which allows you to achieve granular access control . Access control rules, which populate an access control policy , may perform simple IP address matching, or may characterize complex connections involving different users, applications , ports, and URLs. The access control rule action determines how the system handles traffic that meets the rule's conditions. Other rule settings determine how (and whether) the connection is logged, and whether an intrusion policy or file policy inspects matching traffic.

access control rule action	A setting that determines how the system handles network traffic that meets the conditions of an access control rule . You can <i>block</i> matching traffic (with or without resetting the connection); for HTTP traffic you can provide users with the option to bypass the block. You can also <i>trust</i> traffic to pass without further inspection, <i>allow</i> matching traffic, which optionally can be inspected with an intrusion policy and file policy , or continue to <i>monitor</i> the traffic with additional access control rules.
access-controlled user	A user whose network use you can control using access control . You specify the LDAP groups that access-controlled users must belong to when you configure a connection between a Microsoft Active Directory server and the Defense Center . When the Sourcefire User Agent reports logins by access-controlled users, those users are associated with IP addresses, which in turn allows access control rules with user conditions to trigger. Compare with non-access-controlled user .
access list	A list of IP addresses, configured in the system policy , that represents the hosts that can access an appliance . By default, anyone can access the web interface of an appliance using port 443 (HTTPS), as well as the command line using port 22 (SSH). You can also add SNMP access using port 161.
active detection	The discovery of host , application , and user information using active sources. Active sources include scanners such as Nmap , user input to the system's web interface, or host input to the network map using the command line or third-party application API calls. Compare with passive detection .
adaptive profile	An intrusion policy profile that uses discovery data to determine the operating system for the target host of a packet. Profiles within an intrusion policy then automatically adapt to cause preprocessors to defragment IP packets and reassemble streams in the same way as the operating system on the target host, and to cause Snort to analyze the data in the same format as that used by the destination host.
advanced malware protection	Abbreviated AMP, the Sourcefire 3D System's network-based malware detection and malware blocking feature. Compare this functionality with FireAMP , Sourcefire's endpoint-based AMP tool that requires a FireAMP subscription .
advanced setting	A preprocessor or other intrusion policy feature that requires specific expertise to configure. Advanced settings typically require little or no modification and are not common to every deployment.
alert	A notification that the system has generated a specific event . You can alert based on intrusion events (including their impact flags), discovery events , network-based malware events , correlation policy violations, health status changes, and connections logged by specific access control rules . In most cases, you can alert via email, syslog, or SNMP trap.
alert response	A set of configurations that allows the system to send an alert via email, syslog, or SNMP trap. You can use a single alert response to alert you to multiple types of

[events](#).

appliance	A Defense Center or managed device . An appliance can be physical or virtual.
appliance statistics	Information you can obtain about an appliance , including uptime, system memory usage, load average, disk usage, a summary of system processes, and, on the Defense Center , information about data correlator processes.
application	A detected network asset, communications method, or HTTP content against which you can write access control rules . The system detects three types of application: application protocol , client application , and web application .
application business relevance	The likelihood that an application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from very low to very high.
application category	A general classification for an application that describes its most essential function. Each application belongs to at least one category.
application control	A feature that, as part of access control , allows you to specify which application traffic can traverse your network.
application detector	A tool that the system uses to identify applications on your network. Application detectors identify applications using ASCII or hexadecimal patterns in the packet headers, the port that the traffic uses, or both. Sourcefire may deliver additional detectors via system update, vulnerability database update, or the import/export feature. You can also create your own application protocol detectors.
application filter	One or more applications grouped according to criteria associated with the application risk , business relevance , type, categories, and tags. You can use these filters to constrain access control rules , searches, some dashboard widgets , and reports. You create application filters in the object manager or on the fly in the access control rule editor.
application protocol	A type of application that represents application protocol traffic detected during communications between server and client applications on hosts; for example, SSH or HTTP.
application risk	The likelihood that an application's use may violate your organization's security policy . An application's risk can range from very low to very high.
application tag	Information about an application that is not covered by its application category . For example, video streaming web applications often are tagged "high bandwidth" and "displays ads." An application may have any number of tags, including none.
application type	Whether an application is an application protocol , client application , or web application .

apply	The action you take to have a policy , or changes to that policy, take effect. You apply most policies from the Defense Center to its managed devices ; however, you activate and deactivate correlation policies because they do not involve changes to the configuration of managed devices.
audit event	An event that describes a specific Sourcefire 3D System user interaction with the web interface. Each audit event contains a time stamp, the user name of the user whose action generated the event, a source IP address, and text describing the event. Audit events are recorded in the audit log .
audit log	A record of user interactions with the web interface. The audit log comprises audit events .
authentication object	A collection of settings that allows you to connect to an external authentication server for one of two reasons. One type of authentication object enables external authentication (RADIUS or LDAP) to the Sourcefire 3D System's web interface. Another type specifies the LDAP users and groups you can use in access control rules , and allows you to retrieve metadata for certain users whose activity was detected in network traffic or by the Sourcefire User Agent .
automatic application bypass (AAB)	An advanced device setting that limits the time allowed to process packets through an interface and allows packets to bypass processing if the time is exceeded. Sourcefire recommends use of this feature for inline deployments in non-production environments.
banner	See server banner .
base policy	A selectable set of configurations that may be any default intrusion policy provided by Sourcefire, or a custom policy.
base policy layer	A built-in layer in an intrusion policy comprised of default settings. The base policy selected for the intrusion policy determines the settings in the base policy layer.
blacklist	See health monitor blacklist or Security Intelligence blacklist .
bookmark	A saved link to a specific location and time in an event analysis. Bookmarks retain information about the workflow you are using, the part of the workflow you are viewing, the page number within the workflow you are viewing, the time window you selected, and any columns you disabled, as well as any constraints you imposed.
built-in layer	A read-only layer in an intrusion policy . An intrusion policy always includes a built-in base policy layer and, optionally, can include a built-in FireSIGHT Recommendations layer .
business relevance	The likelihood that an application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from very low to very high.

bypass mode	A characteristic of an inline set that allows traffic to continue flowing if the sensing interfaces in the set fail for any reason.
category	See application category , file category , or URL category .
certificate authority	The certificate issuer used to create server or user certificates. Server and user certificates provide an additional confirmation of a server or a user identity.
certificate revocation list (CRL)	A list of certificates revoked by the certificate authority that issued the user certificates for your appliance. This allows you to restrict access to the Sourcefire 3D System web interface using client browser certificate checking. If the user selects a certificate that is listed in the CRL as a revoked certificate, the browser cannot load the web interface.
change reconciliation report	A detailed report of all system changes in the last 24 hours, based on snapshots taken whenever a new configuration is saved. You can configure the system to email these reports daily at a time that you specify.
CLI	See command line interface .
client	Also called a client application, an application that runs on one host and relies on another host (a server) to perform some operation. For example, email clients allow you to send and receive email. When the system detects that a user on a host is using a specific client to access another host, it reports that information in the host profile and network map , including the name and version (if available) of the client.
client application	See client .
clipboard	A holding area where you can copy up to 25,000 intrusion events that you can later add to incidents .
cloud services	See Sourcefire cloud .
clustering	A feature that allows you to achieve redundancy of networking functionality and configuration data between two peer Series 3 devices or stacks . Clustering provides a single logical system for policy applies, system updates, and registration. Compare with high availability , which allows you to configure redundant Defense Centers .
command line interface	A restricted text-based interface on Series 3 and virtual devices . The commands that CLI users can run depend on the users' assigned level of access.
complex constraint	A constraint set in an event view or event search that constrains an event query using all the criteria from a specific event.
compliance white list	Along with correlation rules , one of the ways you can specify criteria that network traffic must meet in order to violate a correlation policy . You can use the Defense

	<p>Center to configure compliance white lists to specify which operating systems, applications, and protocols are allowed to run on the hosts in a specific subnet. You can also configure the Defense Center to launch a response, such as an alert or remediation, when a white list is violated. Note that compliance white lists are not associated with the other types of whitelist.</p>
compliance white list event	See white list event .
compliance white list violation	See white list violation .
configurable bypass	A characteristic of an inline set that allows you to configure bypass mode .
connection	A monitored session between two hosts . You can log connections detected by managed devices in the access control policy ; you configure NetFlow connection logging in the network discovery policy .
connection event	An event generated when the system detects a connection between a monitored host and any other host. Connection events include information about the detected traffic. For connections detected by managed devices , depending on the access control rule action , default action , or Security Intelligence decision, you can log a connection event at the beginning or end of a connection, or both. You can log these connections to the Defense Center database; depending on the rule or default action, you can also log connection events to the syslog or to an SNMP trap server. NetFlow connections are end-of-connection and are always saved to the database.
connection graph	A way of displaying connection events in graphical form.
connection log	See connection event .
connection summary	Connection data aggregated over a five-minute interval. The system uses connection summaries to build connection graphs and traffic profiles . To be aggregated, multiple connections must represent the end of connections, have the same source and destination IP addresses, and use the same port on the responder (destination) host . They must use the same protocol (TCP or UDP) and application protocol . Finally, they must either be detected by the same Sourcefire managed device , or be exported by the same NetFlow -enabled device.
connection tracker	One or more conditions that constrain a correlation rule so that after the rule's initial criteria are met, the system begins tracking certain connections . The rule then triggers only if the tracked connections meet additional criteria.
Context Explorer	A page that displays detailed, interactive graphical information about your monitored network, using intrusion , connection , file, geolocation , malware, and discovery data . Distinct sections present information in the form of vivid line, bar, pie, and donut graphs, accompanied by detailed lists. You can easily create and

apply custom filters to fine-tune your analysis, and you can examine data sections in more detail by clicking or hovering your cursor over graph areas. Compared with a [dashboard](#), which is highly customizable, compartmentalized, and updates in real time, the Context Explorer is manually updated, designed to provide broader context for its data, and has a single, consistent layout designed for active user exploration.

context menu	A pop-up menu, available on many of the pages in the web interface, that you can use as a shortcut for accessing other features in the Sourcefire 3D System. The contents of the menu depend on several factors, including the page you are viewing, the specific data you are investigating, and your user role . Context menu options include links to intrusion rule , event , and host information; various intrusion rule settings, quick links to the Context Explorer ; options to add a host to the Security Intelligence global blacklist or global whitelist by its IP address; and options to add a file to the global whitelist by its SHA-256 hash value .
Control license	A license that allows you to implement user control and application control by adding user and application conditions to access control rules . It also allows you to configure your managed devices to perform switching and routing (including DHCP relay and NAT), as well as clustering managed devices.
correlation	A feature you can use to build a correlation policy that responds in real time to threats on your network. The remediation component of correlation provides a flexible API that allows you to create and upload your own custom remediation modules to respond to policy violations.
correlation event	An event generated by the Defense Center when a correlation rule triggers. Note that white list events , generated by white list violations , are a special kind of correlation event.
correlation policy	A policy that describes the network activity that constitutes a security policy violation, using correlation rules and compliance white lists . You can specify responses to each rule or white list within a policy.
correlation rule	With compliance white lists , one of the ways you can specify criteria that network traffic must meet in order to violate a correlation policy . You can use the Defense Center to configure correlation rules to trigger (and generate a correlation event) when a specific intrusion event , malware event , discovery event , host input event , or connection event occurs, or when your network traffic deviates from your normal network traffic pattern as characterized in a traffic profile . You can constrain correlation rules with host profile qualifications , connection trackers , snooze periods , and inactive periods . You can also configure the Defense Center to launch a response, such as an alert or remediation , when a correlation rule triggers.
CRL	See certificate revocation list (CRL) .

current identity	The operating system or server identity that the system finds most likely to be correct for a particular network asset. The system uses this data in many ways; for example, to calculate statistics, assign vulnerability information, assess impact of an attack, and evaluate correlation rules .
current user	The user that the system associates with a host . If the user is an access-controlled user , the system can perform user control on traffic to or from that host. If no access-controlled user is associated with the host, a non-access-controlled user can be the current user for the host. However, after an access-controlled user logs into the host, only a login by another access-controlled user changes the current user.
custom fingerprint	See fingerprint .
custom table	A table you can construct that combines fields from two or more of the predefined tables delivered with the Sourcefire 3D System. For example, you could combine the host criticality information from the host attributes table with information from the connection data table to examine connection data in a new context.
custom topology	A feature that allows you to meaningfully organize and identify subnets in the host , mobile device , and network device network maps .
custom user role	A user role with specialized access privileges. Custom user roles may have any set of menu-based and system permissions, and may be completely original or based on a predefined user role.
custom workflow	A workflow that you create to meet the unique needs of your organization.
dashboard	A display that provides at-a-glance views of current system status, including data about the events collected and generated by the system. To augment the dashboards delivered with the system, you can create multiple custom dashboards, populated with the dashboard widgets you choose. Compare with the Context Explorer , which offers a broad, brief, and colorful picture of how your monitored network looks and acts.
dashboard widget	A small, self-contained dashboard component that provides insight into an aspect of the Sourcefire 3D System.
data correlator	A program that generates events and creates the network map on the Defense Center , using the data collected by the system.
database access	A feature that allows read-only access to the Defense Center database by a third-party client.
decoder	A component of intrusion detection and prevention that places sniffed packets into a format that can be understood by a preprocessor .

default action
to
discovery event

default action	As part of an access control policy , determines how to handle traffic that does not meet the conditions of any rule in the policy. When you apply an access control policy that does not contain any access control rules or Security Intelligence settings, the default policy action determines how non-fast-pathed traffic on your network is handled. You can set the default action to block or trust traffic without further inspection, or inspect it with a network discovery policy or intrusion policy .
Defense Center	A central management point that allows you to manage devices and automatically aggregate and correlate the events they generate.
defragmentation policy	A policy that describes how the IP defragmentation preprocessor (a component of intrusion detection and prevention) should reassemble fragmented IP packets, based on the target host 's operating system. Note that adaptive profiles use adaptive defragmentation policies.
derived fingerprint	An operating system fingerprint created by the system from all passively collected fingerprints for a host by applying a formula which calculates the most likely identity, using the confidence value of each collected fingerprint and the amount of corroborating fingerprint data between identities.
device	A fault-tolerant, purpose-built appliance available in a range of throughputs. Depending on the licensed capabilities you enable on your devices, you can use them to passively monitor traffic to build a comprehensive map of your network assets, application traffic, and user activity , perform intrusion detection and prevention , perform access control , and configure switching and routing. You must manage devices with a Defense Center .
device clustering	See clustering .
device stacking	See stacking .
discovery	A component of the Sourcefire 3D System that uses managed devices to monitor your network and provide you with a complete, persistent view of your network. Network discovery determines the number and types of hosts (including network devices and mobile devices) on your network, as well as information about the operating systems, active applications , and open ports on those hosts. You can also configure Sourcefire managed devices to monitor user activity on your network, which allows you to identify the source of policy breaches, attacks, or network vulnerabilities.
discovery data	Host, user, and application information that qualifies your network assets and traffic flow, as gathered by the discovery feature.
discovery event	An event that details the discovery of new assets or changes to existing assets. A host input event is a special kind of discovery event. Sometimes, Sourcefire uses the term "discovery event" as a general reference to any discovery data or

	vulnerability information.
discovery policy	See network discovery policy .
discovery rule	Within a network discovery policy , specifies the networks and zones you want to monitor and the devices (including NetFlow -enabled devices) or you want to use to monitor them, as well as any ports you want to exclude from monitoring. Each rule also specifies whether you want to discover hosts , users , or applications on the monitored networks.
disposition	See malware disposition .
drill-down page	An intermediate workflow page used to constrain event views. Generally, a drill-down page presents constraints that you can select to advance to a more narrowly constrained page or a table view .
drop event	An intrusion event generated when a drop rule triggers. In the event viewer , drop events are marked with black down arrows.
drop rule	An intrusion rule whose rule state is set to Drop and Generate Events. When a malicious packet triggers the rule in an inline deployment , and the intrusion policy you apply is set to drop when inline, the system drops the packet and generates an intrusion event (specifically, a drop event).
dynamic rule state	A intrusion rule state that is set for a specified period of time in response to a detected rate anomaly in traffic matching the rule.
endpoint	A computer or mobile device where your users install a FireAMP Connector as part of your organization's advanced malware protection strategy.
eStreamer	A component of the Sourcefire 3D System that allows you to stream event data from a Defense Center or managed device to external client applications .
event	A collection of details about a specific occurrence that you can view in the event viewer , using workflows . Events may represent attacks on your network, changes in your detected network assets, violations of your organization's security and network use policies, and so on. The system also generates events that contain information about the changing health status of appliances , your use of the web interface, rule updates , and launched remediations . Finally, the system presents certain other information as events, even though these "events" do not represent particular occurrences. For example, you can use the event viewer to view detailed information about detected hosts , applications , and their vulnerabilities.
Event Streamer	See eStreamer .
event suppression	A feature that allows you to use suppress intrusion events when a specific IP address or range of IP addresses triggers an intrusion rule . Event suppression is useful for eliminating false positives. For example, if you have an email server that

transmits packets that look like a specific exploit, you can suppress events for the rules that are triggered by that server, so you only see the events for legitimate attacks.

event thresholding	A feature that allows you 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. Use event thresholding if you are overwhelmed with a large number of identical events.
event viewer	A component of the system that allows you to view and manipulate events . The event viewer uses workflows to present a broad, then a more focused event view that contains only the events of interest to you. You can constrain the events in an event view by drilling down through the workflow, or by using a search.
export	A method that you can use to transfer various configurations (such as policies) from appliance to appliance. After you export a configuration from one appliance, you can import it onto another appliance of the same type.
external authentication	A method (such as LDAP authentication or RADIUS authentication) that uses externally stored user credentials to authenticate user names and passwords when users log into Sourcefire 3D System appliances . Compare with internal authentication .
failsafe	A characteristic of an inline set that allows packets to bypass processing and continue through the device if internal traffic buffers are full.
fast-path rule	A rule that you configure at a device 's hardware level, using a limited set of criteria, to allow traffic that does not need to be analyzed to bypass processing.
feed	See Security Intelligence feed .
fingerprint	An established definition that the system compares against specific packet header values and other unique data from network traffic to identify a host 's operating system. If the system misidentifies or cannot identify a host's operating system, you can create a custom fingerprint that identifies the host.
file category	A general classification for file types , such as graphics, executables, or archives.
file control	A feature that, as part of access control , allows you to specify and log the types of files that can traverse your network.
file disposition	See malware disposition .
file event	An event that represents a file detected in network traffic by a managed device .
file list	See global whitelist .

file policy	A policy that the system uses to perform file control and advanced malware protection . Populated by file rules , a file policy is invoked by an access control rule within an access control policy .
file rule	A set of criteria within a file policy that the Sourcefire 3D System uses to examine network traffic. If a transmitted file matches the rule criteria, the rule triggers and generates a file event . The rule's file rule action determine whether you block the file (based on file type or malware disposition) or simply allow the file to pass and log the transmission.
file rule action	A setting that determines how the system handles a file that meets the conditions of a file rule . You can detect and alert on specific file types , as well as block the transmission of those files. You can also perform malware cloud lookups on a subset of those file types and block the transmission of those files based on malware disposition .
file trajectory	See network file trajectory .
file type	A specific type of file format, such as PDF, EXE, or MP3.
FireAMP	Sourcefire's enterprise-class, endpoint -based, advanced malware analysis and protection solution that discovers, understands, and blocks malware outbreaks, persistent threats, and targeted attacks. If your organization has a FireAMP subscription , individual users install lightweight FireAMP Connectors on endpoints (computers, mobile devices), which then communicate with the Sourcefire cloud . This allows you to quickly identify and quarantine malware, as well as identify outbreaks when they occur, track their trajectory, understand their effects, and learn how to successfully recover. You can also use the FireAMP portal to create custom protections, block execution of certain applications, and create custom whitelists. Compare with network-based advanced malware protection .
FireAMP Connector	A lightweight agent that users in a subscription-based FireAMP deployment install on endpoints , such as computers and mobile devices. Connectors communicate with the Sourcefire cloud , exchanging information that allow you to quickly identify and quarantine malware throughout your organization.
FireAMP portal	The website, http://amp.sourcefire.com/ , where you can configure your organization's subscription-based FireAMP deployment.
FireAMP subscription	A separately purchased subscription that allows your organization to use FireAMP as an advanced malware protection (AMP) solution. Compare with a Malware license , which you enable on managed devices to perform network-based AMP.
FireSIGHT license	The default license on the Defense Center , which allows you to perform host , application , and user discovery. The FireSIGHT license also determines how many

	<p>individual hosts and users you can monitor with the Defense Center and its managed devices, as well as the number of access-controlled users you can use in access control rules to perform user control.</p>
FireSIGHT Recommendations layer	<p>A built-in layer in an intrusion policy that exists when you allow the system to modify rule states to those recommended by the FireSIGHT recommended rules feature.</p>
FireSIGHT recommended rules	<p>A feature that recommends which rules should be enabled or disabled in your intrusion policy, based on information from your network map. You can choose to allow the system to modify rule states based on recommendations, in which case the system adds a read-only FireSIGHT Recommendations layer.</p>
GeoDB	<p>See geolocation database.</p>
geolocation	<p>A feature that provides data on the geographical source of routable IP addresses detected in traffic on your monitored network, including connection type, internet service provider, and so on. You can see geolocation information, which is stored in the geolocation database, in connection events, intrusion events, file events, and malware events, as well as in host profiles.</p>
geolocation database	<p>Also called the GeoDB, a regularly updated database of known geolocation data associated with routable IP addresses.</p>
GID	<p>Generator ID, a number that indicates which component of the Sourcefire 3D System generated an intrusion event. GIDs help you analyze events more effectively by categorizing the type of event in the same way a rule's SID offers context for the packets that trigger rules.</p>
global blacklist	<p>A Security Intelligence object included by default in every access control policy's Security Intelligence blacklist. The global blacklist applies to all security zones. You can add individual IP addresses to the global blacklist using the IP address context menu in the dashboard, Context Explorer, and many event viewer pages.</p>
global whitelist	<p>A Security Intelligence object included by default in every access control policy's Security Intelligence whitelist. The global whitelist applies to all security zones. You can add individual IP addresses to the global whitelist using the IP address context menu in the dashboard, Context Explorer, and many event viewer pages.</p>
global malware whitelist	<p>A list of files as represented by their SHA-256 hash values. When the system detects a file in the list, it does not perform a malware cloud lookup or block the file as malware, even if the disposition for the file in the Sourcefire cloud is Malware.</p>
HA link interface	<p>Also called the high availability link interface, a physical interface that you configure on each member of a clustered pair of devices to act as a redundant communications channel for sharing health information between the devices.</p>

health event	An event generated when one of the appliances in your deployment meets (or fails to meet) performance criteria specified in a health module . Health events can also generate alerts .
health module	A test of a particular performance aspect, such as CPU usage or available disk space, of the appliances in your deployment. Health modules, which you enable in a health policy , generate health events when the performance aspects they monitor reach a certain level.
health monitor	A feature that continuously monitors the performance of the appliances in your deployment. The health monitor uses health modules within an applied health policy to test the appliances.
health monitor blacklist	A configuration that temporarily disables aspects of health monitoring to prevent the generation of unnecessary health events . You can disable monitoring for a group of appliances , a single appliance, or a specific health module .
health policy	The criteria used when checking the health of an appliance in your deployment. Health policies use health modules to indicate whether your Sourcefire 3D System hardware and software are working correctly. You can use the default health policy or create your own.
high availability	A feature that allows you to configure redundant physical Defense Centers to manage groups of devices . Event data streams from managed devices to both Defense Centers and most configuration elements are maintained on both Defense Centers. If your primary Defense Center fails, you can monitor your network without interruption using the secondary Defense Center. Compare with clustering , which allows you to designate redundant devices.
host	A device that is connected to a network and has a unique IP address. To the Sourcefire 3D System, a host is any identified host that is not categorized as a mobile device , bridge, router , NAT device , or load balancer .
host attribute	A tool you can use to provide information about hosts detected by the system, classifying them in ways that are important to your network environment. The system has two predefined host attributes, host criticality and notes, as well as host attributes that indicate the compliance of each host with each active compliance white list . You can also create your own host attributes.
host criticality	A host attribute that indicates the business criticality (importance) of any given host detected by the system.
host history	A graphical representation of the last 24 hours of a user's activity. The host history, which you can view in a user's user details , displays the IP addresses of the hosts that the user logged into, with approximate login and logout times represented by bar graphs.
host input	A feature that allows you to import data from third-party sources using scripts or command-line files to augment the information in the network map . The web interface also provides some host input functionality; you can modify operating

	system or application protocol identities, validate or invalidate vulnerabilities, and delete various items from the network map, including clients and server ports.
host input event	A kind of discovery event that is generated when you use the host input feature. In general, the system treats host input and passive discovery events identically, though they are distinguished when building correlation rules .
host profile	Collected information about a specific detected host . This includes general host information, such as its name and operating system, as well as the protocols and applications running on the host. The host profile may also include user history , host attributes , VLAN information, applicable white list violations , detected vulnerabilities, and scan results for that host.
host profile qualification	A constraint placed on a traffic profile or correlation rule . A host profile qualification within a correlation rule specifies that the Defense Center should generate a correlation event only if the host involved meets certain criteria. A host profile qualification within a traffic profile limits the hosts that are profiled.
host view	The final page in workflows that display discovery events or network assets. The host view displays the host profiles of the hosts involved in the events or assets you are viewing.
HTTP response page	A web page you can configure the system to display when a user's HTTP request is blocked by an access control policy . You can display a generic Sourcefire-provided response page, or you can provide custom HTML. If the request was blocked by an Interactive Block rule, you can allow users to click a button on the response page to continue to the originally requested site.
hybrid interface	A logical interface on a managed device that allows the system to bridge traffic between a virtual router and a virtual switch .
identity conflict	The conflict that occurs when the system reports a new passive operating system or server identity that conflicts with the current active identity and previously reported passive identities.
impact flag	For intrusion events , a numbered indicator of the correlation between intrusion data, discovery data , and vulnerability information. Impact level 1 (red impact flag) means that the targeted host is <i>vulnerable</i> to the attack represented by the intrusion event, impact level 2 (orange impact flag) means it is <i>potentially vulnerable</i> , and so on. Attacks directed at hosts on networks not monitored by the network discovery policy are impact level 0 (gray impact flag), which indicates that the Defense Center cannot determine the events' impact.
import	A method that you can use to transfer various configurations from appliance to appliance. You can import configurations that you previously exported from another appliance of the same type.

inactive period	An interval during which a correlation rule does not trigger. You can configure the time, frequency, and duration of inactive periods. See also snooze period .
incident	One or more intrusion events that you suspect are involved in a possible violation of your security policy . The system provides incident-handling features that you can use to collect and process information that is relevant to your investigation of the incident.
inline deployment	A deployment of the Sourcefire 3D System where your managed devices are placed inline on a network. In this configuration, devices can affect network traffic flow using switching, routing, access control , and intrusion detection and prevention .
inline interface	A sensing interface configured to handle traffic in an inline deployment . You must add inline interfaces to inline sets in pairs.
inline set	One or more pairs of inline interfaces .
Interactive Block	An access control rule action that allows your users to click a button on an HTTP response page to continue to an initially blocked web site.
internal authentication	An authentication method that stores user credentials in a local database on the appliance . When a user logs into the appliance, the user name and password are checked against the information in the database. Compare with external authentication .
intrusion	A security breach, attack, or exploit that occurs on your network.
intrusion detection and prevention	The monitoring of your network traffic for security policy violations, and, in inline deployments , the ability to block or alter malicious traffic. In the Sourcefire 3D System, you perform intrusion detection and prevention when you associate an intrusion policy with an access control rule or default action.
intrusion event	An event that records an intrusion policy violation. Intrusion event data includes the date, time, and the type of exploit, as well as other contextual information about the attack and its target.
intrusion policy	A variety of components that you can configure to inspect your network traffic for intrusions and security policy violations. These components include intrusion rules that inspect the protocol header values, payload content, and certain packet size characteristics; variables commonly used in intrusion rules; a FireSIGHT recommended rules configuration; advanced settings such as preprocessors and other detection and performance features; and preprocessor rules that allow you to generate events for associated preprocessor options. When your network traffic meets the conditions in an access control rule , you can inspect that traffic with an intrusion policy; you can also associate an intrusion policy with the default action .

intrusion rule	A set of keywords and arguments that, when applied to monitored network traffic, identify potential intrusions , security policy violations, and security breaches. The system compares packets against rule conditions. If the packet data matches the conditions, the rule triggers and generates an intrusion event . Intrusion rules include drop rules and pass rules .
layer	A complete set of intrusion rule , preprocessor rule , and advanced setting configurations within an intrusion policy . You can add custom user layers to the built-in layer or layers in your policy. A setting in a higher layer in an intrusion policy overrides a setting in a lower layer.
LDAP authentication	A form of external authentication that verifies user credentials by comparing them to a Lightweight Directory Access Protocol (LDAP) directory stored on an LDAP directory server.
Lights-Out Management (LOM)	A Series 3 feature that allows you to use an out-of-band Serial over LAN (SOL) management connection to remotely monitor or manage appliances without logging into the web interface of the appliance. You can perform limited tasks, such as viewing the chassis serial number or monitoring such conditions as fan speed and temperature.
link state propagation	An option for inline sets in bypass mode that automatically brings down the second interface in a pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up also. In other words, if the link state of a paired interface changes, the link state of the other interface changes automatically to match it.
list	See Security Intelligence list .
load balancer	A network device that distributes traffic to optimize performance and resource use. Using discovery , the system can identify load balancers .
logical interface	A virtual subinterface that you define to handle traffic with specific VLAN tags as the tagged traffic passes through a physical interface .
malware blocking	A component of Sourcefire's network-based advanced malware protection (AMP) solution. After malware detection yields a malware disposition for a detected file, you can either block the file or allows its upload or download. Compare this functionality with FireAMP , Sourcefire's endpoint-based AMP tool that requires a FireAMP subscription .
malware cloud lookup	A process by which the Defense Center communicates with the Sourcefire cloud to determine the malware disposition of a file detected in network traffic, based on the file's SHA-256 hash value .
malware detection	A component of Sourcefire's network-based advanced malware protection (AMP) solution. File policies applied to managed devices as part of your overall access control configuration inspect network traffic. The Defense Center then performs

	<p>malware cloud lookups for specific detected file types, and generates events that alert you to the files' malware dispositions. AMP malware blocking follows and either blocks the file or allows its upload or download. Compare this functionality with FireAMP, Sourcefire's endpoint-based AMP tool that requires a FireAMP subscription.</p>
malware disposition	<p>A determination by the Sourcefire cloud as to whether a file contains malware, based on the file's SHA-256 hash value.</p>
malware disposition cache	<p>A cache on the Defense Center that stores malware dispositions for files. To improve performance, if the system already knows the disposition for a file based on its SHA-256 hash value, the Defense Center uses the cached disposition rather than performing a malware cloud lookup. Dispositions in the cache time out after a certain period of time so that cache data does not become stale.</p>
malware event	<p>An event generated by one of Sourcefire's advanced malware protection solutions. Network-based malware events are generated when the Sourcefire cloud returns a malware disposition for a file detected in network traffic; retrospective malware events are generated when that disposition changes. Compare with endpoint-based malware events, which are generated when a deployed FireAMP Connector detects a threat, blocks malware execution, or quarantines or fails to quarantine malware.</p>
Malware license	<p>A license that allows you to perform advanced malware protection (AMP) in network traffic. Using a file policy, you can configure the system to perform malware cloud lookups on specific file types detected by managed devices. Compare with FireAMP subscription.</p>
malware protection	<p>See advanced malware protection.</p>
malware whitelist	<p>See global whitelist.</p>
managed device	<p>See device.</p>
management interface	<p>The network interface that you use to administer a Sourcefire 3D System appliance. In most deployments, the management interface is connected to an internal protected network. Compare with sensing interface.</p>
mobile device	<p>In the Sourcefire 3D System, a host identified by the discovery feature as a mobile, handheld device (such as a mobile phone or tablet). The system can often detect whether a mobile device is jailbroken.</p>
monitor	<p>In an access control policy, a way to log traffic that matches a Security Intelligence blacklist or access control rule, but allows the system to continue to evaluate the traffic rather than immediately allowing or blocking it.</p>
NAT	<p>Network address translation, a feature most commonly used to share a single internet connection among multiple hosts on a private network. Using discovery,</p>

the system can identify [network devices](#) as [load balancers](#). In addition, in a Layer 3 deployment of the Sourcefire 3D System, you can configure routing with NAT using a [NAT policy](#).

NAT policy	A policy that uses NAT rules to perform routing with NAT .
NAT rule	A set of configurations and conditions that evaluate network traffic and specify how traffic matching those qualifications is translated. NAT rules are added to an existing NAT policy to perform routing using NAT .
NetFlow	An open but proprietary network protocol for collecting IP traffic information, developed by Cisco Systems to run on Cisco IOS-enabled equipment. You can use the information collected by NetFlow-enabled devices to supplement the discovery and connection data collected by the system and to monitor networks not covered by managed devices .
NetMod	A module that you install in the chassis of a managed device that contains the sensing interfaces for that device.
network device	In the Sourcefire 3D System, a host identified as a bridge, router , NAT device, or load balancer .
network discovery	See discovery .
network discovery policy	A policy that specifies the kinds of discovery data (including host , user, and application data) the system collects for specific network segments, including networks monitored by NetFlow -enabled devices. The network discovery policy also manages identity conflict resolution preferences and active detection source priorities.
network file trajectory	A visual representation of a file's path as hosts transfer it across your network. For any file with an associated SHA-256 hash value , the trajectory map displays the IP addresses of all hosts that have transferred the file, the time the file was detected, the file's malware disposition , associated file events and malware events , and so on.
network map	A detailed representation of your network. The network map allows you to view your network topology in terms of the hosts , mobile devices , and network devices running on your network, as well as their associated host attributes , application protocols , and vulnerabilities.
network object	A reusable object that represents one or more IP addresses, CIDR blocks, or prefix lengths.
Nmap	Network Mapper, an open source active scanner that you can use to detect operating systems and application protocols running on a host. Running an Nmap scan adds the information detected to your network map .

non-access-controlled user
to
physical interface

non-access-controlled user	Any user, whether detected by the User Agent or a managed device , that is not used for access control . A non-access-controlled user can only be the current user for a host if no access-controlled user has ever logged into the host.
non-bypass mode	A characteristic of an inline set that blocks traffic if the sensing interfaces in the set fail for any reason.
object	A reusable configuration that associates a name with a value (for example, an IP address or URL) so that when you want to use that value, you can use the named object instead. You can use objects in various places in the web interface, including discovery rules , access control rules , reports, dashboards , and event searches . See also network object , Security Intelligence object , port object , VLAN tag object , URL object , application filter , global malware whitelist , and security zone .
object, for import or export	A set of configurations, such as a policy or custom workflow , that is created on an appliance and can be exported from that appliance and imported by another appliance.
object manager	The page on the web interface where you manage objects and object groups, including application filters , the global malware whitelist , and security zones .
operating system identity	The operating system vendor and version details for an operating system on a host .
packet view	A type of workflow page that provides detailed information about the packet that triggered an intrusion rule or the preprocessor that generated an intrusion event . The packet view is the final page in workflows based on intrusion events.
pass rule	An intrusion rule that, when triggered, does not generate an intrusion event and does not log the details of the packet that triggered the rule. Pass rules allow you to prevent packets that meet specific criteria from generating an event in specific situations, as an alternative to disabling the intrusion rule. Compare with drop rule .
passive detection	The collection of discovery data through analysis of traffic passively collected by managed devices . Compare with active detection .
passive interface	A sensing interface configured to analyze traffic in a passive deployment.
pending (application protocol)	A designation given to an application protocol identity when the system can neither positively nor negatively identify the application protocol. Most often, the system needs to collect and analyze more data before it can identify a pending application protocol.
physical interface	An interface that represents a physical port on a NetMod .

policy	A mechanism for applying settings, most often to an appliance . See access control policy , correlation policy , file policy , health policy , intrusion policy , network discovery policy , and system policy .
policy target	An appliance or zone where you apply a policy . A policy may have multiple targets.
port object	A reusable object that represents an open port that uses transport layer protocols (for example, TCP, UDP, or ICMP).
preprocessor	A feature that normalizes traffic inspected by an intrusion policy and that helps identify network layer and transport layer protocol anomalies by identifying inappropriate header options, defragmenting IP datagrams, providing TCP stateful inspection and stream reassembly, and validating checksums. Preprocessors can also render specific types of packet data in a format that the system can analyze; these preprocessors are called data normalization preprocessors, or application layer protocol preprocessors. Normalizing application layer protocol encoding allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently and obtain meaningful results. Preprocessors generate preprocessor events whenever packets trigger preprocessor options that you configure.
preprocessor event	A type of intrusion event that is generated when a packet triggers specified preprocessor options. Preprocessor events can help you detect anomalous protocol exploits.
preprocessor rule	An intrusion rule associated with a preprocessor or with the portscan flow detector. You must enable preprocessor rules if you want them to generate events . Preprocessor rules have a preprocessor-specific GID (generator ID).
private search	A named set of search criteria for a specific table, tied to your user account. Only you and users with Administrator access can use your private searches.
protected network	Your organization's internal network that is protected from users of other networks by a device such as a firewall. Many of the intrusion rules delivered with the Sourcefire 3D System use variables to define the protected network and the unprotected (or outside) network.
Protection license	A license for Series 3 and virtual devices that allows you to perform intrusion detection and prevention , file control , and Security Intelligence filtering. Without a license, Series 2 devices automatically have Protection capabilities, with the exception of Security Intelligence.
RADIUS authentication	Remote Authentication Dial In User Service, a service used to authenticate, authorize, and account for user access to network resources. You can create an external authentication object to allow Sourcefire 3D System users to authenticate through a RADIUS server.
rate filtering	A form of anomaly detection that sets a new intrusion rule state for a rule based on the rate of matching traffic.

remediation	An action that mitigates potential attacks on your system. You can configure remediations and, within a correlation policy , associate them with correlation rules and compliance white lists so that when they trigger, the Defense Center launches the remediation. This can not only automatically mitigate attacks when you are not immediately available to address them, but can also ensure that your system remains compliant with your organization's security policy . The Defense Center ships with predefined remediation modules , and you also can use a flexible API to create custom remediations.
remediation instance	A set of configurations for a remediation module . You can configure multiple instances per module, for example, you could respond to different correlation policy violations using the same module, but different instances that have different settings. When a remediation instance triggers, the resulting action is called a remediation .
remediation module	A program that launches a remediation , using sets of configurations called remediation instances . The Sourcefire 3D System ships with several remediation modules that perform various actions; you can also use a flexible API to create your own modules.
remediation status event	An event generated when a remediation is launched.
report template	A template that specifies the data constraints and formats for a report and its sections.
reputation (IP address)	See Security Intelligence .
reputation (URL)	See URL reputation .
response	A reaction to a correlation policy violation — either an alert or a remediation .
retrospective malware event	A network-based malware event generated when the malware disposition for a previously detected file changes. When this occurs, the system also updates the dispositions for files and malware that share the retrospective event's SHA-256 hash value .
risk	See application risk .
routed interface	An interface that routes traffic in a Layer 3 deployment. You can set up physical routed interfaces for handling untagged VLAN traffic, and logical routed interfaces for handling traffic with designated VLAN tags. You can also add static Address Resolution Protocol (ARP) entries to routed interfaces.
router	A network device , located at a gateway, that forwards packets between networks. Using network discovery , the system can identify routers. In addition, you can configure managed devices as virtual routers that route traffic between

	two or more interfaces.
rule	A construct, usually within a policy , that provides criteria against which network traffic is examined.
rule action	A setting that determines how the system handles network traffic that meets the conditions of a rule. See access control rule action and file rule action .
rule state	Whether an intrusion rule is enabled (set to Generate Events or Drop and Generate Events), or disabled (set to Disable) within an intrusion policy . If you enable a rule, it is used to evaluate your network traffic; if you disable a rule, it is not used.
rule update	An as-needed intrusion rule update that contains new and updated standard text rules , shared object rules , and preprocessor rules. A rule update may also delete rules, modify default intrusion policy settings, and add or delete system variables and rule categories.
scheduled task	An administrative task that you can schedule to run once or at recurring intervals.
Security Intelligence	A feature that allows you to specify the traffic that can traverse your network, per access control policy , 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 . Optionally, you can use a monitor setting for Security Intelligence filtering, which allows the system to analyze connections that would have been blacklisted, but also logs the match to the blacklist.
Security Intelligence blacklist	In an access control policy , a list of IP addresses that allows you to deny traffic to and from those hosts, before the traffic is subjected to analysis by access control rules . A blacklist is comprised of Security Intelligence objects , including the global blacklist . An access control policy's Security Intelligence whitelist overrides its blacklist.
Security Intelligence feed	One of the types of Security Intelligence objects , a dynamic collection of IP addresses that the system downloads on a regular basis, at an interval you configure. Because feeds are regularly updated, using them ensures that the system uses up-to-date information to filter your network traffic using the Security Intelligence feature. See also Sourcefire Intelligence Feed .
Security Intelligence list	A simple static collection of IP addresses that you manually upload to the Defense Center as a Security Intelligence object . Use lists to augment and fine-tune Security Intelligence feeds as well as the global blacklist and global whitelist .
Security Intelligence object	A single configuration that represents one or more IP addresses, and that you add to an access control policy 's Security Intelligence blacklist and Security Intelligence whitelist . Security Intelligence objects include Security Intelligence

	<p>lists, Security Intelligence feeds, and network objects and groups. The global blacklist, global whitelist, and the categories in the Sourcefire Intelligence Feed are also considered Security Intelligence objects.</p>
Security Intelligence whitelist	<p>In an access control policy, a list of IP addresses that forces the policy to examine traffic to and from those hosts using access control rules, that is, to not deny the traffic using Security Intelligence. Because a policy's whitelist overrides its Security Intelligence blacklist, you can use it to fine-tune the blacklist. A whitelist is comprised of Security Intelligence objects, including the global whitelist.</p>
security policy	<p>An organization's guidelines for protecting its network. For example, your security policy might forbid the use of wireless access points. A security policy may also include an acceptable use policy (AUP), which provides employees with guidelines of how they may use their organization's systems.</p>
security policy violation	<p>A security breach, attack, exploit, or other misuse of your network.</p>
security zone	<p>A grouping of one or more inline, passive, switched, or routed interfaces that you can use to manage and classify traffic flow in various policies and configurations. The interfaces in a single zone may span multiple devices; you can also configure multiple security zones on a single device. You must assign each interface you configure to a security zone before it can handle traffic, and each interface can belong to only one security zone.</p>
sensing interface	<p>A network interface on a device that you use to monitor a network segment. Compare with management interface.</p>
Series 2	<p>The second series of Sourcefire appliance models. Because of resource, architecture, and licensing limitations, Series 2 appliances support a restricted set of Sourcefire 3D System features. Series 2 devices include the 3D500, 3D1000, 3D2000, 3D2100, 3D2500, 3D3500, 3D4500, 3D6500, and 3D9900. Series 2 Defense Centers include the DC500, DC1000, and DC3000.</p>
Series 3	<p>The third series of Sourcefire appliance models. Series 3 appliances include 7000 Series and 8000 Series devices, as well as the DC750, DC1500, and DC3500 Defense Centers.</p>
server	<p>The server application (compare with client application) installed on a host, identified by application protocol traffic.</p>
server banner	<p>The first 256 bytes of the first packet detected for a server, which can provide additional information that may help you identify the server. The system collects a server banner only once, the first time the server is detected.</p>
server certificate	<p>An encrypted certificate issued by a certificate authority that provides unalterable confirmation of the server identity. You can request a certificate from any certificate authority and upload that custom certificate to your appliance.</p>

server identity	The application protocol type, vendor, and version details for a server on a host .
SFP module	A small form-factor pluggable transceiver that is inserted into a network module on a 71xx Family device. Sensing interfaces on SFP modules do not allow configurable bypass .
SHA-256 hash value	Sometimes abbreviated as SHA256, a 32-bit string that represents a file for which you are performing a malware cloud lookup . The hash value is calculated using a cryptographic hash function so that files with identical SHA-256 values are very likely to have identical contents.
shared layer	An intrusion policy layer that you allow to be used by other intrusion policies. Policies using a shared layer are updated with changes to intrusion rules and advanced settings in the shared layer when you commit those changes. A shared layer can be modified only in the policy that allows it to be shared; it is read-only in policies using it.
shared object rule	An intrusion rule 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. Shared object rules have a GID (generator ID) of 3.
SID	Signature ID (also Snort ID), a unique identifying number assigned to each intrusion rule . When you create a new rule or modify an existing standard text rule , it is given an SID of 1,000,000 or greater. The SIDs for shared object rules and standard text rules delivered with the Sourcefire 3D System are lower than 1,000,000. Also, preprocessors and decoders use SIDs to identify the different types of packets they detect.
snooze period	An interval specified in seconds, minutes, or hours after a correlation rule triggers during which the Defense Center stops firing that rule, even if the rule is violated again during the interval. When the snooze period has ended, the rule can trigger again (and start a new snooze period). See also inactive period .
Snort	An open source intrusion detection system that performs real-time traffic analysis and packet logging on IP networks. Snort can perform protocol analysis, content searching and matching, and can detect a variety of attacks and probes. Snort uses a flexible rules language to describe network traffic that it should collect or pass. The Sourcefire 3D System uses Snort to test packets against decoders , preprocessors , and intrusion rules .
Sourcefire cloud	Sometimes called <i>cloud services</i> , a Sourcefire-hosted external server where the Defense Center can obtain up-to-date, relevant information including malware, Security Intelligence , and URL filtering data. See also malware cloud lookup .

Sourcefire Intelligence Feed	A collection of regularly updated lists of IP addresses determined by the Sourcefire 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 using Security Intelligence . Because the intelligence feed is regularly updated, using it ensures that the system uses up-to-date information to filter your network traffic.
Sourcefire VRT	Sourcefire's Vulnerability Research Team.
stack	Two to four connected devices that share detection resources.
stacking	A feature that allows you to increase the amount of traffic inspected on a network segment by connecting two to four physical devices in a stacked configuration. When you establish a stacked configuration, you combine the resources of each stacked device into a single, shared configuration.
standard text rule	An intrusion rule created based on the identifiers, keywords, and arguments available in the rule editor. You can create your own custom standard text rules and modify Sourcefire-provided standard text rules. A standard text rule has a GID (generator ID) of 1.
state sharing	A feature that allows clustered devices or stacks to synchronize so that if either device or stack fails, the peer can take over with no interruption to traffic flow. State sharing ensures that strict TCP enforcement, unidirectional access control rules , blocking persistence, and dynamic NAT fail over properly.
sub-server	A server called by another server on the same host.
suppression	See event suppression .
SVID	See vulnerability ID .
switch	A network device that acts as a multiport bridge. Using network discovery , the system identifies switches as bridges. In addition, you can configure managed devices as virtual switches , performing packet switching between two or more networks.
switched interface	An interface that you want to use to switch traffic in a Layer 2 deployment. You can set up physical switched interfaces for handling untagged VLAN traffic, and logical switched interfaces for handling traffic with designated VLAN tags.
system policy	Settings that are likely to be similar for multiple appliances in a deployment, such as mail relay host preferences and time synchronization settings. Use the Defense Center to apply a system policy to itself and its managed devices .
table view	A type of workflow page that displays event information, with one column for each of the fields in the database table. When performing event analysis, you can

tag (application)
to
Unified file

	<p>use drill-down pages to constrain the events you want to investigate before moving to the table view that shows you the details about the events you are interested in. The table view is often the next-to-last page in workflows delivered with the system.</p>
tag (application)	See application tag .
tap mode	An advanced inline set option available on 3D9900 and Series 3 devices where a copy of each packet is analyzed and the network traffic flow is undisturbed instead of passing through the device . Because you are working with copies of packets rather than the packets themselves, the device cannot affect the packet stream even if you configure access control and intrusion policies to drop, modify, or block traffic.
target device	See policy target .
task queue	A queue of jobs that the appliance needs to perform. When you apply a policy , install software updates, and perform other long-running jobs, the jobs are queued and their status reported on the Task Status page. The Task Status page provides a detailed list of jobs and refreshes every ten seconds to update their status.
third-party vulnerability	Vulnerability data obtained from a third party. If your organization can write scripts or create command line import files to import network map data from third-party applications , you can use the host input feature to import third-party vulnerability data to augment the system's vulnerability data.
thresholding	See event thresholding .
time window	A time constraint on the events in any event view. Different event views may have different default time windows, depending on your user preferences. Note that not all event views can be constrained by time.
traffic profile	A profile of the traffic on your network, based on connection data logged over a time span that you specify. You can create profiles using all the traffic on a monitored network segment, or you can create more targeted profiles. Then, you can use the correlation feature to detect abnormal network traffic by evaluating new traffic against an existing profile.
transparent inline mode	An advanced inline set option that allows a device to act as a "bump in the wire" and to forward all the network traffic it sees, regardless of its source and destination.
unidentified host	A host whose operating system cannot be identified because the system has not yet gathered enough information about the host. Compare with unknown host .
Unified file	A binary file format that the Sourcefire 3D System uses to log event data.

unknown host	A host whose traffic has been analyzed by the system, but whose operating system does not match any known fingerprints . Compare with unidentified host .
URL category	A general classification for a URL, such as malware or social networking.
URL filtering	A feature that 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 URL category and URL reputation information about those URLs, which is obtained from the Sourcefire cloud by the Defense Center . You can also achieve more granular, custom control over web traffic by specifying individual URLs or groups of URLs to allow or block.
URL Filtering license	A license that allows you to perform URL filtering based on URL category and URL reputation information. URL Filtering licenses may expire.
URL object	A reusable object that represents an individual URL.
URL reputation	A representation of how likely a URL is to be used for purposes that might be against your organization's security policy , as determined by the Sourcefire cloud .
user	A user whose network activity has been detected by a managed device or User Agent .
user activity	An event generated when the system detects a user login (optionally, including some failed login attempts) or the addition or deletion of a user record from the Defense Center database.
User Agent	An agent you install on a server to monitor users as they log into the network or when they authenticate against Active Directory credentials for any other reason. Activity by access-controlled users is only used for access control only when a User Agent reports it.
user awareness	A feature that allows your organization to correlate threat, endpoint, and network intelligence with user identity information, and that allows you to perform user control .
user certificate	An encrypted certificate that identifies a user's browser to the Sourcefire 3D System web server, allowing the server to do a secondary verification of user identity. The certificate must be issued by the same certificate authority that issued the server certificate for your appliance.
user control	A feature that, as part of access control , allows you to specify and log the user-associated traffic that can enter your network, exit it, or cross from within without leaving it.
user details	The final page in user identity and user activity workflows . Along with general information about the user, the user details also display a host history , which is a graphical representation of the last twenty-four hours of the user's activity.

user history	A graphical representation of the last twenty-four hours of user activity for a host . The user history, which you can view in a host's host profile , displays the user names of the users detected logging into the host, with approximate login and logout times represented by bar graphs.
user identity	See user .
user layer	A layer in an intrusion policy where you can modify settings in the policy.
user role	The level of access granted to a user of the Sourcefire 3D System. For example, you can grant different access privileges to the web interface for event analysts, the administrator managing the Sourcefire 3D System, users accessing the Defense Center database using third-party tools, and so on. You can also create custom roles with specialized access privileges.
user role escalation	A privilege you can give to custom user roles that allows users to enter a password to gain the permissions of another user role for the duration of a login session.
UTC time	Coordinated Universal Time. Also known as Greenwich Mean Time (GMT), UTC is the standard time common to every place in the world. The Sourcefire 3D System uses UTC, although you can set the local time using the Time Zone feature.
variable	A representation of a value that is commonly used in intrusion rules . The Sourcefire 3D System uses preconfigured variables to define networks and port numbers. Rather than hard-coding these values in multiple rules, to tailor a rule to accurately reflect your network environment, you can change the variable value.
VDB	See vulnerability database .
virtual Defense Center	A Defense Center that you can deploy on your own equipment in a virtual hosting environment.
virtual device	A managed device that you can deploy on your own equipment in a virtual hosting environment. Virtual devices do not support hardware-based features, such as high availability , clustering , stacking , NAT , VPN , and fast-path rules , and you cannot configure a virtual device as a virtual switch or virtual router .
virtual router	A group of routed interfaces that route Layer 3 traffic. In a Layer 3 deployment, you can configure virtual routers to route packets by making packet forwarding decisions according to the destination IP address. You can define static routes, configure Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) dynamic routing protocols, as well as implement Network Address Translation (NAT).
virtual switch	A group of switched interfaces that process inbound and outbound traffic through your network. In a Layer 2 deployment, you can configure virtual switches on managed devices to operate as standalone broadcast domains, dividing your

	network into logical segments. A virtual switch uses the media access and control (MAC) address from a host to determine where to send packets.
VLAN	Virtual local area network. VLANs map hosts not by geographic location, but by some other criterion, such as by department or primary use. A monitored host's host profile shows any VLAN information associated with the host. VLAN information is also included in intrusion events , as the innermost VLAN tag in the packet that triggered the event. You can filter intrusion policies by VLAN and target compliance white lists by VLAN. In Layer 2 and Layer 3 deployments, you can configure virtual switches and virtual routers on managed devices to appropriately handle VLAN-tagged traffic.
VLAN tag object	A reusable object that represents an individual VLAN tag.
VPN	A feature that allows you to build secure VPN tunnels among the virtual routers on Sourcefire managed devices , or from managed devices to remote devices or other third-party VPN endpoints.
VPN license	A license that allows you to build secure VPN tunnels among the virtual routers on Sourcefire managed devices , or from managed devices to remote devices or other third-party VPN endpoints.
VRT	See Sourcefire VRT .
vulnerability	A description of a specific compromise to which a host is susceptible. The Defense Center provides information on the vulnerabilities to which each of your hosts is vulnerable in the hosts' host profiles . In addition, you can use the vulnerabilities network map to obtain an overall view of the vulnerabilities that the system has detected on your entire monitored network. If you deem a host or hosts no longer vulnerable to a specific compromise, you can deactivate, or mark as invalid, a specific vulnerability.
vulnerability database	Also called the VDB, a database of known vulnerabilities to which hosts may be susceptible. The system correlates the operating system, application protocols , and clients detected on each host with the VDB to help you determine whether a particular host increases your risk of network compromise. VDB updates may contain new and updated vulnerabilities, as well as new and updated application detectors .
vulnerability detail	The final page in vulnerability workflows . Vulnerability details provide information about a specific vulnerability , including technical details and known solutions.
vulnerability ID	An identification number associated with a particular vulnerability . The Sourcefire vulnerability database and third-party vulnerability databases, such as Bugtraq and CVE, have different vulnerability ID numbering schemes.
vulnerability mapping	The association of vulnerability information with discovery data , so that you can perform impact flag correlation .

web application
to
zone

Glossary

web application	A type of application that represents the content of, or requested URL for, HTTP traffic.
whitelist	A compliance white list , a Security Intelligence whitelist , the global malware whitelist , or a list of IP addresses that you can configure within a remediation to exempt IP addresses from some kind of action.
white list event	An event generated when the system detects that a valid target host has become non-compliant with a compliance white list . White list events are a special kind of correlation event .
white list violation	Information that you can view in the event viewer that details how a host is non-compliant with a compliance white list .
widget	See dashboard widget .
workflow	A series of pages you can use to view and evaluate events by moving from a broad view of event data to a more focused view that contains only the events of interest to you. Workflows can include three types of pages, each of which performs a unique function: drill-down pages , table views , and a final page. Depending on the workflow type, the final page may be a table view , packet view , host view , vulnerability detail , or user details .
zone	See security zone .

Index

Symbols

%U encoding (HTTP Inspect option) 862
\$AIM_SERVERS 699
\$DNS_SERVERS 699
\$EXTERNAL_NET 699
\$FILE_DATA_PORTS 700
\$GTP_PORTS 700
\$HOME_NET 700
\$HTTP_PORTS 700
\$HTTP_SERVERS 700
\$ORACLE_PORTS 700
\$SHELLCODE_PORTS 700
\$SIP_PORTS 700
\$SIP_SERVERS 700
\$SMTP_SERVERS 700
\$SNMP_SERVERS 700
\$SQL_SERVERS 700
\$SSH_PORTS 701
\$SSH_SERVERS 701
\$TELNET_SERVERS 701

Numerics

3D9900s, hardware alert details 2179
3-way handshake timeout (stream option) 948

A

Access Admin 57
access control policies
 advanced settings 447
 applying 466
 comparing 462
 configuring 426
 copying 460
 creating 458
 custom user roles 433
 default action 428, 520
 default action logging 431
 detecting files and malware 448
 editing 459
 filtering rules by device 454
 HTTP response page 436
 introduction 424
 licensing 425
 managing 456
 organizing access control rules 449
 policy preferences 1970
 reports 461
 rule categories 451
 Security Intelligence 438
 targets 434
 URL length 447
access control rules 472
 adding 474

- adding objects 491
- application conditions 498
- categories 451
- comments 521
- detecting intrusions 511
- editing 474
- filtering by device 454
- inspecting files 511
- licensing 473
- logging connections 515
- logging file and malware events 515, 518
- network conditions 493
- organizing 449
- port conditions 504
- rule actions 478
- rule condition mechanics 482
- rule condition types 491
- rule conditions 483, 485, 490
- rule ordering 455
- searching 453, 489
- URL conditions 506
- user conditions 497
- VLAN tags 495
- warnings 455
- zone conditions 492
- access conventions 56
- access list 50
 - configuring 1971
 - external database access 2007
 - SNMP polling 1989
- access-controlled users 1260
- accessing the appliance 58, 60
- ack (rule keyword) 1106
- active responses
 - configuring 1164
 - initiating with drop rules 941
 - initiating with keywords 1160
- active scanning 1692
 - analyzing results 1719
 - monitoring scans 1719
 - Nmap 1693, 1702
 - scan results 1716
 - searching for scan results 1721
- adaptive profiles
 - configuring 1006
 - FireSIGHT recommended rules 1005
 - introduction 1003
 - preprocessors 1004
 - understanding 1004
- add client (discovery event type) 1475
- add host (discovery event type) 1475
- add host (host input event type) 1395
- add port (discovery event type) 1475
- add protocol (discovery event type) 1475
- add protocol (host input event type) 1395
- add scan result (discovery event type) 1395, 1475
- add scan result (host input event type) 1395
- adding devices to a Defense Center 208
- Additional Information (vulnerability details) 1366
- additional MAC detected (discovery event type) 1390, 1471
- Admin access 57, 1906
- administrator rules (access control) 451
- Advanced Malware Protection
 - access control 1201
 - event logging 1202
 - health monitoring 2112, 2121
 - introduction 1197, 1199
 - licensing 1198
 - vs FireAMP 1204
- advanced settings (intrusion policy)
 - automatically enabling 787
 - layers 800
 - modifying 773
- alerting
 - deleting alert responses 534
 - discovery event alerting 536
 - editing alert responses 534
 - email alerts 527
 - email alerts for intrusion events 1042
 - enabling and disabling alert responses 534
 - impact flag alerts 535
 - introduction 524
 - intrusion events 761
 - malware alerting 537
 - retrospective malware alerting 537
 - SNMP alerts 528, 728, 1034
 - syslog alerts 530, 1038, 1040
- any, in variables and intrusion rules 713
- APIs
 - eStreamer 49
 - external database access 49, 1907
 - host input 49, 1269, 1680
 - remediations 50, 1606
- appliance heartbeat monitoring 2112, 2122
- appliance information 1999
- appliance status widget 76
- application details, Context Explorer 130
- application detectors 1663
 - activating 1678
 - creating 1666
 - deleting 1680
 - editing 1679
 - filtering 1675
 - managing 1673
 - navigating 1678

- sorting 1675
 - testing 1673
 - understanding 1263
- application filters
 - creating 176
 - searching 1775
- application information, Context Explorer 126
- application layer decoders 809
- application protocol breakdown 1382
- application protocols 1431
 - connection trackers 1493
 - correlation rules 1473
 - server details 1349
 - traffic profiles 1590
 - white lists 1552
- application statistics dashboard 67
- applications 1424
 - application details 1429
 - application filters 176
 - deleting from host profile 1356
 - host profiles 1354
 - in encrypted traffic 1269
 - introduction 1424, 1429
 - misidentified 1644
 - network map 1321
 - searching 1427, 1432
 - understanding 1425, 1430
 - viewing 1425, 1429
- ASCII encoding (HTTP Inspect option) 862
- asn1 (rule keyword) 1117
- asynchronous network (stream option) 949
- audit logging
 - external streaming 1973
 - intrusion policy 680
 - time window 2219
- auditing
 - audit records 2187
 - field descriptions 2196
 - introduction 2187
 - searching 2197
 - understanding 2196
 - viewing 2188
- authentication objects 1854
 - creating 1866
 - deleting 1897
 - editing 1884
 - LDAP 1866, 1867
 - RADIUS 1885
 - user awareness 1298, 1301
- authentication profiles
 - configuring 1975
 - virtual routers 339
- authoritative user logins 1256

- automatic application bypass 251
 - monitoring 2123
 - status 2112
 - automatically enabling IPS features 787
 - available exploits (vulnerability details) 1366, 1437

B

- Back Orifice
 - detecting 959
 - generator ID 785
- backup and restore 2204
 - backup profiles 2208
 - remote backups 2209
 - scheduling backups 1932
- backup files
 - creating 2205
 - location 2208
 - restoring 2211
- Balanced Security and Connectivity 692
- banners, capturing 1288
- bare byte UTF-8 (HTTP Inspect option) 862
- base (intrusion) policy
 - allowing rule updates to modify 694
 - selecting 695
- base64_data (rule keyword) 1179
- base64_decode (rule keyword) 1179
- blacklists
 - health monitoring 2155, 2158
 - Security Intelligence 163, 166, 438, 441, 1814
- blocking malware, file rule actions 1210
- bookmarks 1838
- both ports (stream reassembly option) 952
- both services (stream reassembly option) 952
- browser requirements 58
- browser timeout
 - configuring 1994
 - configuring exemptions 1905
- Bugtraq ID
 - network map 1323
 - vulnerability details 1365, 1437
- business relevance
 - application details 1431
 - application detectors 1666
 - applications 1426
 - servers 1420
- bypass (AAB)
 - configuring 251
 - status 2112
- bypass mode for fail open fiber interfaces 279

byte_extract (rule keyword) 1156
 byte_jump (rule keyword) 1080
 byte_test (rule keyword) 1083

C

card reset monitoring 2112, 2125
 categories
 application details 1431
 applications 1426
 servers 1420
 setting in an application 1666
 URL category data 507
 certificate revocation lists (CRL) 2006
 scheduling downloads 1934
 change reconciliation 2024
 checksum verification 915
 CIDR notation 57
 Cisco IOS routers
 adding a Cisco IOS instance 1609
 configuring remediations 1608
 IOS block 1612, 1614, 1615, 1616
 Cisco PIX firewalls
 adding a Cisco PIX instance 1618
 Cisco PIX block 1620, 1622
 remediations 1617
 classification, rule 1058
 classtype (rule keyword) 1058
 client fingerprints 1650
 client ports (stream reassembly option) 951
 client requirements 58
 client services (stream reassembly option) 951
 client timeout (discovery event type) 1390, 1471
 client update (discovery event type) 1390, 1471
 clients
 application details 1431
 connection trackers 1493
 correlation rules 1468, 1471, 1473, 1477
 host profile 1354
 host profile qualification 1484
 implied by application protocol 1482
 traffic profiles 1590
 white lists 1554, 1556
 clipboard
 clipboard reports 652
 copying events 618, 623
 deleting intrusion events 654
 introduction 652
 clustered devices, configuring 225
 clustered stacks
 configuring 226
 replacing a stacked device 229
 command line reference (CLI) 2242
 comments
 rules in an access control policy 521
 rules in an intrusion policy 729, 762, 1187
 committing intrusion policy 679
 communication ports 50
 comparison report, intrusion policy 687
 comparison view, intrusion policy 686
 compound constraints 1832
 conditions (in correlation rules) 1503, 1506
 confidence
 host view 1406
 server details 1351
 connection flood, detecting 974
 connection summaries 542, 563, 578
 Connection Summary dashboard 67
 connection summary page 578
 external responders 543
 connection trackers 1486
 syntax 1492, 1495
 connections 539
 access control rule logging 515
 aggregated connection data 563
 Block rule logging 517
 connection summaries 542, 563
 connection summary page 578
 custom connection data workflows 1843
 datasets for graphs 560
 default action logging 431, 520
 detaching connection graphs 569
 exporting graph data as CSV 569
 file information 517, 519, 550, 572, 1224
 graphs 555, 557, 565
 intrusion events 550
 intrusion information 517, 519, 547, 574, 603,
 610
 licensing for connection logging 541
 line graphs 565
 malware information 519
 Monitor rule logging 518, 548, 571
 searching 575
 Security Intelligence logging 440, 445, 520, 544,
 550
 understanding event data 543, 550
 URL length 447
 viewing connection tables and graphs 555
 workflow graphs 563, 564
 Connectivity Over Security 692
 consecutive small segments (stream option) 947
 constraining regular expressions 1028
 content (rule keyword) 1063

- Context Explorer
 - application details list 130
 - application information 126
 - compared with dashboard 118
 - connections by access control action 124
 - context menu 64, 148, 155
 - drilling down 148
 - files by disposition 137
 - files information 135
 - filters 150, 151, 155, 156
 - hosts by risk/business relevance and application 129
 - introduction 117
 - intrusion events by impact 132
 - intrusion events by priority 133
 - intrusion events by risk/business relevance and application 128
 - intrusion events details list 135
 - intrusion information 131
 - malware information 135
 - network information 120
 - operating systems 121
 - refreshing 146
 - setting the time range 147
 - top attackers 132
 - top file names 137
 - top file types 136
 - top hosts receiving files 139
 - top hosts sending files 138
 - top ingress/egress security zones 134
 - top malware detection 140
 - top targets 134
 - top users 133
 - traffic and intrusion events 119
 - traffic by destination IP 124
 - traffic by ingress/egress security zone 125
 - traffic by risk/business relevance and application 127
 - traffic by source IP 122
 - traffic by source user 122
 - traffic by URL 144
 - traffic by URL category 145
 - traffic by URL reputation 145
 - understanding 119
 - URL information 143
- context menus 63, 148, 166
 - adding to malware whitelist 182
- correlation events
 - dashboard widgets 77
 - field descriptions 1524
 - introduction 1521
 - searching 1526
 - understanding 1524
 - viewing 1521
- correlation policies 1459
 - activating and deactivating 1520
 - adding responses 1517
 - adding rules 1515
 - adding white lists 1515
 - creating 1513
 - deleting 1520
 - editing 1520
 - introduction 1459
 - managing 1519
 - remediations 1607
 - response groups 1510
 - responses 1606
 - setting rule and white list priorities 1516
- correlation responses
 - adding to rules and white lists 1517
 - high availability 198
 - remediations 1606
 - response groups 1510
- correlation rules 1459
 - adding to correlation policies 1515
 - basic information 1464
 - building rules 1499
 - conditions 1503, 1506
 - connection data triggers 1476
 - connection trackers 1486
 - creating 1461
 - creating rule groups 1509
 - deleting 1509
 - discovery event triggers 1470
 - editing 1508
 - host input event triggers 1475
 - host profile qualifications 1481, 1483
 - inactive periods 1498
 - introduction 1459
 - intrusion event triggers 1467
 - malware event triggers 1469
 - managing 1508
 - rule groups 1509
 - snooze periods 1498
 - traffic profile triggers 1479
 - triggers 1464
 - user activity event triggers 1474
 - user qualifications 1496, 1497
- correlation white lists, *see* white lists
- cover pages 1759
- CPU usage monitoring 2112, 2124
- criticality, host view 1404
- current sessions widget 78
- current time 104
- current user
 - application details 1432

- applications 1426
 - host profile 1340
 - custom
 - application detectors 1663
 - fingerprints 1649
 - HTTPS server certificates 2002
 - login banner 1987
 - OS mappings 1352, 1353, 1654, 1659, 1684, 1686
 - tables 1779
 - topologies 1326
 - workflows 1807, 1840
 - custom analysis widget
 - configuring 82
 - enabling 1977
 - presets 85
 - understanding 79
 - custom tables
 - creating 1783
 - custom workflows 1794
 - deleting 1786
 - editing 1786
 - introduction 1779
 - searching 1788
 - table combinations 1780
 - understanding 1780
 - workflows 1787
 - custom topology 1326
 - creating 1327
 - importing 1329, 1330
 - managing 1332
 - manually editing 1331
 - topology information 1329
 - custom user roles
 - access control policies 433
 - deleting 1911
 - introduction 1905
 - managing 1908
 - using a predefined user role 1911
 - CVE ID
 - network map 1323
 - vulnerability details 1365
 - cvs (rule keyword) 1155
- D**
- dashboards 66, 106
 - adding widgets 113
 - application statistics 67
 - Connection Summary 67
 - custom dashboards 106
 - default dashboard 66, 2223
 - deleting 116
 - detailed dashboard 68
 - files dashboard 68
 - home page 67
 - malware information 68
 - modifying 111
 - properties 111
 - settings 1977
 - summary dashboard 66
 - tabs 112, 113
 - URL statistics 68
 - user statistics 68
 - viewing 109
 - widgets 70, 73
 - data link layer 585, 635
 - database
 - discovery data preferences 1293
 - external access 49, 2007
 - limits 1978
 - datasets (for connection graphs) 560
 - date published (vulnerability details) 1366, 1437
 - dce_iface (rule keyword) 1121
 - dce_opnum (rule keyword) 1123
 - dce_stub_data (rule keyword) 1124
 - DCE/RPC preprocessor 810
 - generator ID 786
 - global options 811
 - RPC over HTTP transport 817
 - target-based policies 813
 - target-based policy options 818
 - transports 814
 - decoders
 - decoding packets 584, 934
 - FTP Telnet 832
 - HTTP Inspect 850
 - RPC 869
 - decoy portscan 963
 - Defense Center
 - adding devices 208
 - configuring user agent 1307
 - deleting devices 212
 - high availability 193, 194
 - introduction 39
 - licenses 2050
 - updating 2062
 - delete client (discovery event type) 1395, 1475
 - delete host/network (discovery event type) 1395, 1476
 - delete port (discovery event type) 1395, 1475
 - delete protocol (discovery event type) 1395, 1475
 - depth content option (rule keyword) 1068

- derived fingerprints 1645
 - description (vulnerability details) 1366
 - destination IPs in intrusion rules 1051
 - destination ports in intrusion rules 1053
 - detailed dashboard 68
 - detect connectionless DCE/RPC traffic (DCE/RPC option) 822
 - detecting ICMP traffic 490
 - detection_filter (rule keyword) 975, 1165
 - device clustering 220
 - configuring 223
 - configuring interfaces 227
 - editing 224
 - HA link 262
 - maintenance mode 229
 - routed interfaces 303
 - separating a cluster 236
 - switching the active peer 228
 - device groups
 - adding 218
 - creating 218
 - deleting 219
 - editing 218
 - managing 217
 - device list 206
 - DHCP
 - IP address changed (discovery event type) 1390, 1471
 - IP address reassigned (discovery event type) 1390, 1471
 - DHCP relay 310
 - DHCPv4 relay 310
 - DHCPv6 relay 312
 - differentiated services (DS), in rules 1103
 - directional operators, in rules 1055
 - directory transversal (HTTP Inspect option) 863
 - disable pipeline decoding (HTTP Inspect option) 865
 - Discovery Admin 57
 - discovery database purging 2236
 - discovery events 1377, 1382
 - alerting 536
 - analyzing 1388
 - application protocol breakdown 1382
 - event breakdown 1381
 - event statistics 1378
 - event types 1389, 1390, 1394
 - field descriptions 1397
 - monitoring 2112, 2126
 - OS breakdown 1383
 - protocol breakdown 1382
 - searching 1399
 - statistics summary 1379
 - viewing 1396
 - workflows 1386
 - disk status monitoring 2112, 2127
 - disk usage monitoring 2112, 2127
 - disk usage widget 96
 - dispositions, see malware dispositions
 - distance content option (rule keyword) 1068
 - distributed portscans 963
 - DNP3 preprocessor
 - configuring 911
 - generator ID 787
 - dnp3_data (intrusion rule keyword) 1148
 - dnp3_func (intrusion rule keyword) 1149
 - dnp3_ind (intrusion rule keyword) 1151
 - dnp3_obj (intrusion rule keyword) 1152
 - DNS preprocessor
 - configuring 830
 - experimental options 830
 - generator ID 786
 - name server responses 827
 - obsolete resource record types 829
 - overflows in RData text fields 829
 - resource record inspection 827
 - DNS, configuring the cache 1981
 - document attributes, in reporting 1756
 - documentation
 - access conventions 56
 - license conventions 55
 - resources 54
 - double encoding (HTTP Inspect option) 863
 - downloading patches 1367, 1368
 - drill-down pages 1792
 - connection graphs 564
 - intrusion events 613
 - drop rules, definition 1046
 - dsize (rule keyword) 1153
 - duplicate connections 265
 - dynamic routing 316
 - dynamic rule states
 - setting 727, 758
 - understanding 757
- ## E
- email alerting
 - configuring 527, 1044
 - intrusion event response 1042
 - email notification 1983
 - email relay host 1983
 - enable defragmentation (DCE/RPC option) 812
 - enabling features automatically 787

- ERSPAN 585
- ESP 585
- eStreamer 49
- Ethernet 585
- event classifications 1058
- event database limits 1978
- event graphs 601
- event logging, discovery events by type 1295
- event preferences 2218
- event queues, configuration 1016
- event statistics
 - discovery events 1378
 - intrusion events 598
- event summary
 - application protocol breakdown (discovery) 1382
 - event breakdown (discovery) 1381
 - intrusion events 595
 - OS breakdown (discovery) 1383
 - protocol breakdown (discovery) 1382
 - statistics summary (discovery) 1379
- event viewer
 - refresh intervals 1827
 - time windows 1820
- events
 - acknowledging 611
 - copying to the clipboard 618, 623
 - deleting 618, 623
 - event graphs 601
 - generating 586
 - impact evaluation 641
 - intrusion event summary 595
 - reports 1725
 - searching 1770
 - time window 2219
 - time windows 1820
- excessive length value (packet decoder) 936
- expanding time window 2219
- experimental DNS options 830
- experimental TCP options 936
- exporting
 - connection graph data as CSV 569
 - introduction 2225
 - objects 2226
- extended ASCII encoding (HTTP Inspect option) 865
- external alerting
 - email alerting for intrusion events 1042
 - SNMP alerting 1034
 - syslog alerting 1038
- external authentication 1854, 1885
- External Database access 57, 1907

F

- failed logins 1904
- failsafe, enabling 274
- fast_pattern content option (rule keyword) 1074
- fast-path rules 253
 - adding IPv4 rules 253
 - adding IPv6 rules 255
 - deleting 257
 - traffic inspection 425
- feeds
 - RSS feeds 53, 103
 - Security Intelligence 162, 439
 - Sourcefire Intelligence Feed 167
- file control
 - access control 1201, 1206
 - introduction 1197
 - licensing 1198
 - reset connection 1210
- file events 1222
 - logging 518, 1211
 - malware events 1202, 1222, 1230
 - searching 1227
 - understanding event data 1224
 - viewing 1223
- file policies
 - applying 467, 1212
 - comparing 1217
 - creating 1213
 - file rules 1209
 - in access control policies 479, 511, 1206
 - introduction 1206
 - intrusion policies 1207
 - managing 1212
- file rules
 - actions 1209
 - components 1209
 - creating 1214
 - evaluation order 1209
- file_data (rule keyword) 1177
- files
 - adding SHA to whitelist 184
 - adding to whitelist 181
 - Context Explorer 135
 - editing whitelist 185
 - events 1222
 - file categories 1209
 - file types 1209
 - files dashboard 68
 - uploading to whitelist 182
 - viewing associated connections 572, 1224
- filtering intrusion rules
 - in a policy 729

- keywords (in a policy) 730
- keywords (rule editor) 1192
- rule editor 1192, 1195
- strings (in a policy) 739
- strings (rule editor) 1194, 1195
- filters, Context Explorer 150, 151, 155, 156
- fingerprints
 - activating 1660
 - client fingerprints 1650
 - deactivating 1661
 - deleting 1661
 - derived 1645
 - editing 1662
 - introduction 1649
 - managing 1660
 - OS mappings 1352, 1353, 1654, 1659, 1684, 1686
 - server fingerprints 1655
- FireAMP
 - agents 48, 1230
 - FireAMP Connectors 1203, 1219
 - FireAMP portal 1203
 - high availability 199, 1219
 - integrating with the system 1202
 - internet access 1219
 - Sourcefire cloud connections 1219
 - status monitoring 1219, 2112, 2128
 - subscription 1198, 1219
 - vs network-based AMP 1204
- FireSIGHT 46
 - host license monitoring 2113, 2129
 - introduction 46
- FireSIGHT recommended rules 764
 - advanced recommendations 766
 - basic recommendations 765
 - configuring 768
 - generating recommendations 768
 - intrusion policy layers 799
 - scheduling 1943
- fixes (vulnerability details) 1367, 1368
- flags
 - impact 641
 - rule keyword 1107
- flow (rule keyword) 1108
- flowbits (rule keyword) 1168
- flush factor (stream option) 947
- fragbits (rule keyword) 1101
- fragmentation, in IP datagrams 927
- fragoffset (rule keyword) 1155
- frame information (packet view) 634
- FTP decoder
 - client-level options 846
 - configuring 832

- configuring client-level options 847
- generator ID 786
- global options 832, 833
- server-level options 838, 842

G

- generating events 586
- generator IDs (GIDs) 784
- global host profiles 1535, 1549
- global malware whitelist 180
 - enabling in file policy 1213
- global rule thresholding
 - configuring 1012
 - disabling 1014
- graphs
 - detaching connection graphs 569
 - health monitoring 2170
 - intrusion events 601
 - performance statistics 599
- GRE 585
- groups
 - correlation response groups 1510
 - correlation rule groups 1509
 - device groups 217
 - object groups 159
- GTP packets (decoded) 585
- GTP preprocessor
 - configuring 878
 - generator ID 787
- gtp_info (intrusion rule keyword) 1137
- gtp_type (intrusion rule keyword) 1129
- gtp_version (intrusion rule keyword) 1128

H

- HA link, see high availability
- hardware
 - host profile 1341
 - host view 1405
 - monitoring 2130
- hardware alerts for 3Dx900s 2179
- hardware monitoring 2113
- health alerts 2159
 - creating 2159
 - deleting 2163

- editing 2162
 - understanding 2161
- health events 2174
 - field descriptions 2183
 - searching 2184
 - table view 2183
 - understanding 2174
 - viewing 2175
- health modules 2112
 - blacklisting 2158
 - running all modules 2167
 - running specific modules 2168
- health monitor process monitoring 2113
- health monitoring
 - alerts 2159
 - appliance monitoring 2166
 - blacklist 2155, 2158
 - configuring 2115
 - creating alerts 2159
 - deleting alerts 2163
 - editing alerts 2162
 - events 2174
 - graphs 2170
 - health modules 2112
 - health monitor 2163
 - health policies 2111
 - introduction 2109
 - link state propagation 2134
 - power supplies 2136
 - running health modules 2167, 2168
 - status indicators 2165
 - time window 2219
 - understanding 2110
 - understanding alerts 2161
- health policies 2111
 - applying 2145
 - configuring 2116
 - creating 2118
 - defaults 2117
 - deleting 2154
 - editing 2147
- health status monitoring 2131
- high availability 50, 193, 194, 196, 197
 - advanced malware detection 199
 - communication port 200
 - configuration guidelines 199
 - configuring 200
 - correlation responses 198
 - disabling 204
 - FireAMP 199, 1219
 - HA link 262
 - health monitoring blacklist 2156
 - licensing 198, 2045
 - monitoring 202
 - monitoring status 202
 - pausing communications 205
 - restarting communications 205
 - Security Intelligence 164, 198
 - shared configurations 196
 - shared policies 197
 - understanding 195
 - URL filtering 198
 - user agents 199
- hits (server details) 1351
- home page
 - dashboards 67
 - setting 2217
- hops
 - hops change (discovery event type) 1390
 - host profile 1339
 - host view 1404
- host attribute add (discovery event type) 1395
- host attribute delete (discovery event type) 1395
- host attribute delete value (discovery event type) 1395, 1476
- host attribute set value (discovery event type) 1396, 1476
- host attribute update (discovery event type) 1396
- host attributes
 - assigning attribute values 1358
 - attribute types 1371
 - built-in host attributes 1369
 - creating 1372
 - deleting 1375
 - editing 1374
 - host profiles 1358
 - introduction 1412
 - network map 1325
 - predefined 1369
 - remediations 1629
 - searching 1416
 - setting attributes 1415
 - understanding 1413
 - user-defined 1370
 - viewing 1412
- host criticality, host view 1404
- host data collection 1252
- host deleted (discovery event type) 1390, 1471
- host dropped (discovery event type) 1391
- host history 1449
- host input 49
 - custom product mappings 1689
 - event types 1394
 - events 1388, 1396
 - patch tools 1685
 - services 1688

- third-party tools 1682
 - vulnerabilities 1687
 - host license limits 2113
 - host profile qualifications
 - correlation rules 1481
 - in traffic profiles 1590
 - syntax 1591
 - host profiles
 - activating vulnerabilities 1367, 1368
 - applications 1354
 - basic information 1338
 - creating white lists 1361
 - host attributes 1358
 - host protocols 1359
 - identity conflicts 1345, 1353
 - impact qualification 1366
 - introduction 1334
 - malware 1362
 - network map 1316
 - operating system 1341
 - predefined host attributes 1369
 - scan results 1375
 - scanning a host 1376
 - servers 1347
 - shared host profiles 1537
 - systems 1343
 - third-party vulnerabilities 1363
 - user history 1357
 - viewing 1337
 - VLAN tags 1357
 - vulnerabilities 1363
 - vulnerability details 1365
 - white list violations 1360
 - white lists 1534, 1548
 - host protocols 1359
 - host statistics 2096
 - host statistics, intrusion event statistics 596
 - host timeout (discovery event type) 1391, 1471
 - host type (host view) 1405
 - host type changed router/bridge (discovery event type) 1391
 - host type changed to network devices (discovery event type) 1471
 - host view, hardware 1405
 - hosts
 - current user (host profile) 1340
 - host attributes 1412
 - host criticality 1404
 - host history 1449
 - host name (host profile) 1338
 - host type (host profile) 1339
 - identity conflicts 1345, 1353
 - introduction 1401
 - IP address 1403
 - last seen 1340, 1403
 - MAC addresses 1403, 1404
 - misidentification 1643
 - mobile devices 1319, 1343
 - NETBIOS name 1404
 - new host (discovery event type) 1392, 1471
 - searching 1408
 - source type 1406
 - traffic profiles 1407
 - understanding 1403
 - viewing 1402
 - viewing on the network map 1317
 - VLAN ID 1404
 - white lists 1408
 - HTTP decoder, generator ID 785
 - HTTP Inspect
 - additional rules 868
 - configuring 865
 - decoding 850
 - generator ID 785
 - normalization encoding options 862
 - normalization options 851, 853, 854
 - HTTP response page, access control policies 436
 - http_client_body content option (rule keyword) 1073
 - http_cookie content option (rule keyword) 1073
 - http_header content option (rule keyword) 1072
 - http_method content option (rule keyword) 1072
 - http_raw_cookie content option (rule keyword) 1073
 - http_raw_header content option (rule keyword) 1072
 - http_raw_uri content option (rule keyword) 1072
 - http_status_code content option (rule keyword) 1074
 - http_status_encode (rule keyword) 1175
 - http_status_message content option (rule keyword) 1074
 - http_uri content option (rule keyword) 1071
 - HTTPS certificates 2002
 - generating a server certificate request 2003
 - uploading server certificates 2004
 - viewing 2002
 - hybrid interfaces
 - adding logical interfaces 352
 - deleting 356
 - introduction 352
- I
- ICMP
 - header values 1104
 - ICMPv4 585

- ICMPv6 585
- icmp_all (resp keyword) 1162
- icmp_host (resp keyword) 1162
- icmp_id (rule keyword) 1105
- icmp_net (resp keyword) 1162
- icmp_port (resp keyword) 1162
- icmp_seq (rule keyword) 1105
- icode (rule keyword) 1105
- ID (rule keyword) 1102
- identities
 - adding identity sources 1295
 - conflict settings 1289
 - current identity 1646
 - identity conflicts 1648
 - operating system conflicts 1345
 - server conflicts 1353
- identity conflict (discovery event type) 1391, 1472
- identity timeout (discovery event type) 1391, 1472
- IIS backslash (HTTP Inspect option) 863
- IIS encoding (HTTP Inspect option) 863
- IMAP preprocessor
 - configuring 880
 - generator ID 786
 - options 881
- impact (vulnerability details) 1366
- impact flags 641
 - alerting 535
 - descriptions 642
 - impact qualification 1366
- impact qualification 1367
- import logging 2082, 2084, 2089
- importing
 - introduction 2225
 - objects 2231
- inactive periods 1498, 1499
- incidents
 - common processes 657
 - creating 661
 - default incident types 660
 - definition 657
 - editing 663
 - incident handling basics 657
 - incident types 665
 - introduction 656
 - reports 664
- ing 262
- inline interfaces 269
- inline IPS deployments 269
- inline normalization
 - configuring 922
 - introduction 918
 - preprocessor 918
- inline sets
 - adding 271
 - advanced options 275
 - creating 271
 - deleting 279
 - IPS deployments 269
 - strict TCP enforcement 277
 - tap mode 275
 - transparent inline mode 276
- input parameters 1750
- installing system software 2056
- integer (host attribute type) 1371, 1373
- interface status widget 77
- interface traffic widget 97
- interfaces
 - configuring 258
 - configuring a clustered device 227
 - configuring a stacked device 243
 - interface status widget 77
 - interface traffic widget 97
- internet access 50
- introduction 37
- Intrusion Admin 57
- intrusion event responses 588
 - email alerting 1042
 - introduction 1033
 - SNMP alerting 1034
 - syslog alerting 1038, 1040
- intrusion events
 - analyzing 587
 - associated connections 547, 574, 603, 610
 - clipboard 652
 - constraining events 620
 - Context Explorer 119, 131, 133, 135
 - drill-down pages 613
 - event overview 597
 - event statistics 598
 - field descriptions 604
 - impact flags 641
 - introduction 593
 - intrusion event responses 1033
 - packet view 615, 622
 - portscan events 968
 - rate monitoring 2113, 2132
 - reviewing 611
 - searching 644
 - suppressing 753
 - table view of events 613
 - thresholding and suppressing 746
 - understanding 604
 - unreviewing 611
 - viewing 602, 616
 - widget 97
 - workflows 612

- intrusion policies
 - access control default action 429
 - adaptive profiles 1003
 - advanced settings 773
 - alerting 761
 - automatically applying 1938
 - base policy 691
 - base policy (allowing rule updates to modify) 694
 - base policy (selecting) 695
 - benefits 591
 - committing 679
 - comparing 685, 687
 - creating 672
 - default policies 692
 - detection challenges 781
 - drop rule behavior (setting) 689
 - dynamic rule states 756
 - editing 675
 - event filtering 746
 - file policies 1207
 - filtering rules 741
 - FireSIGHT recommended rules 764
 - ignoring VLAN headers 917
 - importing local intrusion rules 2080
 - importing rule updates 694, 2072, 2074
 - in access control policies 479, 511
 - introduction 667, 773
 - layers 792
 - managing 670
 - navigation panel 678
 - normalizing inline traffic 918
 - planning 668
 - predefined variables 699
 - preferences 1985
 - preprocessors 773, 780, 1015
 - rate-based attack prevention 756
 - reapplying 680
 - reporting 682
 - rules 717
 - setting rule state 743
 - suppressing 753
 - thresholding 1009
 - thresholding (understanding) 747
 - troubleshooting options 790
 - using rules in layers 795
 - variables 698, 701
 - viewing rules 719
- intrusion prevention, *see* IPS
- intrusion rules
 - alerting 761
 - comments 1187
 - comments (in a policy) 729
 - creating 1182
 - destination IPs 1051
 - directional operators 1055
 - dynamic rule states 727, 756
 - editing 1184
 - event details 1057
 - filtering (in a policy) 729
 - filtering (rule editor) 1192
 - FireSIGHT recommended rules 764
 - ICMP header values 1104
 - importing local rule files 2080
 - importing rule updates 694, 2072, 2074
 - introduction 1046
 - IP header values 1101
 - keywords 1055
 - managing 717
 - metadata 1096
 - parts of a rule 1047
 - replacing content 1078
 - rule actions 1050
 - rule details 723
 - rule headers 1049
 - searching for 1189
 - setting rule state 743
 - SNMP alerting 728
 - sorting (in a policy) 722
 - source ports 1053
 - specifying IP addresses 1051
 - specifying ports 1053
 - suppressing 753
 - suppression 726
 - TCP header values 1106
 - thresholding (configuring) 725
 - thresholding (understanding) 747
 - invalid IP options (packet decoder) 936
 - invalid RFC delimiters (HTTP Inspect option) 864
 - invalid widgets 71
 - IP addresses
 - excluding 714
 - host view 1403
 - in intrusion rules 1051
 - in rules 1051
 - in variables and rules 713
 - original client IP 646
 - syntax for searches and reports 1776
 - IP defragmentation 927
 - configuring 932
 - exploits 928
 - generator ID 785
 - target-based policies 929
 - IP header values 1101
 - IP layer (packet view) 636
 - IP_Proto (rule keyword) 1103
 - IPopts (rule keyword) 1102

- IPS
 - deployments 588
 - inline deployments 269
 - introduction 581
 - intrusion event responses 588
 - intrusion events 593
 - intrusion policies 667, 773
 - IPS-only deployment 473
 - passive deployments 266
 - preprocessors 780
 - setting up IPS devices 266
 - IPv4 585
 - conventions 57
 - fast-path rules 253
 - IPv6 585
 - conventions 57
 - fast-path rules 255
 - IPV6 Teredo tunneling, detecting 935
 - isdataat (rule keyword) 1153
 - itype (rule keyword) 1105
- J**
- jailbreaking
 - allow in white lists 1542, 1545
 - mobile devices 1342
- K**
- keyword
 - ack 1106
 - asn1 1117
 - base64_data 1179
 - base64_decode 1179
 - byte_extract 1156
 - byte_jump 1080
 - byte_test 1083
 - classtype 1058
 - content 1063
 - content depth option 1068
 - content distance option 1068
 - content fast_pattern option 1074
 - content http_client_body option 1073
 - content http_cookie option 1073
 - content http_header option 1072
 - content http_method option 1072
 - content http_raw_cookie option 1073
 - content http_raw_header option 1072
 - content http_raw_uri option 1072
 - content http_status_code option 1074
 - content http_status_message option 1074
 - content http_uri option 1071
 - content nocase option 1065
 - content offset option 1068
 - content rawbytes option 1065
 - content within option 1069
 - cvs 1155
 - dce_iface 1121
 - dce_opnum 1123
 - dce_stub_data 1124
 - detection_filter 1165
 - dnp3_data 1148
 - dnp3_func 1149
 - dnp3_ind 1151
 - dnp3_obj 1152
 - dsize 1153
 - file_data 1177
 - flags 1107
 - flow 1108
 - flowbits 1168
 - fragbits 1101
 - fragoffset 1155
 - gtp_info 1137
 - gtp_type 1129
 - gtp_version 1128
 - http_encode 1175
 - icmp_id 1105
 - icmp_seq 1105
 - icode 1105
 - ID 1102
 - IP_Proto 1103
 - IPopts 1102
 - isdataat 1153
 - itype 1105
 - metadata 1096
 - modbus_data 1146
 - modbus_func 1146
 - modbus_unit 1147
 - msg 1058
 - pcre 1086
 - pkt_data 1178
 - priority 1058
 - react 1162
 - reference 1062
 - resp 1161
 - rpc 1116
 - sameip 1155
 - seq 1110
 - sip_body 1125

- sip_header 1125
 - sip_method 1126
 - sip_stat_code 1127
 - ssl_state 1113
 - ssl_version 1115
 - stream_reassemble 1112
 - stream_size 1111
 - tag 1166
 - threshold (deprecated) 1166
 - tos 1103
 - ttl 1103
 - urilen 1118
 - window 1110
 - keywords, in rules 1055
- L**
- language 1986
 - last seen (host view) 1340, 1403
 - last successful login 59
 - latency thresholding troubleshooting options 791
 - layer (host profile) 1359
 - layer 2 switches 281
 - layer 3 routers 295
 - layers
 - adding 804
 - advanced settings 800
 - configuring 803
 - FireSIGHT recommendations 797
 - rule settings (removing) 797
 - sharing 794
 - understanding 792
 - working with 792, 795
 - LDAP 1851, 1854, 1879
 - authentication profiles 1975
 - authentication server 1867
 - connecting Defense Centers 1298
 - creating an authentication object 1859
 - enabling and disabling user awareness 1306
 - examples 1880
 - for user access control 1301
 - for user awareness 1258
 - group access settings 1875
 - logging in 60
 - managing user accounts 1902
 - shell access 1878
 - user awareness 1298
 - legacy reassembly (stream option) 948
 - license monitoring 2113, 2133
 - licenses 55
 - access control policies 425
 - access control rules 473
 - access-controlled user limit 2048
 - adding to a Defense Center 2050
 - changing licensed capabilities 2052
 - Control 2042
 - deleting 2052
 - FireSIGHT 2039
 - high availability 198
 - host and user limits 2045
 - license types 2037
 - Malware 2043
 - managing 2036
 - monitoring 2113, 2133
 - product licensing widget 101
 - Protection 2041
 - RNA Host 2039
 - RUA User 2039
 - third-party products 2285
 - URL filtering 2043
 - viewing 2048
 - lights out management 2025, 2028, 2030
 - line graphs (connection data) 565
 - link state propagation
 - configuring 276
 - monitoring 2113, 2134
 - list (host attribute type) 1371, 1373
 - listeners 578
 - load balancers, excluding from monitoring 1276
 - local configurations, certificate revocation lists (CRL) 2006
 - local settings 1998
 - appliance information 1999
 - change reconciliation 2024
 - Cloud Services 2032
 - custom HTTPS certificates 2002
 - external database access 2007
 - licenses 2036
 - lights out management 2025, 2028, 2030
 - network settings 2009
 - remote access 2025
 - remote management 213
 - remote storage 2017
 - shutting down or restarting 2014
 - time 2015
 - local time 104
 - logging into the appliance 58
 - using LDAP 60
 - logging multiple packets per stream 1016
 - logging out of the appliance 62
 - login banner 1987
 - logos 1760

M

- MAC addresses
 - additional MAC detected (discovery event type) 1390, 1471
 - host profile 1339
 - host view 1403, 1404
 - information change (discovery event type) 1392, 1472
- mail relay host 1983
- Maint 57
- Maintenance access 1907
- malware blocking
 - connection event Reason 1211
 - file rule action 1210
 - introduction 1199
 - reset connection 1210
- malware cloud lookup
 - connection event Reason 1211
 - file rule action 1210
 - health monitoring 2112, 2121
 - high availability 1212
 - internet access 1212
 - introduction 1199
- malware detection, high availability 199
- malware dispositions
 - caching 1200
 - changing 1202, 1211, 1229, 1230
 - introduction 1199
- malware events 1229
 - alerting 537
 - Context Explorer 135
 - endpoint-based 1230
 - event types 1238
 - file events 1202, 1222, 1230
 - files dashboard 68
 - logging 518, 1211
 - network vs endpoint 1204, 1212, 1229
 - network-based 1230
 - retrospective 537, 1202, 1211, 1229, 1230
 - searching 1239
 - understanding event data 1233
 - viewing 1232
- malware whitelist
 - adding files by context menu 181
 - adding files by SHA 184
 - adding files by uploading 182
 - editing files 185
 - enabling in file policy 1213
 - managing 180
- malware, host profiles 1362
- managed devices
 - adding to a Defense Center 208
 - automated application bypass 251
 - comparison reports 212
 - configuring interfaces 258
 - configuring the management interface 260
 - deleting 212
 - device clustering 220
 - Device Management page 206
 - disabling interfaces 264
 - editing 244
 - HA link 262
 - introduction 39, 206
 - management concepts 191
 - stacking 237
 - time sync 1993
 - updating 2066
- maximum active responses (TCP stream option) 943
- maximum chunk size (HTTP Inspect option) 864
- maximum fragment size (DCE/RPC option) 812
- maximum response seconds (TCP stream option) 943
- maximum TCP window (stream option) 946
- maximum transmission unit (MTU) 927
- memory cap reached (DCE/RPC option) 812
- memory usage monitoring 2113, 2134
- metacharacters 1088
- metadata (rule keyword) 1096
- Microsoft
 - %U encoding (HTTP Inspect option) 862
 - Active Directory 1300
 - IIS encoding (HTTP Inspect option) 863
- Microsoft Active Directory 1855
- mobile devices
 - behind network device 1319
 - jailbreaking 1342, 1542, 1545
 - network map 1319
- Modbus preprocessor
 - configuring 909
 - generator ID 787
- modbus_data (intrusion rule keyword) 1146
- modbus_func (intrusion rule keyword) 1146
- modbus_unit (intrusion rule keyword) 1147
- MPLS 585
- msg (react keyword) 1163
- msg (rule keyword) 1058
- MTU 927
- multi-rule search engine 586
- multi-slash obfuscation (HTTP Inspect option) 863

N

NAT

- adding NAT entries 344
 - deleting NAT entries 350
 - excluding devices from monitoring 1276
 - managing devices in NAT environments 193
 - NAT entry list 343
 - virtual routers 341
- NAT policies
 - applying 401
 - comparing 397
 - configuring 385
 - copying 395
 - creating 392
 - editing 394
 - introduction 383
 - managing 391
 - organizing NAT rules 389
 - reports 396
 - targets 387
- NAT rules
 - adding 404
 - adding objects 415
 - destination network conditions 419
 - editing 404
 - NAT rule condition mechanics 409
 - NAT rule conditions 411
 - NAT rule types 406
 - network conditions 418
 - organizing 389
 - port conditions 421
 - rule condition types 415
 - rule conditions 410, 414
 - rule ordering 390
 - searching 413
 - warnings 390
 - zone conditions 415
- NetBIOS name
 - host profile 1339
 - host view 1404
- NetBIOS name change (discovery event type) 1392, 1471
- NetFlow 1271
 - adding devices 1292
 - comparing with FireSIGHT data 1272
- Network (TCP stream option) 945
- network address translation, *see* NAT
- Network Admin 57
- network behavioral analysis 1585
- network compliance widget 99
- network devices (on the network map) 1318
- network discovery 1250
 - access control default action 429
 - active detection 1645
 - analyzing events 1388
 - application detection 1263
 - applications 1424, 1429
 - correlation events 1521
 - correlation policies 1459
 - correlation policy remediations 1607
 - correlation rules 1459
 - data collection 1251
 - discovery events 1377
 - discovery rules 1276
 - enhancing detection 1641, 1642
 - host attributes 1370
 - host profiles 1334
 - misidentifications 1643, 1644
 - network map 1314
 - passive detection 1645
 - performance statistics 1384
 - remediation events 1633
 - servers 1418
 - Sourcefire Vulnerability ID (SVID) 1365, 1437, 1443
 - third-party data collection 1269
 - third-party vulnerabilities 1441
 - vulnerabilities 1434
 - white lists 1530
- network discovery events
 - event statistics 1378
 - workflows 1386
- network discovery policies
 - advanced options 1287
 - applying 1297
 - creating 1275
 - discovery rules 1276
 - monitored networks 1279
 - port exclusions 1280
 - restricting user logging 1285
 - zones 1279
- network file trajectories 1241
 - data points 1247
 - events tables 1249
 - SHA-256 hash values 1243
 - summaries 1244
 - trajectory maps 1246
- network layer 585, 636
- network map
 - applications 1321
 - Bugtraq ID 1323
 - custom topology 1326
 - CVE ID 1323
 - host attributes 1325
 - hosts 1316
 - introduction 1314
 - mobile devices 1319
 - network devices 1318

- SIDs 1323
 - understanding 1315
 - viewing hosts 1317
 - vulnerabilities 1323
- network objects 161
 - as Security Intelligence objects 444
 - in variables and intrusion rules 714
- network settings, configuring 2009
- network surveys (white lists) 1543
- Networks (DCE/RPC option) 818
- new client (discovery event type) 1392, 1471
- new host (discovery event type) 1392, 1471
- new network protocol (discovery event type) 1392, 1471
- new operating system (discovery event type) 1393
- new OS (discovery event type) 1472
- new TCP port (discovery event type) 1393, 1472
- new UDP port (discovery event type) 1472
- Nmap
 - adding an Nmap instance 1623
 - introduction 1693
 - managing 1710
 - on-demand scans 1713
 - remediations 1623, 1625, 1705, 1712
 - sample scanning profiles 1699
 - scan instances 1702, 1710
 - scan targets 1704
 - scheduling scans 1936
 - setting up 1702
- no_alert_large_fragments (RPC decoder) 870
- nocase content option (rule keyword) 1065
- non-access-controlled users 1260
- non-RFC characters (HTTP Inspect option) 864
- normalizing inline traffic 918
- notification
 - email alerting for intrusion events 1042
 - SNMP alerting 1034
 - syslog alerting 1038

O

- objects
 - application filters 176
 - exporting 2226
 - groups 159
 - importing 2231
 - in searches 1775
 - introduction 158
 - managing 159
 - network objects 161

- object manager 159
- port objects 173
- Security Intelligence 162, 441
- sorting and filtering 161
- URL objects 175
- VLAN tag objects 174
- obsolete TCP options 937
- offset content option (rule keyword) 1068
- on-demand scans 1713
- open shortest path first (OSPF) 316
- OpenLDAP 1300, 1305, 1855, 1863, 1873
- operating system identities
 - Context Explorer 121
 - editing 1344
 - resolving conflicts 1345
 - viewing 1343
- order of execution, of preprocessors 782
- original client IP address 646
- OS (host view) 1405
- OS breakdown 1383
- OS custom fingerprints 1649
- OS mappings 1352, 1353, 1654, 1659, 1684, 1686
- OS names (host view) 1406
- OS vendor (host view) 1405
- OS version (host view) 1406
- OS, TCP policy option (stream) 946
- OSPF (open shortest path first)
 - configuring 323
 - export filters 334
 - import filters 333
 - routing areas 323
- overlap limit (stream option) 946

P

- packet bytes (packet view) 641
- packet decoder
 - capturing packets 584
 - configuring 938
 - excessive length value 936
 - experimental TCP options 936
 - generator ID 785
 - invalid IP options 936
 - invalid TCP options 938
 - obsolete TCP options 937
 - other protocol header anomalies 938
 - T/TCP options 938
 - Teredo traffic (detecting) 935
 - understanding 934
- packet latency thresholding

- settings 1019
 - understanding 1017
- packet latency, GID 786
- packet size performance boost (TCP stream option) 948
- packet type performance boost (TCP stream option) 943
- packet view 622
 - copying events to the clipboard 618, 623
 - datalink layer 635
 - definition 615
 - deleting events 618, 623
 - event information 625
 - frame information 634
 - network layer 636
 - packet bytes 641
 - portscans 969
 - rule actions 629
 - setting thresholds 631
 - suppression options 633
 - TCP packets 639
 - transport layer 638
 - UDP packets 640
- packets
 - capturing 584
 - data link layer 585
 - decoding 584, 934
 - packet bytes 641
 - processing 585
- passive interfaces 267
- passive IPS deployments 266
- passwords
 - changing 1913, 2216
 - expiration warnings 1904
 - failed logins 1904
 - force reset 1904
 - minimum length 1904
 - password options 1904
 - strength check options 1905
- patches (for vulnerabilities) 1367, 1368
- payloads, detection 1390
- PCRE 1088
 - character classes 1089
 - examples 1094
 - in intrusion rules 1086
 - metacharacters 1088
 - modifier options 1090
- pcre (rule keyword) 1086
- peer manager 201
- performance boost (UDP stream option) 958
- performance graphs 1384
- performance monitor 1026
- performance statistics 598
- Perl-compatible regular expressions 1086, 1088
 - options 1090
- pkt_data (rule keyword) 1178
- policy (DCE/RPC option) 819
- POP preprocessor
 - configuring 884
 - generator ID 787
 - options 885
- port objects 173
- ports
 - exclusion from discovery monitoring 1280
 - in variables 711
 - in variables and intrusion rules 715
 - server details 1348
- ports ignoring small segments (stream option) 948
- portscans 961
 - configuring 965
 - decoy portscan 963
 - distributed portscans 963
 - generator ID 785
 - packet view 969
 - portscan events 968
 - portscan types 963
 - portsweep 963
 - sensitivity levels 964
- portsweep 963
- power supply monitoring 2113, 2136
- PPP 585
- PPPoE 585
- preferences 2215
 - changing passwords 2216
 - event preferences 2218
 - home page 2217
 - intrusion policies 1985
 - time zone 2222
 - widgets 73
- prefix lengths 57
- preprocessor rules 719
 - setting rule states 743
- preprocessors
 - adaptive profiles 1004
 - application layer 809
 - Back Orifice detection 959
 - DCE/RPC 810
 - DNS 827
 - events 784
 - execution order 782
 - FTP Telnet 832
 - generator IDs 784
 - IMAP preprocessor 880
 - inline normalization 918
 - introduction 773, 1015
 - intrusion policies 780

- IP defragmentation 927
 - order of execution 782
 - performance monitor 1026
 - POP preprocessor 884
 - processing packets 585
 - reading events 784
 - sensitive data detection 983
 - SIP preprocessor 872
 - SSH preprocessor 899
 - SSL preprocessor 905
 - primary Defense Center 195, 199
 - priority (rule keyword) 1058
 - process status monitoring 2114, 2137
 - processing packets 585
 - product licensing widget 101
 - product updates widget 101
 - profile conditions (in traffic profiles) 1596
 - profiling time window (PTW) 1585, 1593
 - protocol breakdown 1382
 - protocols
 - host profile 1359
 - server details 1348
 - white lists 1557
 - PTW 1585, 1593
 - purging the discovery database 2236
- R**
- RADIUS 1885
 - connection settings 1886
 - creating authentication objects 1885
 - custom attributes 1891
 - editing authentication objects 1896
 - examples 1894
 - shell access 1891
 - testing authentication 1893
 - user roles 1888
 - rate-based attack prevention
 - configuring 981
 - dynamic rule states 756
 - intrusion rules 727
 - policywide setting 971
 - thresholding (dynamic rule states) 976
 - thresholding (policy-wide) 978
 - understanding 971
 - with detection_filter keyword 975
 - with other filters 974
 - rawbytes content option (rule keyword) 1065
 - react (rule keyword) 1162
 - reassembly threshold (DCE/RPC option) 812
 - redundant Defense Centers 193
 - reference (rule keyword) 1062
 - refreshing, Context Explorer 146
 - registration ID 201
 - registration key 193, 201
 - regular expressions 1028
 - regular expressions, in intrusion rules 1086
 - remediation events 1633
 - field descriptions 1636
 - searching 1638
 - understanding 1636
 - viewing 1633
 - remediations 50, 1606, 1607
 - adding a Cisco IOS instance 1609
 - Cisco IOS routers 1608
 - Cisco PIX 1617
 - deleting modules 1608
 - installing new modules 1608
 - Nmap remediations 1705, 1712
 - status events 1633
 - remote (vulnerability details) 1366, 1437
 - remote access 2025
 - lights out management 2028, 2030
 - serial over LAN 2029
 - remote backups 2209
 - remote management 213
 - changing the management port 217
 - editing 216
 - remote storage 2017
 - local storage 2018
 - NFS 2019
 - reporting 1765
 - SMG 2022
 - SSH 2020
 - replace rules 1078
 - report designer 1740
 - report templates
 - creating 1740
 - deleting 1768
 - introduction 1736
 - reports
 - clipboard reports 652
 - cover pages, in reports 1759
 - creating report templates 1740
 - deleting 1769
 - document attributes 1756
 - downloading 1768
 - emailing reports automatically 1763
 - event graphs 601
 - from event views 1725
 - generating reports 1763
 - incidents 664
 - input parameters 1750

- introduction 1724
 - intrusion policy, comparing 685, 687
 - intrusion policy, using 682
 - logos 1760
 - managing report templates 1766
 - remote storage 1765
 - report templates 1736
 - scheduling reports 1763, 1940
 - require TCP handshake (stream option) 948
 - reset connection 1210
 - reset_dest (resp keyword) 1161
 - reset_source (resp keyword) 1161
 - resp (rule keyword) 1161
 - response groups 1510
 - activating and deactivating 1512
 - creating 1511
 - deleting 1512
 - editing 1512
 - restarting the appliance 2014
 - restoring from backup files 2211
 - retrospective malware events 1202, 1211, 1229
 - alerting 537
 - understanding 1230
 - right-click menus 63, 166
 - adding to malware whitelist 182
 - RIP (routing information protocol)
 - adding export filters 322
 - adding import filters 320
 - adding interfaces 316
 - advanced settings 319
 - authentication settings 318
 - configuring 316
 - risk
 - applications 1426
 - in application detectors 1666
 - root rules (access control) 451
 - routed interfaces
 - introduction 296
 - logical interfaces 300
 - physical interfaces 296
 - routing information protocol (RIP) 316
 - rpc (rule keyword) 1116
 - RPC decoder
 - configuring 869
 - generator ID 785
 - RPC proxy traffic only (DCE/RPC option) 820
 - RRD server process monitoring 2114, 2138
 - RSS feed widget 53, 103
 - rst_both (resp keyword) 1162
 - rule categories 739
 - rule classification 1058
 - rule comments (in an intrusion policy) 729
 - rule details 723
 - rule editor 1182
 - rule headers 1049, 1055
 - building 1049
 - specifying IPs 1051
 - specifying ports 1053
 - types 1050
 - rule latency
 - configuring thresholding 1025
 - GID 786
 - understanding thresholding 1021
 - rule processing configuration 1030
 - rule states, changing 743
 - rule updates 694, 2072
 - automatic one-time imports 2076
 - detail view 2088
 - import logging 2082, 2084
 - manual one-time imports 2074
 - recurring imports 2077
 - searching import logs 2089
 - rules
 - categories 739
 - creating 1182
 - directional operators 1055
 - filtering (in an intrusion policy) 729
 - keywords 1055
 - managing in an intrusion policy 717
 - setting rule state 743
 - tuning 591
 - Rules page, intrusion policy 795
- ## S
- sameip (rule keyword) 1155
 - saved searches 1774
 - deleting 1774
 - loading 1774
 - SCADA preprocessor 909
 - scan results in host profiles 1375
 - scan targets 1714
 - deleting scan targets 1716
 - editing scan targets 1715
 - scanning 1692
 - analyzing results 1719
 - importing scan results 1720
 - monitoring scans 1719
 - Nmap 1693, 1702
 - scan results 1716
 - searching for scan results 1721
 - scheduling tasks
 - automatically applying intrusion policies 1938

- automating URL filtering updates 1955
- backups 1932
- CRL downloads 1934
- deleting 1959
- editing 1959
- FireSIGHT recommended rules 1943
- introduction 1929
- Nmap scans 1936
- recurring tasks 1930
- reports 1940
- software downloads 1946
- software installs 1949
- software pushes 1948
- software updates 1945
- using the calendar 1957
- using the task list 1958
- VDB downloads 1952
- VDB installs 1953
- VDB updates 1951
- viewing 1957
- searching
 - active scan results 1721
 - application filters 1775
 - applications 1427, 1432
 - audit records 2197
 - connection data 575
 - correlation events 1526
 - custom tables 1788
 - discovery events 1399
 - events 1770, 1771
 - file events 1227
 - health events 2184
 - host attributes 1416
 - hosts 1408
 - intrusion events 644
 - intrusion rules 1189
 - loading a saved search 1774
 - malware events 1239
 - remediation events 1638
 - rule update import logs 2089
 - servers 1422
 - specifying ports 1777
 - third-party vulnerabilities 1444
 - time constraints 1775
 - user activity events 1456
 - vulnerabilities 1439
 - white list violations 1582
 - white lists 1576
 - wildcards 1775
- secondary Defense Center 195, 199
- SecurID 59, 61
- Security Analyst 57
- Security Approver 57
- Security Intelligence
 - access control policies 438
 - blacklisted host icons 1814
 - creating custom lists 170
 - high availability 164, 198
 - lists and feeds 162, 439
 - logging 440, 445, 520, 544, 550
 - monitoring 441, 2114, 2139
 - objects 162, 441
 - Sourcefire Intelligence Feed 167
- Security Over Connectivity 693
- security zones
 - access control rules 492
 - Context Explorer 134
 - introduction 158
 - NAT rules 415
 - Security Intelligence filtering 440
 - working with 187
- sensitive data
 - configuring custom data types 999
 - configuring detection 989
 - data type options 987
 - defining custom data patterns 996
 - deploying 985
 - detecting 983
 - detecting in FTP traffic (special case) 995
 - generator ID 786
 - global detection options 985
 - monitoring application protocols 993
 - predefined data types 988
- seq (rule keyword) 1110
- server fingerprints 1655
- server identities
 - editing 1352
 - resolving conflicts 1353
- server ports (stream reassembly option) 951
- server services (stream reassembly option) 951
- servers
 - host profiles 1347
 - introduction 1418
 - searching 1422
 - source types 1421
 - understanding 1419
 - viewing 1322, 1418, 1419
- serving time to managed devices 1993
- session logouts 59, 62
- sessions, current sessions widget 78
- set host criticality (discovery event type) 1476
- set operating system definition (discovery event type) 1396, 1476
- set server definition (discovery event type) 1475
- set service definition (discovery event type) 1396
- shared host profiles 1564

- adding 1558
 - creating 1565
 - deleting 1571
 - editing 1567
- shared object rules 719, 1046
 - generator ID 785
- shell access 61, 1878, 1891
- shutting down the appliance 2014
- SIDs
 - mapped to vulnerabilities 1437
 - network map 1323
 - vulnerability details 1365
- simultaneous connections, excessive 974
- SIP preprocessor 872
 - additional rules 876
 - configuring 875
 - generator ID 786
 - options 873
- sip_body (rule keyword) 1125
- sip_header (rule keyword) 1125
- sip_method (rule keyword) 1126
- sip_stat_code (rule keyword) 1127
- sliding time window 2219
- small segment size (stream option) 948
- SMB invalid shares (DCE/RPC option) 819
- SMB maximum AndX chain (DCE/RPC option) 819
- SMTP decoder, generator ID 786
- SMTP decoding 889
 - configuring 895
 - maximum memory alerting 899
 - options 890
- SNMP alerting 528, 1034
 - configuring 1037
 - intrusion policies 761
 - intrusion rules 728
- SNMP polling 1988
- snooze periods 1498
- Snort ID, *see* SID
- software updates 2054
 - installing 2056
 - scheduling downloads 1946
 - scheduling installs 1949
 - scheduling pushes 1948
 - scheduling updates 1945
- solution (vulnerability details) 1366
- source type
 - host profiles 1342
 - server details 1351
- Sourcefire cloud
 - enabling communications 2032
 - FireAMP connection 1219
- Sourcefire redundancy protocol (SFRP) 304
- Sourcefire Vulnerability ID (SVID) 1437, 1443
 - vulnerability details 1365
- SSH preprocessor 899
 - configuring 903
 - generator ID 786
 - options 901
- SSL clients 1269
- SSL preprocessor 905
 - additional rules 907
 - configuring 907
 - understanding 905
- SSL rule keywords 1113
- ssl_state (rule keyword) 1113
- ssl_version (rule keyword) 1115
- stacking
 - cabling 239
 - configuring an individual device 242
 - configuring interfaces 243
 - editing a device stack 241
 - establishing a device stack 239
 - eStreamer 238
 - licensing 2045
 - managed devices 237
 - primary devices 238
 - secondary devices 238
 - separating stacked devices 243
 - traffic handling 238
- standard rules (access control) 451
- standard text rules 1046, 1181
 - definition 719
 - generator ID 785
- stateful inspection anomalies (stream option) 947
- static routes 313
 - adding 314
 - viewing 313
- static time window 2219
- statistics
 - events 595
 - host 596
 - performance 598
- status, task queues 2238
- stream preprocessor
 - configuring TCP stream preprocessing 952
 - configuring UDP session tracking 957
 - generator ID 786
 - state-related TCP exploits 941
 - stream reassembly options 950
 - target-based policies 943
 - TCP global options 943
 - TCP policy options 945
 - TCP session tracking 940
 - UDP session tracking 956
- stream reassembly 949
 - generator ID (stream) 786

- reassembly options (stream) 950
- stream-based attacks 950
- stream_reassemble (rule keyword) 1112
- stream_size (rule keyword) 1111
- summary dashboard 66
- Sun directory server 1300
- suppression
 - configuring 753
 - creating 753
 - deleting 755
 - in the packet view 633
 - introduction 746
- switched interfaces 282
 - logical interfaces 285
 - physical interfaces 283
- switches 281
- SYN attacks, preventing 974
- syslog
 - alerting 530, 1040
 - filer examples 2202
 - filter syntax 2201
 - priority levels 1040
 - severity levels 532
 - syslog facilities 531, 1039
 - viewing 2200
- system daemons 2103
- system default variables 699
- system load widget 104
- system management
 - backup and restore 2204
 - IPS performance statistics 598
 - local settings 1998
 - software updates 2054
 - system policies 1961
 - uninstalling software updates 2068
 - VDB updates 2054
- system monitoring
 - disk usage 2099
 - host statistics 2096
 - introduction 2095
 - system load widget 104
 - system processes 2100
 - system status 2099
- system policies 1961
 - access control policy preferences 1970
 - access list 1971
 - applying 1965
 - audit log streaming 1973
 - authentication profiles 1975
 - creating 1962
 - custom login banner 1987
 - dashboard settings 1977
 - database limits 1978

- deleting 1969
- DNS cache 1981
- editing 1963
- language 1986
- mail relay host 1983
- SNMP polling 1988
- time sync 1991
- user interface settings 1994
- vulnerability mapping 1996
- system processes
 - understanding 2103
 - viewing 2100
- system time widget 104
- system utilities and executables 2105
- systems (host profile) 1343

T

- T/TCP options 938
- tab obfuscation (HTTP Inspect option) 864
- tab URI delimiter (HTTP Inspect option) 864
- table view of events 1792
 - features 1818
 - intrusion events 613
- tag (rule keyword) 1166
- tagged packets
 - generator ID 785
 - tag keyword 1166
- tags
 - application details 1431
 - application detectors 1666
 - applications 1426
 - servers 1420
- tap mode 275
- task queue 2238
 - managing 2240
 - viewing 2238
- TCP 585
 - experimental options 936
 - invalid options 938
 - new TCP port (discovery event type) 1393, 1472
 - obsolete options 937
 - port closed (discovery event type) 1393
 - port timeout (discovery event type) 1393
 - server information update (discovery event type) 1393
 - session tracking (stream) 940
 - T/TCP options 938
 - TCP header values 1106
 - TCP port closed (discovery event type) 1471

- TCP port timeout (discovery event type) 1471
 - TCP ports (discovery event type) 1393
 - TCP server confidence update (discovery event type) 1393
 - TCP server information update (discovery event type) 1472
 - technical description (vulnerability details) 1366
 - Telnet decoder
 - configuring 832
 - configuring global options 833
 - configuring options 836
 - generator ID 786
 - global options 832
 - understanding options 835
 - Teredo traffic, detecting (packet decoder) 585, 935
 - testing authentication 1879
 - text (host attribute type) 1371
 - third-party product mappings 1682, 1687
 - third-party tools 1682
 - identifying hosts 1643
 - identifying services 1688
 - patch tools 1685
 - vulnerabilities 1644, 1687
 - threshold (deprecated rule keyword) 1166
 - thresholding
 - global rule 1009
 - global rule (configuring) 1012
 - global rule (disabling) 1014
 - intrusion rule (adding from packet view) 631
 - intrusion rule (adding) 749
 - intrusion rule (configuring) 747
 - intrusion rule (deleting) 751
 - intrusion rules 746
 - packet latency 1017
 - packet latency (configuring) 1020
 - rate-based 975
 - rate-based (configuring) 981
 - rate-based (dynamic rule states) 976
 - rate-based (multiple filters) 979
 - rate-based (policy-wide) 978
 - rule latency 1021
 - rule latency (configuring) 1025
 - time constraints (in searches) 1775
 - time series data monitoring 2114, 2140
 - time sync 1991, 1993
 - monitoring 2114, 2141
 - servicing time 1993
 - setting time manually 2015
 - system time widget 104
 - time window 1821, 2219
 - changing the default 1826
 - configuring 1820
 - Context Explorer 147
 - default setting 2219
 - pausing 1829
 - time zone 2222
 - timeout
 - defragmentation 931
 - DNS cache option 1982
 - host timeout 1391, 1471
 - session logouts 59, 62
 - TCP policy option (stream) 946
 - TCP port timeout (discovery event type) 1393, 1471
 - UDP port timeout (discovery event type) 1394, 1471
 - UDP session tracking (stream) 958
 - timestamp
 - host view 1340, 1403
 - server view 1420
 - tos (rule keyword) 1103
 - traffic profiles
 - activating 1595
 - creating 1585
 - editing 1596
 - host profile qualifications 1590
 - host view 1407
 - profile conditions 1588, 1596
 - saving 1595
 - specifying conditions 1588
 - viewing 1603
 - traffic status monitoring 2114, 2142
 - Transmission Control Protocol, *see* TCP
 - transparent inline mode 276
 - transport layer 585
 - preprocessor 940
 - viewing information 638
 - transport protocol (discovery event type) 1393, 1471
 - trend graphs 601
 - troubleshooting intrusion policies 790
 - TTL
 - host profile 1339
 - rule keyword 1103
- U**
- UDP 585
 - and Back Orifice 959
 - new port (discovery event type) 1393
 - packet view 640
 - port closed (discovery event type) 1394
 - port timeout (event type) 1394
 - session tracking 956

- UDP port closed (discovery event type) 1471
- UDP port timeout (discovery event type) 1471
- UDP server confidence update (discovery event type) 1394
- UDP server information update (discovery event type) 1394
- UDP service information update (discovery event type) 1472
- unauthorized activity 60
- unauthorized widgets 71
- uninstalling software updates 2068
- unique NAT ID 193
- unreviewing events 611
- update interval, discovery policy 1288
- updating
 - Defense Centers 2062
 - intrusion rule updates 2072
 - managed devices 2066
 - system software 2054
 - vulnerability database (VDB) 2070
- URI parsing (HTTP Inspect option) 865
- urilen (rule keyword) 1118
- URL (host attribute type) 1371
- URL filtering
 - access control rules 506
 - automating updates 1955
 - high availability 198
 - licensing 2043
 - monitoring 2115
 - URL statistics dashboard 68
- URL Filtering Monitoring 2142
- URL information
 - Context Explorer 143, 144, 145
 - URL length 447
- URL objects 175
- URL reputations, Context Explorer 145
- URL statistics dashboard 68
- user accounts
 - browser timeout exemption 1905
 - command line access 1901
 - deleting 1914
 - editing 1912
 - externally authenticated user accounts 1902
 - LDAP users 1902
 - login settings 1903
 - managing 1849, 1897
 - menu access per access type 1914
 - password options 1904
 - privileges 1852
 - shell access 1878, 1891
 - user authentication 1849
- user activity
 - event searches 1456
 - user activity events 1456
- user activity events 1453, 1455
 - viewing 1454
- user agents
 - configuring 1307
 - high availability 199
 - installing 1310
 - monitoring 2115, 2143
 - understanding 1255
- user authentication 1849
 - external authentication 1851
 - internal authentication 1851
- user awareness
 - Context Explorer 133
 - LDAP 1298
 - searching 1456
 - user activity events 1453, 1455
 - user agents 1307
 - user identity events 1445, 1447, 1449, 1451
- user data collection 1252
- User Datagram Protocol, *see* UDP
- user history 1357
- user identity
 - correlation rules 1496
 - FireAMP users 1814
 - in the host view 1404
 - login types 1446
 - user qualification 1496
- user identity events 1445, 1447, 1449
 - searching 1451
- user preferences 2215
 - changing event preferences 2218
 - changing passwords 2216
 - home page 2217
 - time zone 2222
- user roles
 - access conventions 56
 - activating and deactivating 1908
 - custom user role escalation 1926
 - escalating 1925
 - escalation targets 1926
 - predefined 1906
 - using with custom user roles 1911
- user statistics dashboard 68
- users
 - configuring access control 1303
 - database 1260
 - restricting logging in discovery policy 1285
- UTF-8 (HTTP Inspect option) 862

V

- variables
 - creating 703
 - custom variables 710
 - deleting 709
 - editing 701
 - predefined 699
 - system default 699
 - using IP addresses 711
- VDB updates 2070
 - scheduling downloads 1952
 - scheduling installs 1953
 - scheduling updates 1951
- version
 - operating systems 1342
 - server details 1349
 - servers 1351
- violations (white lists) 1579
- virtual routers
 - adding 307
 - authentication profiles 339
 - configuring 306
 - deleting 341
 - DHCP relay 310
 - dynamic routing 316
 - filtering 335
 - hybrid interfaces 352
 - introduction 295
 - NAT 341
 - routed interfaces 296
 - static routes 313
 - statistics 340
 - strict TCP enforcement 307
 - viewing 307
- virtual switches 281
 - adding 289
 - advanced settings 291
 - configuring 288
 - deleting 294
 - hybrid interfaces 352
 - strict TCP enforcement 291
 - switched interfaces 282
 - viewing 288
- VLAN ID (host view) 1404
- VLAN packets (decoded) 585
- VLAN tag objects 174
- VLAN tags 300
 - access control rules 490, 495
 - untagged traffic 296
 - VLAN tag objects 174
- VLANs
 - host profiles 1357
 - ignore VLAN headers (intrusion policy detection setting) 917
 - tag information update (discovery event type) 1394, 1471
 - VLAN tag objects 174
- VPN
 - comparing deployments 381
 - configuring 363
 - configuring advanced settings 374
 - introduction 358
 - managing 362
 - monitoring 2115, 2144
 - understanding deployments 360
 - understanding IKE 359
 - understanding IPSec 359
 - understanding mesh deployments 361
 - understanding point-to-point deployments 360
 - understanding star deployments 360
 - viewing logs 378
 - viewing statistics 378
 - viewing status 377
- vulnerabilities
 - activating and deactivating 1367, 1368
 - deactivating 1323, 1438
 - host profiles 1363
 - impact assessment mappings 1291
 - impact qualification 1367
 - impacts 1366
 - introduction 1434
 - invalidating 1323
 - mapped to SIDs 1437
 - misidentified 1644
 - network map 1323
 - searching 1439
 - third-party tools 1644, 1687
 - understanding 1436
 - viewing 1435
 - vulnerability details 1365
- vulnerabilities, third-party
 - host profiles 1363
 - introduction 1441
 - searching 1444
 - understanding 1442
 - viewing 1441
- vulnerability impact qualification (discovery event type) 1396
- vulnerability mappings 1352, 1353, 1654, 1659, 1684, 1686
 - servers 1996
- vulnerability set invalid (discovery event type) 1396, 1476
- vulnerability set valid (discovery event type) 1396, 1476

W

- web applications 1421
 - application details 1431
 - connection trackers 1493
 - correlation rules 1474, 1478
 - detecting 1390
 - host profile 1354
 - host profile qualifications 1484
 - traffic profiles 1590
- web browser requirements 58
- webroot directory traversal (HTTP Inspect option) 864
- white list events
 - viewing 1573
 - working with 1572
- white list targets 1533
 - configuring network segments 1545
 - deleting 1548
 - editing 1548
- white lists 1565
 - adding to correlation policies 1515
 - allow jailbreaking 1542, 1545
 - application protocols 1552
 - basic information 1545
 - clients 1554, 1556
 - configuring targets 1545
 - creating 1541
 - creating from host profiles 1361
 - definition 1530
 - deleting 1564
 - deleting host profiles 1562
 - editing 1563
 - events 1572
 - field descriptions 1575, 1581
 - global host profiles 1535, 1549
 - host profiles 1534, 1548
 - host view 1408
 - introduction 1530
 - managing 1563
 - modifying host profiles 1559
 - network compliance widget 99
 - operating systems 1535, 1550
 - protocols 1557
 - searching 1576
 - Security Intelligence 163, 166, 439, 441
 - shared host profiles 1537, 1558, 1564
 - surveying your network 1543
 - target evaluations 1538
 - understanding 1532, 1575
 - violations 1360, 1539, 1579, 1582
 - white list events 1539
 - white list events widget 105
 - white list violations 1581
- widgets 70
 - adding to a dashboard 113
 - appliance status 76
 - availability 71
 - correlation events 77
 - current sessions 78
 - custom analysis 79
 - deleting 116
 - disk usage 96
 - interface status 77
 - interface traffic widget 97
 - intrusion events 97
 - invalid widgets 71
 - minimizing and maximizing 115
 - moving 115
 - network compliance 99
 - predefined 71, 74
 - preferences 73
 - product licensing 101
 - product updates 101
 - RSS feeds 53, 103
 - system load 104
 - system time 104
 - unauthorized widgets 71
 - white list events 105
- wildcards (in searches) 1775
- window (rule keyword) 1110
- within content option (rule keyword) 1069
- workflows
 - common functionality 1813
 - components 1792
 - compound constraints 1832
 - connection graphs 563, 564
 - constraining events 1830
 - constraining time 1820
 - creating 1840
 - custom 1791
 - custom connection data workflows 1843
 - custom tables 1787
 - default workflows 2221
 - deleting 1847
 - discovery events 1386
 - editing 1847
 - introduction 1791
 - intrusion events 612
 - navigating between workflows 1836
 - navigating to other pages 1836
 - predefined 1791, 1794
 - predefined vs.custom tables 1794
 - predefined vs.custom workflows 1794
 - saved custom workflows 1807
 - selecting workflows 1809
 - sorting pages 1834, 1835

- table view of events 1818
- toolbar 1812
- using bookmarks 1838
- using workflows 1808
- viewing custom workflows 1846
- would have dropped (event type) 689