Sourcefire 3D System

Sourcefire 3D System
User Guide



SOURCEfire
Security for the real world.

Version 4.10

Terms of Use Applicable to the User Documentation

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 Sourcefire, Inc. or its subsidiaries (collectively, "Sourcefire") or any Sourcefire-provided products. Sourcefire products are available for purchase and subject to a separate license agreement and/or terms of use containing very different terms and conditions.

Terms of Use and Copyright and Trademark Notices

The copyright in the Documentation is owned by Sourcefire 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 Sourcefire'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 Sourcefire. Sourcefire 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.

SOURCEFIRE®, SNORT®, CLAMAV®, SOURCEFIRE DEFENSE CENTER®, SOURCEFIRE 3D®, RNA®, RUA®, SECURITY FOR THE REAL WORLD®, the Sourcefire logo, the Snort and Pig logo, the ClamAV logo, Sourcefire IPS, RAZORBACK, Sourcefire Master Defense Center, DAEMONLOGGER, and certain other trademarks and logos are trademarks or registered trademarks of Sourcefire, Inc. in the United States and other countries. Other company, product and service names may be trademarks or service marks of others.

© 2004 - 2011 Sourcefire, Inc. All rights reserved.

Disclaimers

THE DOCUMENTATION AND ANY INFORMATION AVAILABLE FROM IT MAY INCLUDE INACCURACIES OR TYPOGRAPHICAL ERRORS. SOURCEFIRE MAY CHANGE THE DOCUMENTATION FROM TIME TO TIME. SOURCEFIRE MAKES NO REPRESENTATIONS OR WARRANTIES ABOUT THE ACCURACY OR SUITABILITY OF ANY SOURCEFIRE-CONTROLLED WEBSITE, THE DOCUMENTATION AND/OR ANY PRODUCT INFORMATION. SOURCEFIRE-CONTROLLED WEBSITES, THE DOCUMENTATION AND ALL PRODUCT INFORMATION ARE PROVIDED "AS IS" AND SOURCEFIRE 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 SOURCEFIRE 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 SOURCEFIRE-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 SOURCEFIRE 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.

The Documentation may contain "links" to websites that are not created by, or under the control of Sourcefire. Sourcefire provides such links solely for your convenience, and assumes no responsibility for the availability or content of such other sites.

2011-Mar-16 14:59

Table of Contents

Chapter 1: Introduction to the Sourcefire 3D System 32

Components of the Sourcefire 3D System 33
- 3D Sensors 33
- Defense Centers 34
- IPS 34
- Sourcefire RNA 35
- Sourcefire RUA 36
- PEP Traffic Management 36
- Master Defense Centers 36
- Intrusion Agents 37
- RNA and IPS for Crossbeam Systems 37
- eStreamer 37

Logging into the Appliance 37

Logging into the Appliance to Set Up an Account 40

Logging Out of the Appliance 42

Last Successful Login 43

Specifying Your User Preferences 43
- Changing Your Password 43
- Configuring Event View Settings 45
- Setting Your Default Time Zone 50
- Specifying Your Home Page 51
- Specifying Your Default Dashboard 52

Using the Context Menu 53

Communication Ports 54

Documentation Resources 55

Keysight, EX-1004
Page 3

Documentation Conventions 55
Platform Conventions 56
Component Conventions 56
Access Conventions 57
IP Address Conventions 59

Chapter 2: Using Dashboards 61

Understanding Dashboard Widgets 62
Understanding Widget Availability 63
Understanding Widget Preferences 66
Understanding the Predefined Widgets 67
Understanding the Appliance Information Widget 68
Understanding the Appliance Status Widget 69
Understanding the Compliance Events Widget 69
Understanding the Current Interface Status Widget 70
Understanding the Current Sessions Widget 71
Understanding the Custom Analysis Widget 71
Understanding the Disk Usage Widget 82
Understanding the Interface Traffic Widget 83
Understanding the Intrusion Events Widget 83
Understanding the Network Compliance Widget 85
Understanding the Product Licensing Widget 87
Understanding the Product Updates Widget 88
Understanding the RSS Feed Widget 89
Understanding the System Load Widget 90
Understanding the System Time Widget 90
Understanding the White List Events Widget 91
Working with Dashboards 92
Creating a Custom Dashboard 92
Viewing Dashboards 94
Modifying Dashboards 96
Deleting a Dashboard 102

Chapter 3: Using the Defense Center 103

Management Concepts 104
Understanding the Benefits of Managing Your Sensors 105
What Can Be Managed by a Defense Center? 106
Understanding Software Sensors 110
Beyond Policies and Events 116
Using Redundant Defense Centers 117
Working in NAT Environments 117

Keysight, EX-1004
Page 4

Working with Sensors 118
Understanding the Sensors Page 120
Adding Sensors to the Defense Center 122
Deleting Sensors 126
Resetting Management of a Sensor 127
Managing a 3Dx800 Sensor 131
Adding Intrusion Agents 136
Sensor Attributes - Intrusion Agent Page 136
Managing Sensor Groups 137
Creating Sensor Groups 137
Editing Sensor Groups 138
Deleting Sensor Groups 139
Editing a Managed Sensor's System Settings 139
Viewing a Sensor's Information Page 141
Stopping and Restarting a Managed Sensor 143
Managing Communication on a Managed Sensor 144
Setting the Time on a Managed Sensor 145
Managing a Clustered Pair 146
Establishing a Clustered Pair 149
Separating a Clustered Pair 151
Configuring High Availability 151
Using High Availability 152
Guidelines for Implementing High Availability 157
Setting Up High Availability 158
Monitoring the High Availability Status 160
Disabling High Availability and Unregistering Sensors 162
Pausing Communication between Paired Defense Centers 163
Restarting Communication between Paired Defense Centers 163

Chapter 4: Using the Master Defense Center 165

Understanding Event Aggregation 166
Aggregating Intrusion Events 167
Aggregating Compliance Events 168
Limitations on Event Aggregation 168
Understanding Global Policy Management 170
Managing Global Intrusion Policies 171
Using RNA Detection Policies on a Master Defense Center 171
Using Health Policies on a Master Defense Center 172
Using System Policies on a Master Defense Center 172
Master Defense Center Policy Management Limitations 173

Keysight, EX-1004
Page 5

Adding and Deleting Defense Centers 174
Adding a Master Defense Center 174
Adding a Defense Center 178
Deleting a Defense Center 181
Resetting Management of a Defense Center 182
Using the Appliances Page 184
Editing Settings for a Managed Defense Center 185
Viewing the Defense Center Information Page 186
Editing the Event Filter Configuration 187
Editing or Disabling Remote Management Communications 189
Managing the Health Blacklist 189
Managing High Availability Defense Centers 190
Managing Appliance Groups 191
Creating Appliance Groups 191
Editing Appliance Groups 192
Deleting Appliance Groups 193
Editing Master Defense Center System Settings 193
Listing Master Defense Center Information 193
Viewing a Master Defense Center License 194
Configuring Network Settings 194
Shutting Down and Restarting the System 195
Configuring Remote Management Networking 195
Setting System Time 196
Blacklisting Health Policies 196

Chapter 5: Using Detection Engines and Interface Sets 198

Understanding Detection Engines 199
Understanding Detection Resources and 3D Sensor Models 202
Understanding Default Detection Engines 205
Managing Detection Engines 206
Creating a Detection Engine 206
Editing a Detection Engine 208
Deleting a Detection Engine 211
Using Detection Engine Groups 211
Creating Detection Engine Groups 211
Editing Detection Engine Groups 212
Deleting Detection Engine Groups 213
Using Variables within Detection Engines 213
Assigning Values to System Default Variables in Detection Engines 215
Creating New Variables for Detection Engines 217
Deleting and Resetting Variables 218
Configuring Custom Variables in Detection Engines 219
Using Portscan-Only Detection Engines 220

Keysight, EX-1004
Page 6

Using Interface Sets 222
Understanding Interface Set Configuration Options 223
Creating an Interface Set 230
Creating an Inline Interface Set 234
Editing an Interface Set 239
Editing an 8000 Series Interface Set 241
Deleting an Interface Set 244
Using Interface Set Groups 244
Creating Interface Set Groups 244
Editing Interface Set Groups 245
Deleting Interface Set Groups 246
Inline Fail Open Interface Set Commands 246
Removing Bypass Mode on Inline Fail Open Fiber Interfaces 246
Forcing an Inline Fail Open Interface Set into Bypass Mode 247
Using Clustered 3D Sensors 250
Identifying Detection Engines and Interface Sets on a Clustered Pair 251

Chapter 6: Introduction to Sourcefire IPS 253

Understanding How Traffic Is Analyzed 256
Capturing and Decoding Packets 257
Processing Packets 258
Generating Events 260
Analyzing Intrusion Event Data 261
Using Intrusion Event Responses 261
Understanding IPS Deployments 262
Using IPS Software Sensors 265
The Benefits of Custom Intrusion Policies 265

Chapter 7: Working with Intrusion Events 267

Viewing Intrusion Event Statistics 269
Host Statistics 271
Event Overview 271
Event Statistics 272
Viewing Intrusion Event Graphs 273
Viewing Intrusion Events 274
Understanding Intrusion Events 276
Reviewing Intrusion Events 278
Understanding Workflow Pages for Intrusion Events 280
Using Drill-Down and Table View Pages 284

Keysight, EX-1004
Page 7

Using the Packet View 290
Viewing Event Information 293
Viewing Frame Information 299
Viewing Data Link Layer Information 300
Viewing Network Layer Information 301
Viewing Transport Layer Information 303
Viewing Packet Byte Information 306
Using Impact Flags to Evaluate Events 307
Searching for Intrusion Events 309
Specifying Rule Classifications in Event Searches 314
Using the Clipboard 316
Generating Clipboard Reports 316
Deleting Events from the Clipboard 318

Chapter 8: Handling Incidents 319

Incident Handling Basics 320
Definition of an Incident 320
Common Incident Handling Processes 321
Incident Types in the Sourcefire 3D System 324
Creating an Incident 324
Editing an Incident 326
Generating Incident Reports 327
Creating Custom Incident Types 329

Chapter 9: Creating Intrusion Policies 331

Planning and Implementing an Intrusion Policy 333
Managing Intrusion Policies 335
Creating an Intrusion Policy 338
Editing an Intrusion Policy 341
Using the Navigation Panel 344
Committing Intrusion Policy Changes 345
Applying an Intrusion Policy 347
Viewing an Intrusion Policy Report 350
Comparing Two Intrusion Policies 354
Setting Drop Behavior in an Inline Deployment 358
Understanding the Base Policy 361
Using Default Intrusion Policies 362
Using a Custom Base Policy 363
Allowing SEUs to Modify the Base Policy 363
Selecting the Base Policy 364
Accepting Rule Setting Changes from a Custom Base Policy 366

Keysight, EX-1004
Page 8

Managing Targets 368
Managing Variables 371
Understanding Existing Variables 372
Modifying Variables 374
Creating New Variables 377
Deleting Unused Variables 389
Understanding Custom Variables 391
Defining VLANs and Subnetworks 393
Limiting Variables to VLANs and Subnetworks 398
Defining IP Addresses and Ports for Your Network 398
Defining IP Addresses in Variables and Rules 399
Defining Ports in Variables and Rules 401

Chapter 10: Managing Rules in an Intrusion Policy 403

Understanding IPS Rule Types 404
Viewing Rules in an Intrusion Policy 406
Sorting the Rule Display 409
Viewing Rule Details 410
Filtering Rules in an Intrusion Policy 417
Understanding Rule Filtering in an Intrusion Policy 419
Setting a Rule Filter in an Intrusion Policy 432
Setting Rule States 435
Filtering Intrusion Event Notification Per Policy 439
Configuring Event Thresholding 439
Configuring Suppression Per Intrusion Policy 447
Adding Dynamic Rule States 452
Understanding Dynamic Rule States 453
Setting a Dynamic Rule State 454
Adding Alerts 458
Adding SNMP Alerts 458
Adding OPSEC Alerts 460
Adding Rule Comments 466
Managing RNA Rule State Recommendations 468
Understanding Basic Rule State Recommendations 469
Understanding Advanced Rule State Recommendations 470
Using RNA Recommendations 472

Keysight, EX-1004
Page 9

Chapter 11: Using Advanced Settings in an Intrusion Policy 477

Modifying Advanced Settings 478

Understanding Preprocessors 484
- Meeting Traffic Challenges with Preprocessors 485
- Understanding Preprocessor Execution Order 486
- Reading Preprocessor Events 488

Automatically Enabling Advanced Settings 490

Understanding Troubleshooting Options 492

Chapter 12: Using Layers in an Intrusion Policy 494

Understanding Intrusion Policy Layers 494
- Sharing Layers 496
- Using Rules in Layers 497
- Removing Multi-Layer Rule Settings 499
- Using the RNA Recommendations Layer 503
- Using Layers with Advanced Settings 504

Configuring User Layers 507

Chapter 13: Using Application Layer Preprocessors 513

Decoding DCE/RPC Traffic 514
- Selecting Global DCE/RPC Options 515
- Understanding Target-Based DCE/RPC Server Policies 517
- Understanding DCE/RPC Transports 518
- Selecting DCE/RPC Target-Based Policy Options 522
- Configuring the DCE/RPC Preprocessor 526

Detecting Exploits in DNS Name Server Responses 531
- Understanding DNS Preprocessor Resource Record Inspection 532
- Detecting Overflow Attempts in RData Text Fields 534
- Detecting Obsolete DNS Resource Record Types 534
- Detecting Experimental DNS Resource Record Types 535
- Configuring the DNS Preprocessor 535

Decoding FTP and Telnet Traffic 538
- Understanding Global FTP and Telnet Options 538
- Configuring Global FTP/Telnet Options 539
- Understanding Telnet Options 542
- Configuring Telnet Options 543
- Understanding Server-Level FTP Options 546
- Configuring Server-Level FTP Options 550
- Understanding Client-Level FTP Options 555
- Configuring Client-Level FTP Options 557

Keysight, EX-1004
Page 10

Decoding HTTP Traffic 560
Selecting Global HTTP Normalization Options 562
Configuring Global HTTP Configuration Options 563
Selecting Server-Level HTTP Normalization Options 564
Selecting Server-Level HTTP Normalization Encoding Options 570
Configuring Server-Level HTTP Configuration Options 574
Enabling Additional HTTP Inspect Preprocessor Rules 577
Using the Sun RPC Preprocessor 578
Configuring the Sun RPC Preprocessor 579
Decoding SMTP Traffic 582
Understanding SMTP Decoding 582
Configuring SMTP Decoding 585
Detecting Exploits Using the SSH Preprocessor 590
Selecting SSH Preprocessor Options 591
Configuring the SSH Preprocessor 594
Using the SSL Preprocessor 596
Understanding SSL Preprocessing 597
Enabling SSL Preprocessor Rules 598
Configuring the SSL Preprocessor 599

Chapter 14: Using Transport & Network Layer Preprocessors 602

Verifying Checksums 602
Normalizing Inline Traffic 604
Understanding Protocol Normalization 605
Configuring Inline Normalization 607
Defragmenting IP Packets 613
Understanding IP Fragmentation Exploits 614
Target-Based Defragmentation Policies 615
Selecting Defragmentation Options 616
Configuring IP Defragmentation 618
Understanding Packet Decoding 621
Configuring Packet Decoding 625
Using TCP Stream Preprocessing 627
Understanding State-Related TCP Exploits 628
Initiating Active Responses with Drop Rules 629
Selecting TCP Global Options 630
Understanding Target-Based TCP Policies 631
Selecting TCP Policy Options 633
Reassembling TCP Streams 637
Configuring TCP Stream Preprocessing 640
Using UDP Stream Preprocessing 645
Configuring UDP Stream Preprocessing 645

Keysight, EX-1004
Page 11

Chapter 15: Detecting Specific Threats 648

Detecting Back Orifice 648

Detecting Portscans 651
- Configuring Portscan Detection 655
- Understanding Portscan Events 658

Preventing Rate-Based Attacks 662
- Understanding Rate-Based Attack Prevention 662
- Rate-Based Attack Prevention and Other Filters 666
- Configuring Rate-Based Attack Prevention 673

Detecting Sensitive Data 677
- Deploying Sensitive Data Detection 678
- Selecting Global Sensitive Data Detection Options 678
- Selecting Individual Data Type Options 680
- Using Predefined Data Types 681
- Configuring Sensitive Data Detection 683
- Selecting RNA Services to Monitor 686
- Special Case: Detecting Sensitive Data in FTP Traffic 689
- Using Custom Data Types 690

Chapter 16: Using Adaptive Profiles 697

Understanding Adaptive Profiles 698
- Using Adaptive Profiles with Preprocessors 698
- Using Adaptive Profiles with Intrusion Rules 700
- Adaptive Profiles and RNA Recommended Rules 702

Configuring Adaptive Profiles 702

Chapter 17: Using Global Rule Thresholding 705

Understanding Thresholding 705
- Understanding Thresholding Options 706

Configuring Global Thresholds 708
- Disabling the Global Threshold 711

Chapter 18: Using Performance Settings in an Intrusion Policy 713

Event Queue Configuration 714

Understanding Packet Latency Thresholding 716
- Setting Packet Latency Thresholding Options 718
- Configuring Packet Latency Thresholding 719

Understanding Rule Latency Thresholding 720
- Setting Rule Latency Thresholding Options 723
- Configuring Rule Latency Thresholding 724

Keysight, EX-1004
Page 12

Performance Statistics Configuration 726
Constraining Regular Expressions 728
Rule Processing Configuration 731

Chapter 19: Configuring External Responses to Intrusion Events 734

Using Check Point OPSEC Responses 735
Creating the OPSEC SAM Application 736
Connecting the 3D Sensor and the SAM Server 739
Selecting a SAM Server 743
Clearing All OPSEC SAM Responses for a Detection Engine 745
Adding and Deleting Firewall Objects 745
Viewing Sourcefire SAM Client System Information 746
Viewing OPSEC Debug Messages 747
Using SNMP Responses 748
Configuring SNMP Responses 749
Using Syslog Responses 752
Configuring Syslog Responses 754
Understanding Email Alerting 756
Configuring Email Alerting 759

Chapter 20: Understanding and Writing Intrusion Rules 761

Understanding Rule Anatomy 762
Understanding Rule Headers 764
Specifying Rule Actions 765
Specifying Protocols 766
Specifying Source and Destination IP Addresses 766
Specifying Source and Destination Ports 768
Specifying Direction 769
Understanding Keywords and Arguments in Rules 769
Defining Intrusion Event Details 771
Searching for Content Matches 776
Constraining Content Matches 778
Replacing Content in Inline Deployments 792
Using Byte_Jump and Byte_Test 794
Searching for Content Using PCRE 800
Adding Metadata to a Rule 810
Inspecting IP Header Values 812
Inspecting ICMP Header Values 816
Inspecting TCP Header Values and Stream Size 818

Keysight, EX-1004
Page 13

Enabling and Disabling TCP Stream Reassembly 824
Extracting SSL Information from a Session 825
Inspecting Application Layer Protocol Values 828
Inspecting Packet Characteristics 837
Reading Packet Data into Keyword Arguments 840
Initiating Active Responses with Rule Keywords 844
Filtering Events 849
Evaluating Post-Attack Traffic 850
Detecting Attacks That Span Multiple Packets 851
Generating Events on the HTTP Encoding Type and Location 855
Pointing to a Specific Payload Type 857
Decoding and Inspecting Base64 Data 859

Constructing a Rule 861
Writing New Rules 861
Modifying Existing Rules 864
Adding Comments to Rules 867
Deleting Custom Rules 868

Searching for Rules 869

Filtering Rules on the Rule Editor Page 872
Using Keywords in a Rule Filter 873
Using Character Strings in a Rule Filter 875
Combining Keywords and Character Strings in a Rule Filter 875
Filtering Rules 876

Chapter 21: Introduction to Sourcefire RNA 878

Understanding RNA 879
What Data Does RNA Collect? 880
What Information Does RNA Provide? 880
How Can You Use RNA? 886
What are the Prerequisites for RNA Data Collection? 887
What is an RNA Detection Policy? 889
How Do I Manage my RNA Deployment with Subnet Detection? 897

Creating RNA Detection Policies 900
Configuring RNA Detection Policy Settings 902
Specifying Networks to Monitor with 3D Sensors 905
Specifying Ports to Exclude 907
Specifying Networks to Monitor with NetFlow Devices 909

Managing RNA Detection Policies 910
Managing Subnet Detection 911
Applying an RNA Detection Policy 919
Editing an RNA Detection Policy 921
Deleting an RNA Detection Policy 921

Keysight, EX-1004
Page 14

Chapter 22: Using the Network Map 923

Understanding the Network Map 924
Working with the Hosts Network Map 926
Working with the Network Devices Network Map 928
Working with the Services Network Map 930
Working with the Vulnerabilities Network Map 932
Working with the Host Attributes Network Map 935
Working with Custom Network Topologies 936
- Creating Custom Topologies 937
- Managing Custom Topologies 942

Chapter 23: Using Host Profiles 944

Viewing Host Profiles 948
Working with Basic Host Information in the Host Profile 948
Working with Operating Systems in the Host Profile 951
- Viewing Operating System Identities 954
- Editing an Operating System 955
- Resolving Operating System Identity Conflicts 956

Working with Services in the Host Profile 958
- Service Detail 962
- Editing Service Identities 965
- Resolving Service Identity Conflicts 967

Working with Client Applications in the Host Profile 969
- Viewing Client Applications in the Host Profile 970
- Deleting Client Applications from the Host Profile 972

Working with VLAN Tags in the Host Profile 972
Working with User History in the Host Profile 973
Working with Host Attributes in the Host Profile 973
- Assigning Host Attribute Values 974

Working with Host Protocols in the Host Profile 974
Working with White List Violations in the Host Profile 976
- Creating a White List Host Profile from a Host Profile 977

Working with Vulnerabilities in the Host Profile 978
- Viewing Vulnerability Details 980
- Setting the Vulnerability Impact Qualification 982
- Downloading Patches for Vulnerabilities 983
- Setting Vulnerabilities for Individual Hosts 984

Working with the Predefined Host Attributes 985
- Assigning a Host Criticality Value to a Host 986
- Adding Notes to a Host 986

Keysight, EX-1004
Page 15

Working with User-Defined Host Attributes 987
Restrictions on User-Defined Host Attributes 988
Creating User-Defined Host Attributes 989
Editing a User-Defined Host Attribute 991
Deleting a User-Defined Host Attribute 991
Working with Scan Results in a Host Profile 992
Scanning a Host from the Host Profile 992

Chapter 24: Working with RNA Events 994

Viewing RNA Event Statistics 996
Statistics Summary 997
Event Breakdown 998
Protocol Breakdown 1000
Service Breakdown 1000
OS Breakdown 1001
Understanding RNA Event Workflows 1002
Working with Network Discovery and Host Input Events 1004
Understanding RNA Network Discovery Event Types 1005
Understanding RNA Host Input Event Types 1010
Viewing RNA Network Discovery and Host Input Events 1012
Understanding the RNA Events Table 1014
Searching for RNA Network Discovery Events 1016
Working with Hosts 1019
Viewing Hosts 1020
Understanding the Hosts Table 1022
Creating a Traffic Profile for Selected Hosts 1025
Creating a Compliance White List Based on Selected Hosts 1026
Locating Whois Information for a Host 1027
Searching for Hosts 1027
Working with Host Attributes 1033
Viewing Host Attributes 1034
Understanding the Host Attributes Table 1035
Setting Host Attributes for Specific Hosts 1036
Searching for Host Attributes 1038
Working with Services 1040
Viewing Services 1041
Understanding the Services Table 1043
Searching for Services 1045
Working with Client Applications 1050
Viewing Client Applications 1051
Understanding the Client Applications Table 1053
Searching for Client Applications 1054

Keysight, EX-1004
Page 16

Working with Sourcefire Vulnerabilities 1057
Viewing Sourcefire Vulnerabilities 1057
Understanding the Sourcefire Vulnerabilities Table 1059
Deactivating Sourcefire Vulnerabilities 1061
Searching for Sourcefire Vulnerabilities 1062
Working with Third-Party Vulnerabilities 1064
Viewing Third-Party Vulnerabilities 1065
Understanding the Third-Party Vulnerabilities Table 1067
Searching for Third-Party Vulnerabilities 1068

Chapter 25: Working with Flow Data and Traffic Profiles 1071

Understanding Flow Data 1072
Understanding Flow Summary Data 1076
Understanding Flow Data Graphs 1078
Understanding Flow Data Tables 1086
Viewing Flow Data 1093
Searching for Flow Data 1094
Manipulating Flow Data Graphs 1100
Viewing Aggregated Flow Data 1101
Manipulating a Flow Data Graph on a Workflow Page 1103
Drilling Down through a Flow Data Workflow 1105
Recentering and Zooming on Line Graphs 1108
Changing the Graph Type 1109
Selecting Data to Graph 1109
Selecting Datasets 1113
Detaching Flow Data Graphs 1114
Exporting Flow Data 1115
Viewing the Flow Summary Page 1116
Creating Traffic Profiles 1118
Providing Basic Profile Information 1121
Specifying Traffic Profile Conditions 1122
Adding a Host Profile Qualification 1124
Setting Profile Options 1127
Saving a Traffic Profile 1129
Activating and Deactivating Traffic Profiles 1129
Editing a Traffic Profile 1130
Understanding Condition-Building Mechanics 1130
Viewing Traffic Profiles 1138

Keysight, EX-1004
Page 17

Chapter 26: Introduction to NetFlow 1140

Understanding NetFlow 1141
- What Information Does NetFlow Collect and Provide? 1141
- How is NetFlow Data Different from RNA Flow Data? 1142
- What are the Prerequisites for NetFlow Data Collection? 1146

Planning your NetFlow Deployment 1148

Configuring NetFlow Data Collection 1150
- Licensing NetFlow 1150
- Configuring the Sourcefire 3D System to see NetFlow Devices 1151
- Using NetFlow in an RNA Detection Policy 1152

Chapter 27: Using RNA as a Compliance Tool 1156

Understanding Compliance White Lists 1158
- Understanding White List Targets 1159
- Understanding White List Host Profiles 1160
- Understanding White List Evaluations 1164
- Understanding White List Violations 1165

Creating Compliance White Lists 1167
- Surveying Your Network 1168
- Providing Basic White List Information 1170
- Configuring Compliance White List Targets 1171
- Configuring Compliance White List Host Profiles 1175

Managing Compliance White Lists 1188
- Modifying a Compliance White List 1188
- Deleting a Compliance White List 1189

Working with Shared Host Profiles 1189
- Creating Shared Host Profiles 1190
- Modifying a Shared Host Profile 1192
- Deleting a Shared Host Profile 1195
- Resetting Built-In Host Profiles to Their Factory Defaults 1195

Working with White List Events 1197
- Viewing White List Events 1197
- Understanding the White List Events Table 1200
- Searching for Compliance White List Events 1201

Working with White List Violations 1204
- Viewing White List Violations 1204
- Understanding the White List Violations Table 1207
- Searching for White List Violations 1208

Keysight, EX-1004
Page 18

Chapter 28: Configuring Compliance Policies and Rules 1212

Creating Rules for Compliance Policies 1214
- Providing Basic Rule Information 1216
- Specifying Compliance Rule Trigger Criteria 1216
- Adding a Host Profile Qualification 1233
- Adding a Flow Tracker 1237
- Adding a User Qualification 1248
- Adding Snooze and Inactive Periods 1250
- Understanding Rule-Building Mechanics 1251

Managing Rules for Compliance Policies 1260
- Modifying a Rule 1260
- Deleting a Rule 1261
- Creating a Rule Group 1261
- Deleting a Rule Group 1262

Creating Compliance Policies 1262
- Providing Basic Policy Information 1264
- Adding Rules and White Lists to a Compliance Policy 1265
- Setting Rule and White List Priorities 1266
- Adding Responses to Rules and White Lists 1267

Managing Compliance Policies 1269
- Activating and Deactivating Compliance Policies 1270
- Modifying a Compliance Policy 1270
- Deleting a Compliance Policy 1271

Working with Compliance Events 1271
- Viewing Compliance Events 1272
- Understanding the Compliance Events Table 1274
- Searching for Compliance Events 1276

Chapter 29: Configuring Responses for Compliance Policies 1281

Configuring Alerts 1282
- Creating Alerts 1285
- Modifying Alerts 1292
- Deleting Alerts 1292
- Activating and Deactivating Alerts 1293

Assigning Notification Alerts to RNA Events 1293

Configuring External Alerting on Impact Flags 1295

Configuring Remediations 1296
- Configuring Remediations for Check Point Firewalls 1298
- Configuring Remediations for Cisco IOS Routers 1319
- Configuring Remediations for Cisco PIX Firewalls 1327
- Configuring Nmap Remediations 1333
- Configuring Set Attribute Remediations 1340

Keysight, EX-1004
Page 19

Working with Remediation Status Events 1343
Viewing Remediation Status Events 1344
Working with Remediation Status Events 1346
Understanding the Remediation Status Table 1347
Searching for Remediation Status Events 1349
Configuring Response Groups 1351
Creating a Response Group 1352
Modifying a Response Group 1353
Deleting a Response Group 1353
Activating and Deactivating Response Groups 1353

Chapter 30: Enhancing RNA Detection 1355

Assessing Your Detection Strategy 1356
Are Your Sensors Correctly Placed? 1356
Do Unidentified Operating Systems Have a Unique TCP Stack? 1357
Can RNA Identify All Services? 1358
Have You Applied Patches that Fix Vulnerabilities? 1358
Do You Want to Track Third-Party Vulnerabilities? 1358
Enhancing Your Network Map 1358
Understanding Passive Detection 1359
Understanding Active Detection 1359
Understanding Current Identities 1360
Understanding Identity Conflicts 1362
Using Custom Fingerprinting 1363
Fingerprinting Clients 1364
Fingerprinting Servers 1370
Managing Fingerprints 1377
Activating Fingerprints 1377
Deactivating Fingerprints 1378
Deleting Fingerprints 1378
Editing Fingerprints 1379
Working with Service and Application Detectors 1381
Understanding RNA Service Detection 1384
Creating a User-Defined Service Detector 1387
Managing RNA Detectors 1394
Working with RNA Detector Groups 1402
Importing Host Input Data 1406
Enabling Use of Third-Party Data 1407
Managing Third-Party Product Mappings 1408
Mapping Third-Party Vulnerabilities 1412
Managing Custom Product Mappings 1414

Keysight, EX-1004
Page 20

Chapter 31: Configuring Active Scanning 1418

Understanding Nmap Scans 1419
- Understanding Nmap Remediations 1419
- Creating an Nmap Scanning Strategy 1422
- Sample Nmap Scanning Profiles 1424

Setting up Nmap Scans 1428
- Creating an Nmap Scan Instance 1428
- Creating an Nmap Scan Target 1430
- Creating an Nmap Remediation 1431

Managing Nmap Scanning 1437
- Managing Nmap Scan Instances 1437
- Managing Nmap Remediations 1439
- Running an On-Demand Nmap Scan 1440

Managing Scan Targets 1442
- Editing a Scan Target 1442
- Deleting a Scan Target 1443

Working with Active Scan Results 1443
- Monitoring Scans 1444
- Importing Results 1445
- Viewing Scan Results 1445
- Understanding the Scan Results Table 1448
- Searching for Scan Results 1449
- Analyzing Nmap Results 1451

Chapter 32: Using Sourcefire RUA 1452

Understanding RUA 1453
- How Does RUA Work? 1454
- How Do I Choose an RUA Implementation? 1463
- What Information Does RUA Provide? 1466
- How Can You Use RUA? 1468
- What are the Limitations of Sourcefire RUA? 1473

Configuring RUA 1474
- Licensing RUA 1475
- Restricting RUA Logging 1476
- Connecting to LDAP Servers 1478
- Managing RUA LDAP Connections 1482
- Configuring an RUA Agent on an Active Directory Server 1483
- Setting up 3D Sensors and RUA 1487

Working with RUA Users 1489
- Viewing RUA Users 1490
- Understanding the RUA Users Table 1493
- Understanding User Details and Host History 1495
- Searching for RUA Users 1497

Keysight, EX-1004
Page 21

Working with RUA Events 1500
Viewing RUA Events 1501
Understanding the RUA Events Table 1505
Searching for RUA Events 1506

Chapter 33: Using PEP to Manage Traffic 1510

Understanding PEP Traffic Management 1511
Creating PEP Policies 1513
Selecting and Changing Target PEP Interface Sets 1514
Building and Editing IPv4 Fast Path Rules 1516
Building and Editing IPv6 Fast Path Rules 1518
Building and Editing IPv4 PEP Rules 1520
Building and Editing IPv6 PEP Rules 1522
Working with PEP Policies 1525
Editing PEP Policies 1526
Applying PEP Policies 1527
Deleting PEP Policies 1529
Viewing Active PEP Policy Details 1529

Chapter 34: Working with Event Reports 1531

Generating Reports from an Event View 1531
Generating a Basic Report 1532
Viewing Generated Reports 1534
Downloading Generated Reports 1535
Deleting Generated Reports 1535
Running Remote Reports 1536
Using Remote Storage for Reports 1537
Using Report Profiles 1538
Using a Predefined Report Profile 1538
Creating a Report Profile 1542
Setting Up the Report Options 1547
Editing Report Profiles 1549
Deleting Report Profiles 1549

Chapter 35: Searching for Events 1550

Performing and Saving Searches 1551
Performing a Search 1551
Loading a Saved Search 1554
Deleting a Saved Search 1554
Using Wildcards and Symbols in Searches 1555
Specifying Time Constraints in Searches 1556

Keysight, EX-1004
Page 22

Specifying IP Addresses in Searches 1557
Specifying Ports in Searches 1558
Specifying Detection Engine/Appliance Names 1559

Chapter 36: Using Custom Tables 1561

Understanding Custom Tables 1562
Understanding Possible Table Combinations 1563
Creating a Custom Table 1565
Modifying a Custom Table 1568
Deleting a Custom Table 1568
Viewing a Workflow Based on a Custom Table 1569
Searching Custom Tables 1569

Chapter 37: Understanding and Using Workflows 1573

Components of a Workflow 1574
Comparing Predefined and Custom Workflows 1575
Comparing Workflows for Predefined and Custom Tables 1576
Predefined Intrusion Event Workflows 1576
Predefined RNA Workflows 1580
Predefined Third-Party Vulnerabilities Workflows 1582
Predefined Flow Data Workflows 1582
Predefined RUA Workflows 1584
Predefined Compliance and White List Workflows 1585
Predefined System Workflows 1586
Saved Custom Workflows 1587
Using Workflows 1589
Selecting Workflows 1590
Understanding the Workflow Toolbar 1592
Using Workflow Pages 1593
Setting Event Time Constraints 1597
Constraining Events 1606
Using Compound Constraints 1609
Sorting Table View Pages and Changing Their Layout 1611
Sorting Drill-down Workflow Pages 1612
Selecting Rows on a Workflow Page 1612
Navigating to Other Pages in the Workflow 1613
Navigating between Workflows 1614
Using Bookmarks 1615

Keysight, EX-1004
Page 23

Using Custom Workflows 1618
Creating Custom Workflows 1618
Creating Custom Flow Data Workflows 1621
Viewing Custom Workflows 1623
Editing Custom Workflows 1624
Deleting Custom Workflows 1625

Chapter 38: Managing Users 1626

Understanding Sourcefire User Authentication 1626
Understanding Internal Authentication 1628
Understanding External Authentication 1628
Understanding User Privileges 1629

Managing Authentication Objects 1631
Understanding LDAP Authentication 1631
Preparing to Create an LDAP Authentication Object 1636
Quick Start to LDAP Authentication 1637
Creating Advanced LDAP Authentication Objects 1642
LDAP Authentication Object Examples 1657
Editing LDAP Authentication Objects 1661
Understanding RADIUS Authentication 1662
Creating RADIUS Authentication Objects 1662
RADIUS Authentication Object Examples 1671
Editing RADIUS Authentication Objects 1674
Deleting Authentication Objects 1674

Managing User Accounts 1675
Viewing User Accounts 1675
Adding New User Accounts 1676
Managing Command Line Access on 8000 Series Sensors 1678
Managing Externally Authenticated User Accounts 1680
Managing User Password Settings 1681
Configuring User Roles 1682
Modifying User Privileges and Options 1684
Modifying Restricted Event Analyst Access Properties 1685
Modifying User Passwords 1689
Deleting User Accounts 1690
User Account Privileges 1690

Keysight, EX-1004
Page 24

Chapter 39: Scheduling Tasks 1699

Configuring a Recurring Task 1700
Automating Backup Jobs 1702
Automating Software Updates 1704
- Automating Software Downloads 1705
- Automating Software Pushes 1706
- Automating Software Installs 1708

Automating Vulnerability Database Updates 1710
- Automating VDB Update Downloads 1711
- Automating VDB Update Pushes 1712
- Automating VDB Update Installs 1714

Automating SEU Imports 1716
Automating Intrusion Policy Applications 1719
Automating Reports 1721
Automating Nmap Scans 1722
- Preparing Your System for an Nmap Scan 1723
- Scheduling an Nmap Scan 1723

Automating Recommended Rule State Generation 1725
Viewing Tasks 1727
- Using the Calendar 1727
- Using the Task List 1729

Editing Scheduled Tasks 1730
Deleting Scheduled Tasks 1730
- Deleting a Recurring Task 1731
- Deleting a One-Time Task 1731

Chapter 40: Managing System Policies 1732

Creating a System Policy 1733
Editing a System Policy 1735
Applying a System Policy 1736
Deleting System Policies 1737

Keysight, EX-1004
Page 25

Configuring the Sections of Your System Policy 1738
Configuring the Access List for Your Appliance 1738
Configuring Audit Log Settings 1740
Configuring Authentication Profiles 1741
Configuring Command Line Interface Settings 1745
Configuring Dashboard Settings 1746
Configuring Database Event Limits 1747
Configuring DNS Cache Properties 1750
Configuring a Mail Relay Host and Notification Address 1751
Configuring Intrusion Policy Preferences 1753
Specifying a Different Language 1754
Adding a Custom Login Banner 1755
Configuring RNA Settings 1756
Configuring RUA Settings 1764
Configuring SNMP Polling 1766
Synchronizing Time 1768
Mapping Vulnerabilities for Services 1773

Chapter 41: Configuring System Settings 1775

Viewing and Modifying the Appliance Information 1777
Understanding Licenses 1780
Licensing an 8000 Series Sensor 1781
Understanding Feature Licenses 1784
Verifying Your Product License 1786
Managing Feature Licenses 1787
Using Custom HTTPS Server Certificates 1794
Viewing the Current HTTPS Server Certificate 1795
Generating a Certificate Request 1796
Uploading Certificates 1797
Enabling Access to the Database 1799
Configuring Network Settings 1800
Editing Network Interface Configurations 1803
Shutting Down and Restarting the System 1806
Configuring the Communication Channel 1807
Setting Up the Management Virtual Network 1808
Editing the Management Virtual Network 1809
Configuring Remote Access to the Defense Center 1810
Setting the Time Manually 1813
Blacklisting Health Modules 1815
Specifying NetFlow-Enabled Devices 1816

Keysight, EX-1004
Page 26

Managing Remote Storage 1817
Using Local Storage 1817
Using NFS for Remote Storage 1818
Using SSH for Remote Storage 1820
Using SMB for Remote Storage 1821
Lights-Out Management 1823
Understanding a Serial Over LAN Connection 1823
Understanding Lights-Out Management 1824
Enabling Lights-Out Management on the 3D Sensor 1827
Enabling Lights-Out Management User Access 1828
Using Serial Over LAN or Lights-Out Management 1830

Chapter 42: Updating System Software 1831

Understanding Update Types 1832
Installing Software Updates 1833
Updating a Defense Center or Master Defense Center 1835
Updating Managed Sensors 1838
Updating Unmanaged 3D Sensors 1840
Monitoring the Status of Major Updates 1843
Uninstalling Software Updates 1845
Updating the Vulnerability Database 1846
Importing SEUs and Rule Files 1848
Using One-Time SEU Imports 1851
Using Recurring SEU Imports 1855
Importing Local Rule Files 1857
Viewing the SEU Import Log 1859

Chapter 43: Monitoring the System 1869

Viewing Host Statistics 1870
Monitoring System Status and Disk Space Usage 1874
Viewing System Process Status 1875
Understanding Running Processes 1877
Understanding System Daemons 1878
Understanding Executables and System Utilities 1880
Viewing IPS Performance Statistics 1883
Generating IPS Performance Statistics Graphs 1884
Saving IPS Performance Statistics Graphs 1885
Viewing RNA Performance Statistics 1886
Generating RNA Performance Statistics Graphs 1887
Saving RNA Performance Statistics Graphs 1889

Keysight, EX-1004
Page 27

Chapter 44: Using Health Monitoring 1890

Understanding Health Monitoring 1891
- Understanding Health Policies 1892
- Understanding Health Modules 1893
- Understanding Health Monitoring Configuration 1895

Configuring Health Policies 1895
- Predefined Health Policies 1896
- Creating Health Policies 1902
- Applying Health Policies 1928
- Editing Health Policies 1931
- Deleting Health Policies 1934

Using the Health Monitor Blacklist 1935
- Blacklisting Health Policies or Appliances 1936
- Blacklisting a Health Policy Module 1938

Configuring Health Monitor Alerts 1940
- Preparing to Create a Health Alert 1941
- Creating Health Monitor Alerts 1942
- Interpreting Health Monitor Alerts 1943
- Editing Health Monitor Alerts 1944
- Deleting Health Monitor Alerts 1945

Using the Health Monitor 1945
- Interpreting Health Monitor Status 1947

Using Appliance Health Monitors 1947
- Interpreting Appliance Health Monitor Status 1949
- Viewing Alerts by Status 1950
- Running All Modules for an Appliance 1950
- Running a Specific Health Module 1951
- Generating Health Module Alert Graphs 1953
- Using the Health Monitor to Troubleshoot 1954

Working with Health Events 1957
- Understanding Health Event Views 1958
- Viewing Health Events 1958
- Understanding the Health Events Table 1963
- Searching for Health Events 1965

Chapter 45: Auditing the System 1968

Managing Audit Records 1969
- Viewing Audit Records 1969
- Suppressing Audit Records 1973
- Understanding the Audit Log Table 1977
- Searching Audit Records 1978

Viewing the System Log 1981
- Filtering System Log Messages 1982

Keysight, EX-1004
Page 28

Chapter 46: Using Backup and Restore 1985

Creating Backup Files 1986
Creating Backup Profiles 1990
Performing Sensor Backup with the Defense Center 1991
Uploading Backups from a Local Host 1992
Restoring the Appliance from a Backup File 1993

Appendix A: Importing and Exporting Objects 1997

Exporting Objects 1998
- Exporting a Custom Table 1999
- Exporting a Custom Workflow 1999
- Exporting a Dashboard 2000
- Exporting a Health Policy 2001
- Exporting an Intrusion Policy 2001
- Exporting a PEP Policy 2002
- Exporting an RNA Detection Policy 2003
- Exporting a System Policy 2003
- Exporting a Third-Party Product Mapping 2004
- Exporting a Third-Party Vulnerability Mapping 2005
- Exporting an RNA Detector 2005
- Exporting Multiple Objects 2006

Importing Objects 2009

Appendix B: Purging the RNA and RUA Databases 2014

Appendix C: Viewing the Status of Long-Running Tasks 2016

Viewing the Task Queue 2016
Managing the Task Queue 2018

Appendix D: Detected Operating Systems and Services 2019

Detected Operating Systems 2020
- Desktops, Servers, and NAS Devices 2020
- Printers 2023
- Gaming Systems, Music, and Television 2024
- Telephones, Communications, VoIP, and Videoconferencing 2024
- Routers, Switches, Access Points, and Network Appliances 2026

Keysight, EX-1004
Page 29

Detected Services 2027
Service Detector Types 2028
Special Considerations for Service Detection 2029
Services Detected by RNA 2030
Detected Client Applications 2036
Detected Payloads 2039

Appendix E: Command Line Reference 2044

Basic CLI Commands 2045
configure password 2045
end 2046
exit 2046
help 2046
history 2047
logout 2047
? (question mark) 2048
?? (double question marks) 2048
Show Commands 2049
alarms 2050
clustering 2050
cpu 2050
detection-engines 2052
disk 2052
dns 2052
failopen 2053
gui 2053
hostname 2053
hyperthreading 2054
interfaces 2054
interface-sets 2054
lcd 2055
link-state 2055
managers 2055
memory 2056
model 2056
network 2056
network-modules 2057

Keysight, EX-1004
Page 30

ntp 2057
portstats 2057
processes 2058
serial-number 2058
summary 2059
time 2059
user 2059
users 2060
version 2061
Configuration Commands 2061
gui 2062
lcd 2062
hyperthreading 2062
manager 2063
network 2064
user 2066
system file 2069
system generate-troubleshoot 2071
system reboot 2071
system restart 2071
system shutdown 2072

Appendix F: Third-Party Products 2073

Glossary 2074

Index 2099

Keysight, EX-1004
Page 31

Chapter 1
Introduction to the Sourcefire 3D System

The Sourcefire 3D System® provides you with real-time network intelligence for real-time network defense. The Sourcefire 3D System has the tools you need to:

- discover the changing assets and vulnerabilities on your network
- determine the types of attacks against your network and the impact they have to your business processes
- defend your network in real time

The topics that follow introduce you to the Sourcefire 3D System and describe some of the key components that contribute to its value as a part of any security strategy for your network.

- Components of the Sourcefire 3D System on page 33 provides descriptions of each of the components that may be in your Sourcefire 3D System.
- Logging into the Appliance on page 37 explains how to access the web interface on your appliance and log in using one of the user accounts.
- Logging into the Appliance to Set Up an Account on page 40 explains how to set up an association between a external user account and a set of credentials on the appliance.
- Logging Out of the Appliance on page 42 explains how to log out of the web interface.
- Specifying Your User Preferences on page 43 explains how to 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.
- Using the Context Menu on page 53 explains how to display a context-specific menu of shortcuts on certain pages in the web interface.

Keysight, EX-1004
Page 32

- Documentation Resources on page 55 explains where to locate specific information about using the Defense Center.
- Documentation Conventions on page 55 explains conventions used throughout the guide to convey specific types of information.
- IP Address Conventions on page 59 explains how the Sourcefire 3D System treats IP address ranges specified using Classless Inter-Domain Routing (CIDR) notation.

Components of the Sourcefire 3D System

The topics that follow introduce you to Sourcefire Sourcefire 3D System and describe some of the key components that contribute to its value as a part of any security strategy for your network.

- 3D Sensors on page 33
- Defense Centers on page 34
- Master Defense Centers on page 36
- IPS on page 34
- Sourcefire RNA on page 35
- Sourcefire RUA on page 36
- Intrusion Agents on page 37
- RNA and IPS for Crossbeam Systems on page 37
- eStreamer on page 37

3D Sensors

Sourcefire 3D Sensors™ are fault-tolerant, purpose-built appliances available with throughputs from 5Mbps up to 10Gbps. 3D Sensors passively aggregate network and user intelligence while defending the network against internal and external threats. Each 3D Sensor is capable of running Sourcefire IPS™, RNA® (Real-time Network Awareness), and RUA® (Real-time User Awareness).

Each Sourcefire 3D Sensor is capable of running any combination of the following software components:

Sourcefire IPS™ (Intrusion Prevention System) provides best-in-class intrusion detection and prevention by harnessing the power of the industry-standard Snort® rules-based detection engine. Backed by the acclaimed Sourcefire Vulnerability Research Team™ (VRT), Sourcefire delivers its customers unrivaled protection against known and unknown threats.

Sourcefire RNA® (Real-time Network Awareness) passively monitors networks 24x7 to deliver real-time, comprehensive network intelligence, including operating systems, services, applications, protocols, and potential vulnerabilities. RNA automates key IPS functions while fueling additional Sourcefire network

Keysight, EX-1004
Page 33

security solutions, including Network Visibility, Network Behavior Analysis (NBA), and IT Policy Compliance.

Sourcefire RUA® (Real-time User Awareness) correlates Active Directory and LDAP usernames with host IP addresses involved in security and compliance events. RUA dramatically reduces the time needed to uncover user identity and contact information by 95% or more. Security teams can resolve security and compliance incidents more quickly, when time is of the essence.

Defense Centers

The Defense Center provides a centralized management interface and database repository for the Sourcefire 3D System. You can analyze and respond to events from all your sensors consistently by doing the analysis through an interface where you can see all the data collected by the managed sensors. You can also push policies created on the Defense Center and software updates to managed sensors. If you have software sensors or Intrusion Agents on your network, you must use the Defense Center to manage them. Note that a 3D Sensor running the IPS component includes its own local web interface, but if you want to use RNA on the sensor, you must manage the sensor with a Defense Center.

If you use your Defense Center to manage 3D Sensors that run RNA and IPS (either on the same sensor or different sensors that monitor the same network segments), the Defense Center correlates intrusion events from IPS with host vulnerabilities from RNA and assigns impact flags to the intrusion events. Impact correlation lets you focus in on attacks most likely to damage high priority hosts.

If you deploy Real-time User-Awareness (RUA), the Defense Center correlates threat, endpoint, and network intelligence with user identity information so that you can identify the source of policy breaches, attacks, or network vulnerabilities.

IPS

The Sourcefire Intrusion Prevention System (also called IPS) is one of the components of the Sourcefire 3D System that you can run on the 3D Sensor. IPS allows you to monitor your network for attacks that might affect the availability, integrity, or confidentiality of hosts on the network. By placing 3D Sensors with IPS on key network segments, you can examine the packets that traverse your network for malicious activity. Each 3D Sensor with IPS uses rules, decoders, and preprocessors to look for the broad range of exploits that attackers have developed.

3D Sensors that are licensed to use IPS include a set of *intrusion rules* developed by the Sourcefire Vulnerability Research Team (VRT). You can choose to enable rules that would detect the attacks you think most likely to occur on your network. You can also create custom intrusion rules tuned to your environment. In addition, 3D Sensors with IPS run *preprocessors* against detected network traffic to normalize traffic and detect malicious packets.

Keysight, EX-1004
Page 34

When a 3D Sensor 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.

In a Sourcefire 3D System deployment that includes 3D Sensors with IPS and a Defense Center, the sensors transmit events and sensor statistics to the Defense Center where you can view the aggregated data and gain a greater understanding of the attacks against your network assets.The Defense Center can also push health, system, and intrusion policies to your sensors. You can push software updates from the Defense Center to sensors as well. For more information, see What Can Be Managed by a Defense Center? on page 106.

If your 3D Sensor is running IPS, you can also use a local web interface to create intrusion policies and review the resulting intrusion events. Note that if you do manage your 3D Sensors with a Defense Center, Sourcefire recommends that you use only the Defense Center's web interface to interact with the sensor and its data.

If you deploy your 3D Sensor inline on your network and create what is called an inline detection engine, you can configure your 3D Sensor to drop or replace packets that you know to be harmful.

Sourcefire RNA

Sourcefire Real-time Network Awareness (also called RNA) is one of the components of the Sourcefire 3D System that you can use on your 3D Sensor. RNA monitors traffic on your network, using information from detected packets to build a comprehensive map of the devices on the network. You can set up compliance policies, compliance white lists, and traffic profiles to protect your company's infrastructure by monitoring network traffic for unusual patterns or behavior and automatically responding as needed. You **must** use a Defense Center to manage a 3D Sensor that is running RNA.

As RNA passively observes traffic, listening to the network segments you specify, it compiles the following information:

- the number and types of network devices running on your network
- the operating systems running on monitored network devices
- the active services and open ports on monitored network devices
- the vulnerabilities and exploits to which monitored network devices may be susceptible
- *flow data*, which are records of active sessions involving monitored network devices including the frequency and size of the session, as well as the service and protocol used and, if applicable, the client application and URL involved in the session

You can access event views and graphs to analyze this collected data. RNA builds a host profile for each host it detects, containing host details such as detected

Keysight, EX-1004
Page 35

operating system, services, and protocols, and assigned host attributes. RNA assigns vulnerabilities to the host based on the operating system vendor and version detected for the host. You can access host profiles by browsing the network map or through one of the workflows Sourcefire provides to aid your analysis.

3D Sensors running RNA transmit the network map, event and flow data, and sensor statistics to the Defense Center so you can see a consolidated view of events. The Defense Center can also push health, system, and RNA detection policies to your sensors. You can push vulnerability database (VDB) and software updates from the Defense Center as well. For more information, see What Can Be Managed by a Defense Center? on page 106.

Sourcefire RUA

The Real-time User-Awareness component (also called RUA) allows you to create policies and response rules that are user-based. You can apply these policies and rules across the Sourcefire 3D System. As a result, RUA enables you to implement and enforce policies specific to individuals, departments, or other user characteristics.

The network protocol used by your organization to provide user authentication largely determines the amount of data and efficiency of RUA. See Using Sourcefire RUA on page 1452 for more information about RUA.

PEP Traffic Management

PEP is a technology based on the hardware capabilities of the 3D9900 Sensors. PEP allows you to create rules to block, analyze, or send traffic directly through the 3D9900 with no further inspection. PEP traffic management enhances the sensor's efficiency by allowing you to pre-select traffic to cut through or to drop instead of analyzing.

Master Defense Centers

The Sourcefire Master Defense Center is a key component in the Sourcefire 3D System. You can use the Master Defense Center to aggregate and analyze intrusion events, compliance events, and white list events from up to ten Defense Centers within your Sourcefire 3D System deployment. The Master Defense Center can also aggregate events related to the health of the Defense Centers it is managing. In this way, you can view the current status of the Defense Centers across your enterprise from a single web interface.

See Using the Master Defense Center on page 165 for more information about managing your Defense Centers with a Master Defense Center.

Keysight, EX-1004
Page 36

Intrusion Agents

If you have an existing installation of Snort®, you can install an Intrusion Agent to forward intrusion events to a Defense Center. You can then analyze the events detected by Snort alongside your other data. Although you cannot manage policies or rules for an Intrusion Agent from the Defense Center, you can do analysis and reporting on those events. If the network map on the Defense Center has entries for the target host in a given event, the Defense Center assigns impact flags to the events. You can continue to manually tune Snort rules and preprocessors with the Intrusion Agent in place.

IMPORTANT! When using Intrusion Agents registered to Defense Centers configured for high availability and managed by a Master Defense Center, register all Intrusion Agents to the primary Defense Center.

RNA and IPS for Crossbeam Systems

The Sourcefire 3D System currently supports software-only versions of RNA and IPS for Crossbeam Systems X-Series security switches. RNA and IPS data received by a Defense Center from a Crossbeam-based software sensors is treated in a similar way to data received from a 3D Sensor. Separate installation and configuration guides are available for the 3D Sensor Software for X-Series.

IMPORTANT! Because the 3D Sensor Software for X-Series does not have a web interface, you must use a Defense Center to manage it.

eStreamer

You can access event data within your own applications through the eStreamer Application Programming Interface (API). eStreamer integration requires custom programming, but allows you to request specific data from a Defense Center. 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. See the *eStreamer Integration Guide* for more information.

Logging into the Appliance

Platform: DC/MDC/3D Sensor

Component: None

The Defense Center and many 3D Sensor models have a web-based interface that you can use to perform administrative, management, and analysis tasks.

Keysight, EX-1004
Page 37

If your Series 2 3D Sensor **is** licensed for IPS, you are presented with a more complete web interface that you can use to perform additional configuration and event analysis.If your Series 2 3D Sensor is **not** licensed for IPS, there is a limited web interface that you can use to set up the appliance for the first time and to register the sensor with a Defense Center. The 8000 Series 3D Sensors have the same limited web interface as Series 2 sensors.

Software sensors such as Crossbeam-based software sensors, Intrusion Agents, and Virtual 3D Sensors do **not** have a web interface. You must use the Defense Center's web interface to manage these sensors.

You can access the web interface by logging into the appliance using a web browser. The current version of the web interface is tested with the browsers listed in the following table.

Browser Requirements

Browser	Required Enabled Options and Settings
Firefox 3.6.x	• JavaScript • cookies • Secure Sockets Layer (SSL) v3
Microsoft Internet Explorer 7	• JavaScript • cookies • Secure Sockets Layer (SSL) v3 • 128-bit encryption • Active scripting security setting • Set **Check for newer versions of stored pages** to **Automatically**
Microsoft Internet Explorer 8	• JavaScript • cookies • Secure Sockets Layer (SSL) v3 • 128-bit encryption • Active scripting security setting • Compatibility View • Set **Check for newer versions of stored pages** to **Automatically**

If you are the first user to log into the appliance after it is installed, you must log in using the `admin` user account. The initial setup process is described in the installation guide for your appliance.

Keysight, EX-1004
Page 38

After you log into the appliance, the features that 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.

TIP! 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.

When the appliance was installed, the user who performed the installation created a single administrative user account and password. The first time you log into the appliance, you should use this account. After you create other user accounts as described in Adding New User Accounts on page 1676, you and other users should use those accounts to log into the appliance.

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! Because the Defense Center and the 3D Sensor audit user activity based on user accounts, you should make sure that users log into the system with the correct account.

Your session automatically logs you out after 3.5 hours of inactivity, unless you are viewing a page (such as an unpaused dashboard) that periodically communicates with the web server on the appliance.

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.

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.

   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.

Keysight, EX-1004
Page 39

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 51 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 an 8000 Series or Virtual 3D Sensor via the command line:

Access: CLI Basic or 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 2044 for more information on available CLI commands.

Logging into the Appliance to Set Up an Account

Platform: DC/MDC/3D Sensor

Component: None

Some user accounts may be authenticated through an external authentication server. If this is the case, the first time you log into the Defense Center or 3D Sensor 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.

If the default role for external user accounts is set to a specific access role, externally authenticated users can log into the appliance using their external account credentials without any additional configuration by the system administrator. If an account is externally authenticated and by default receives no

Keysight, EX-1004
Page 40

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.

LDAP usernames can include underscores (_), periods (.), and hyphens (-) but otherwise only alphanumeric characters are supported.

Note that when a shell access user logs into the appliance, it does not create a local user account. Shell access is controlled entirely through 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.

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! The 3Dx800 sensor models do not have a web interface. Instead, use the Defense Center's web interface to manage policies and view events.

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. In the **Username** and **Password** fields, type your user name 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.

Keysight, EX-1004
Page 41

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 51 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 re-appears, with the following error message:

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

4. If you do not have access, contact your system administrator and ask them to modify your account privileges or login as a user with Administrator access and modify the privileges for the account. For more information, see Modifying User Privileges and Options on page 1684.

Logging Out of the Appliance

Platform: DC/MDC/3D Sensor

Component: None

Make sure you log out of the appliance, 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.

Note that your session automatically logs you out after 3.5 hours of inactivity, unless you are viewing a page (such as an unpaused dashboard) that periodically communicates with the web server on the appliance.

To log out of the appliance:

Access: Any

▶ Click **Logout** on the toolbar.

Keysight, EX-1004
Page 42

Last Successful Login

Platform: DC/MDC/3D Sensor

Component: None

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

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

Specifying Your User Preferences

Platform: DC/MDC/3D Sensor

Component: None

Users can specify certain preferences for their 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 43 explains how to change the password for your user account.
- Configuring Event View Settings on page 45 describes how the event preferences affect what you see as you view events.
- Setting Your Default Time Zone on page 50 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 Home Page on page 51 explains how to use one of the existing pages as your default home page. After setting this value, this is the first page you see upon logging into the appliance.
- Specifying Your Default Dashboard on page 52 explains how to choose which of the dashboards you want to use as your default dashboard.

Changing Your Password

Platform: DC/MDC/3D Sensor

Component: None

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 44.

Keysight, EX-1004
Page 43

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:

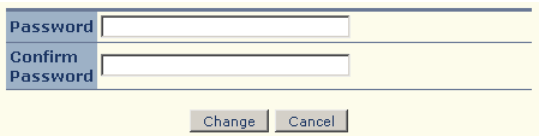
Access: Any

1. In the toolbar, click **Preferences**.

   The User Preferences page appears.

2. Click **Change Password**.

   The Change Password page appears.



3. In the **Current Password** field, type your current password and click **Change**.
4. In the **New Password** and **Confirm** fields, type your new password.
5. Click **Change**.

   A success message appears on the page when your new password is accepted by the system.

Changing an Expired Password

Platform: DC/MDC/3D Sensor

Component: None

Depending on the settings for your user account, your password can 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.

Keysight, EX-1004
Page 44

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.

Configuring Event View Settings

Platform: DC/MDC/3D Sensor

Component: None

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. In the toolbar, click **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 46.
4. Configure the default time window or windows.

   For more information, see Default Time Windows on page 47.
5. Configure default workflows.

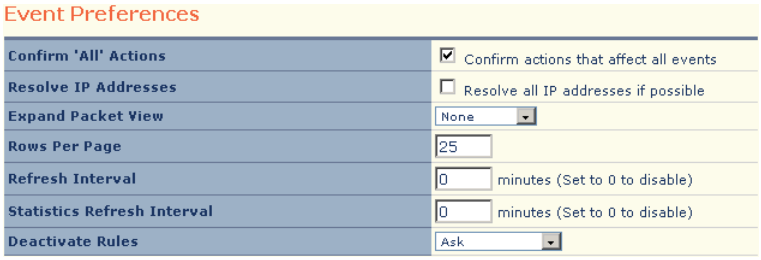
   For more information, see Default Workflows on page 50.
6. Click **Save**.

   Your changes are implemented.

Keysight, EX-1004
Page 45

Event Preferences

Use the Event Preferences section of the Event View Settings page to configure basic characteristics of event views in the Sourcefire 3D System.



The following fields appear in the Event Preferences section.

Confirm "All" Actions

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.

Resolve IP Addresses

Whenever possible, allows the appliance to display host names instead of IP addresses in event views.

Note that an event view can 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 1800.

Expand Packet View

Allows you to configure how the packet view for intrusion events appears. By default, the appliance displays a collapsed version of the packet view.

- **None** - collapse all subsections of the Packet Information section of the packet view
- **Packet Text** - expand only the Packet Text subsection
- **Packet Bytes** - expand only the Packet Bytes subsection
- **All** - expand all sections

Regardless of the default setting, you can always manually expand the sections in the packet view to view detailed information about a captured packet. For more information on the packet view, see Using the Packet View on page 290.

Keysight, EX-1004
Page 46

Rows Per Page

Controls how many rows of events per page you want to appear in drill-down pages and table views.

Refresh Interval

Sets the refresh interval for event views, in minutes. Entering zero disables the refresh option. Note that this interval does not apply to dashboards.

Statistics Refresh Interval

Sets the refresh interval for event summary pages such as the Intrusion Event Statistics and RNA Statistics pages. Entering zero disables the refresh option. Note that this interval does not apply to dashboards.

Deactivate Rules

Controls which links appear on the packet view for intrusion events generated by standard text rules.

- **All Policies** - a single link that deactivates the standard text rule in all the locally defined custom intrusion policies
- **Current Policy** - a single link that deactivates the standard text rule in only the currently applied intrusion policy. Note that you cannot deactivate rules in the default policies.
- **Ask** - links for each of these options

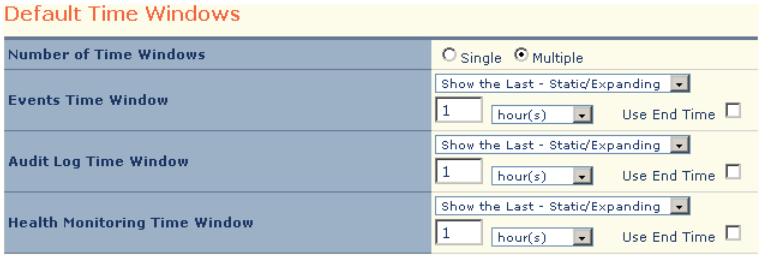
To see these links on the packet view, your user account must have either Administrator access or both Intrusion Event Analyst and Policy & Response Administrator access.

Default Time Windows

Platform: DC/MDC/3D Sensor

Component: None

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. The following graphic shows the Defense Center version of the page.



Keysight, EX-1004
Page 47

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 1597.

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 (depending on the appliance) intrusion events, RNA events, flow data, RUA events, compliance events, remediation status events, white list events, the SEU import log, and event views for custom tables 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, RNA event analysts, and IPS event 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 RNA hosts, host attributes, services, client applications, vulnerabilities, RUA users, 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.



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).

Keysight, EX-1004
Page 48

The following fields appear in the Time Window Settings section.

Show the Last - Sliding

This setting allows you to 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.

Show the Last - Static/Expanding

This setting 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.

Current Day - Static/Expanding

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.

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.

Current Week - Static/Expanding

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.

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.

Keysight, EX-1004
Page 49

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

Platform: DC/MDC/3D Sensor

Component: None

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, depending on the type of analysis you are performing, you can choose between 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 not only on the appliance you are using, but also on your user role. For example, on a 3D Sensor without an IPS license, you can only configure the default workflow for the audit log. As another example, on the Defense Center, intrusion event analysts cannot set default RNA workflows. For general information on workflows, see Understanding and Using Workflows on page 1573.

Setting Your Default Time Zone

Platform: DC/MDC/3D Sensor

Component: None

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 sensors, see Synchronizing Time on page 1768.

Keysight, EX-1004
Page 50

To change your time zone:

Access: Any

1. In the toolbar, click **Preferences**.

   The User Preferences page appears.

2. Click **Time Zone Settings**.

   The Time Zone Preference page appears.



3. From the box on the left, 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 box on the right, 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 Home Page

Platform: DC/MDC/3D Sensor

Component: None

You can specify a page within the web interface as your home page for the appliance. The default home page is the dashboard (**Analysis & Reporting > Dashboards**), except for user accounts with External Database User or Restricted Event Analyst, which use the Welcome page.

To specify your home page:

Access: Any

1. In the toolbar, click **Preferences**.

   The User Preferences page appears.

Keysight, EX-1004
Page 51

2. Click **Home Page**.

   The Home Page page appears.



3. Select the page you want to use as your home page from the **Opening Screen** 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 1690.

4. Click **Save**.

   Your home page preference is saved.

Specifying Your Default Dashboard

Platform: DC/MDC/3D Sensor

Component: None

You can specify one of the dashboards on the appliance as the default dashboard. The default dashboard appears when you select **Analysis & Reporting > 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 61.

To specify your default dashboard:

Access: Any except Restricted or Database

1. In the toolbar, click **Preferences**.

   The User Preferences page appears.

2. Click **Dashboard Settings**.

   The Dashboard Settings page appears.



3. Select the dashboard you want to use as your default from the **Default Dashboard** drop-down list.

   If you select **None**, when you select **Analysis & Reporting > Dashboards**, the Dashboard List page appears. You can then select a dashboard to view.

4. Click **Save**.

   Your default dashboard preference is saved.

Keysight, EX-1004
Page 52

Using the Context Menu

Platform: DC/MDC/3D Sensor

Component: None

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. As the name implies, the contents of the menu depend on the context where you access it. For example, if you access the menu while viewing an RNA event, the context menu provides you with the option to view the event in a separate browser window. However, if you access the context menu while viewing an intrusion event that was triggered by an intrusion rule, you have a range of options that includes enabling, disabling, suppressing, and thresholding the rule. You can also view the rule documentation and edit the rule.

You can access the context menu on the following pages.

- Event pages (drill-down pages and table views) contain hotspots over each event.
- The Rule Editor page for intrusion rules contains a hotspot over each intrusion rule.

Note that if you try to access the context menu for a web page or location that doesn't support the Sourcefire-specific menu, the normal context menu for your browser appears.

To access the context menu:

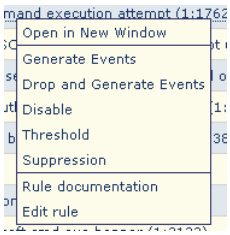
Access: Any

1. On one of the hotspot-enabled pages in the web interface, hover your pointer over one of the hotspots.

   A "Right-click for menu" message appears.

2. Right-click your pointing device.

   A pop-up context menu appears with options that are appropriate for the hotspot. For example, the following menu appears if you right-click over an intrusion event.



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.

Keysight, EX-1004
Page 53

Communication Ports

The Sourcefire 3D System uses ports 443 and 8305 to communicate internally and externally between the Defense Center and sensors. Open other ports to enable optional functionality within your deployment.

Communication Ports

Ports	Description	Protocol	Direction	Open the port to...
22	ssh/ssl	TCP	Bidirectional	allow a secure remote connection to the appliance.
25	smtp	TCP	Outbound	send email notices from the appliance (for example, if a process has completed).
53	dns	TCP	Outbound	use DNS.
67, 68	dhcp	UDP	Outbound	use DHCP. Default is disabled.
80	http	TCP	Outbound	allow the RSS Feed dashboard widget to connect to a remote web server; use for auto-update.
162	snmp	UDP	Bidirectional	provide access if you enabled SNMP polling (inbound) and SNMP traps (outbound).
389, 636	ldap	TCP	Outbound	use RUA and for authentication.
443	https	TCP	Inbound	access the appliance. **Required**.
514	syslog	UDP	Outbound	use for the remote syslog server.
1500, 2000	database access	TCP	Inbound	access the Defense Center or Master Defense Center if external database access is enabled.
1812, 1813	RADIUS	UDP	Outbound	use RADIUS. Open both ports to ensure that RADIUS functions correctly.
3306	RUA Agent	TCP	Inbound	allow communication between the Defense Center and RUA Agents.
8301	Intrusion Agent	TCP	Bidirectional	allow communication between the Defense Center and Intrusion Agents.
8302	eStreamer	TCP	Bidirectional	use for an eStreamer client.

Keysight, EX-1004
Page 54

Communication Ports (Continued)

Ports	Description	Protocol	Direction	Open the port to...
8305	sensor management	TCP	Bidirectional	communicate between the Defense Center and 3D Sensors. **Required**.
8307	Host Input Client API	TCP	Bidirectional	communicate with the Defense Center during client/server authentication.
18183	OPSEC SAM	TCP	Outbound	use OPSEC for remediation.

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 **Operations > 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 a PDF version of the *Sourcefire 3D System User Guide*, which together include the same content as the online help, but in an easy-to-print format.

The Documentation CD also contains copies of the *Defense Center Installation Guide* and the *3D Sensor Installation Guide*, which includes information about installing the appliance as well as hardware specifications and safety information. The CD also contains copies of various API guides and supplementary material. You can access the most up-to-date versions of the PDF documentation on the Sourcefire Support web site (https://support.sourcefire.com/).

Documentation Conventions

This documentation includes information about which Sourcefire 3D System components are required for each feature and which user roles have permission to complete each procedure, as follows:

- Platform Conventions on page 56 explains the meaning of the Platform statement at the beginning of each section.
- Component Conventions on page 56 explains the meaning of the Component statement at the beginning of each section.
- Access Conventions on page 57 explains the meaning of the Access statement at the beginning of each procedure.

Keysight, EX-1004
Page 55

Platform Conventions

The Platform statement at the beginning of a section in this documentation indicates the appliance platform you need to use the feature described in the section. Platform requirement information for specific aspects of a feature is provided where needed.

The following table defines the terms used to indicate each different platform requirement.

Platform Conventions

Term	Description
3D Sensor	A 3D Sensor with a web interface. This indicates that you can use the feature on any sensor model provided it has its own web interface. You do not need to have an IPS license on the sensor, nor must the sensor be managed by a Defense Center.
DC	Any physical or virtual appliance used as a Defense Center
MDC	Any appliance used as a a Master Defense Center
Standalone Sensor	A 3D Sensor with a web interface and an IPS license, regardless of whether it is currently managed by a Defense Center.

If the Platform statement lists several platforms separated by a forward slash, then you can use the feature on any of those appliances. For example, if the Platform statement is DC/3D Sensor, then you can use the feature on either a Defense Center or a 3D Sensor. Occasionally, a topic contains information that pertains to a specific Defense Center or 3D Sensor model. In these cases, the specific model appears in the Platform statement.

Component Conventions

The Component statement at the beginning of a section indicates the component licenses you need to use the feature. The following table defines the abbreviations used to indicate each different platform requirement:

Component Conventions

Term	Description
None	Does not require a license
IPS	Requires an IPS license

Keysight, EX-1004
Page 56

Component Conventions (Continued)

Term	Description
NetFlow	Requires a NetFlow license (generally paired with an RNA license)
RNA	Requires an RNA license
RUA	Requires an RUA license

If more than one component is required, then the components are listed with a plus sign between them. For example, the impact flags feature includes IPS + RNA. If you can access a feature by having either of several licensed components, then those components are listed with a forward slash between them. For example, traffic status monitoring is applicable to both IPS and RNA, so the Component statement is IPS/RNA.

Access Conventions

The Access statement at the beginning of each procedure in this documentation indicates the user role required to perform the procedure. The following table defines the abbreviations used to indicate each different platform requirement.

Access Conventions

Requires Acronym	Indicates
Admin	User must have the Administrator role
Any	User can have any role
Any Analyst	User can have any analyst role
Any except Database	User can have any role except External Database
Any except Restricted	User can have any role except Restricted Analyst or Restricted Analyst (Read Only)
Any Analyst except Restricted	User can have any analyst role except Restricted Analyst or Restricted Analyst (Read Only)
Any IPS	User must have the Intrusion Event Analyst role or Intrusion Event Analyst (Read Only) role or the Restricted Event Analyst role or Restricted Event Analyst (Read Only) role with rights to that function

Keysight, EX-1004
Page 57

Access Conventions (Continued)

Requires Acronym	Indicates
Database	User must have the External Database role
IPS	User must have the Intrusion Event Analyst role or Restricted Event Analyst role with rights to that function
IPS-RO	User must have the Intrusion Event Analyst (Read Only) roleor Restricted Event Analyst (Read Only) role with rights to that function
Maint	User must have the Maintenance role
P&R Admin	User must have the Policy & Response Administrator role
Any RNA	User must have the RNA Event Analyst or RNA Event Analyst (Read Only)or Restricted Event Analyst or Restricted Event Analyst (Read Only) with rights to that function
RNA	User must have the RNA Event Analyst roleor Restricted Event Analyst role with rights to that function
RNA-RO	User must have the RNA Event Analyst (Read Only) roleor Restricted Event Analyst (Read Only) role with rights to that function

A forward slash separating roles indicates that either of the roles can perform the procedure. For example, to view the Hosts network map, a user must have the RNA Event Analyst or RNA Event Analyst (Read Only) role or the Restricted Event Analyst or Restricted Event Analyst (Read Only) role with **RNA Hosts Data** set to **Show All Data** or to show a specific search. The Access setting for the procedure in the Working with the Hosts Network Map topic is `Any RNA/Admin`.

A plus sign separating user roles means that you must have both those roles to perform the procedure. For example, you must have the Administrator role or have the Policy & Response Administrator role in combination with the Intrusion Event Analyst role or the Restricted Event Analyst role with **Intrusion Events Data** set to **Show All Data** or to show a specific search to access the packet view and set thresholding for a rule from the packet view. As a result, the Access setting for the procedure is `IPS + P&R Admin/Admin`.

Keysight, EX-1004
Page 58

IP Address Conventions

You can use Classless Inter-Domain Routing (CIDR) notation to define IP address ranges in many places in the Sourcefire 3D System, including but not limited to the following:

- RNA detection policies
- custom topologies
- auto-assigned networks for user-defined host attributes
- traffic profiles
- compliance rules and white lists
- active scan targets
- intrusion policies, variables, and standard text rules
- PEP

CIDR notation uses a network IP address combined with a bit mask to define the IP addresses in the specified range. 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

When you use CIDR notation to specify a range of IP addresses, the Sourcefire 3D System uses **only** the masked portion of the network IP address you specified, without changing your user input. For example, if you type 10.1.2.3/8, the Sourcefire 3D System uses 10.0.0.0/8, but the web interface continues to display 10.1.2.3/8.

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

Note that for PEP rules in a PEP policy, you can use CIDR notation to annotate IPv6 IP addresses. Because only PEP supports full IPv6 addresses, and only on 8000 Series sensors, only PEP supports IPv6 CIDR notation.

Keysight, EX-1004
Page 59

For example, if you want to create a PEP rule that analyzes the traffic generated on a subnet using the 2001:db8:3c4d:15::/64 unicast address range, you set the initiator address to 2001:db8:3c4d:15::/64 and set the rule to Analyze. This rule analyzes traffic initiated by hosts in the 2001:db8:3c4d:15:0000:0000:0000:0000 to 2001:db8:3c4d:15:ffff:ffff:ffff:ffff address range. If you want to analyze link-local traffic for the same segment, you create a rule with an initiator address of fe80::/10.

Keysight, EX-1004
Page 60

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 Sourcefire 3D System. You can also use the dashboard to see information about the status and overall health of the appliances in your deployment.

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, current version of the Sourcefire 3D System software running on the appliance, and its remote manager.

The dashboard has a time range that constrains its widgets. You can change the time range to reflect a period as short as the last hour or as long as the last year.

Each type of appliance is delivered with a default dashboard, named Default Dashboard. This dashboard provides the casual user with basic event and system status information for your Sourcefire 3D System deployment. Note that because not all widgets are useful for all types of appliances, the default dashboard differs depending on whether you are using a Master Defense Center, Defense Center, or 3D Sensor.

Keysight, EX-1004
Page 61

By default, the home page for your appliance displays the default 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 **Analysis & Reporting > Dashboards**. For more information, see Viewing Dashboards on page 94.

In addition to the default dashboard, the Defense Center is delivered with two other predefined dashboards:

- The Flow Summary dashboard uses flow data to create tables and charts of the activity on your monitored network; for more information on flow summary data, see Understanding Flow Data on page 1072.

  Note that Restricted Event Analysts use the Flow Summary page instead of the Flow Summary Dashboard; see Viewing the Flow Summary Page on page 1116 for more information. External Database users cannot see either.
- The Detailed Dashboard provides advanced users with detailed information about your Sourcefire 3D System deployment, and includes multiple widgets that summarize collected IPS, RNA, compliance, and system status data.

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.

For more information, see the following sections:

- Understanding Dashboard Widgets on page 62
- Understanding the Predefined Widgets on page 67
- Working with Dashboards on page 92

Understanding Dashboard Widgets

Platform: DC/MDC/3D Sensor

Component: None

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 several predefined dashboard widgets, each of which provides insight into a

Keysight, EX-1004
Page 62

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.

For more information, see:

- Understanding Widget Availability on page 63
- Understanding Widget Preferences on page 66
- Understanding the Predefined Widgets on page 67
- Working with Dashboards on page 92

Understanding Widget Availability

Platform: DC/MDC/3D Sensor

Component: None

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 Appliance Information widget is available on all appliances for all user roles, while the Compliance Events widget is available only on the Defense Center for users with Administrator, Intrusion Event Analyst, or RNA Event Analyst 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, the Master Defense Center cannot access flow data, RUA events,

Keysight, EX-1004
Page 63

RNA events, and so on. If you import a dashboard onto a Master Defense Center that contains a Custom Analysis widget configured to display one of those data types, the widget displays an error message.

Similarly, the content of a widget can differ depending on the type of appliance you are using. For example, the Current Interface Status widget on a 3D Sensor displays the status of its sensing interfaces, but on Defense Centers and Master Defense Centers the widget displays only the status of the management interface. Note than any content generated in table format can be sorted 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 101 and Deleting Widgets on page 101.

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	Master Defense Center	Defense Center	Standalone Sensor	Any 3D Sensor Appliance
Appliance Information	X	X	X	X
Appliance Status	X	X		
Compliance Events	X	X		
Current Interface Status	X	X	X	X
Current Sessions	X	X	X	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		
Product Licensing		X		

Keysight, EX-1004
Page 64

Sourcefire Appliances and Dashboard Widget Availability (Continued)

Widget	Master Defense Center	Defense Center	Standalone Sensor	Any 3D Sensor Appliance
Product Updates	X	X	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		

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. User accounts with Restricted Event Analyst or External Database access cannot use dashboards.

User Roles and Dashboard Widget Availability

Widget	Administrator	Maintenance	P&R Admin	IPS Analyst	RNA Analyst
Appliance Information	X	X	X	X	X
Appliance Status	X	X		X	X
Compliance 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	X
Interface Traffic	X	X		X	X
Intrusion Events	X			X	
Network Compliance	X			X	X
Product Licensing	X	X			

Keysight, EX-1004
Page 65

User Roles and Dashboard Widget Availability (Continued)

Widget	Administrator	Maintenance	P&R Admin	IPS Analyst	RNA Analyst
Product Updates	X	X	X		
RSS Feed	X	X	X	X	X
System Load	X	X	X	X	X
System Time	X	X	X	X	X
White List Events	X			X	X

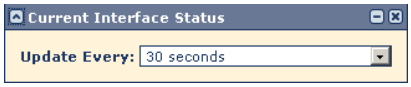
Understanding Widget Preferences

Platform: DC/MDC/3D Sensor

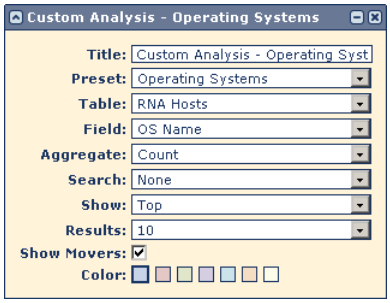
Component: None

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 the network interfaces for the appliance. 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.



Keysight, EX-1004
Page 66

To modify a widget's preferences:

Access: Any except Restricted or Database

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 67.

3. On the widget title bar, click the hide preferences icon () to hide the preferences section.

Understanding the Predefined Widgets

Platform: DC/MDC/3D Sensor

Component: None

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 Sourcefire 3D 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 68
- Understanding the Appliance Status Widget on page 69
- Understanding the Compliance Events Widget on page 69
- Understanding the Current Interface Status Widget on page 70
- Understanding the Current Sessions Widget on page 71
- Understanding the Custom Analysis Widget on page 71
- Understanding the Disk Usage Widget on page 82
- Understanding the Interface Traffic Widget on page 83
- Understanding the Intrusion Events Widget on page 83
- Understanding the Network Compliance Widget on page 85
- Understanding the Product Licensing Widget on page 87
- Understanding the Product Updates Widget on page 88
- Understanding the RSS Feed Widget on page 89
- Understanding the System Load Widget on page 90

Keysight, EX-1004
Page 67

- Understanding the System Time Widget on page 90
- Understanding the White List Events Widget on page 91

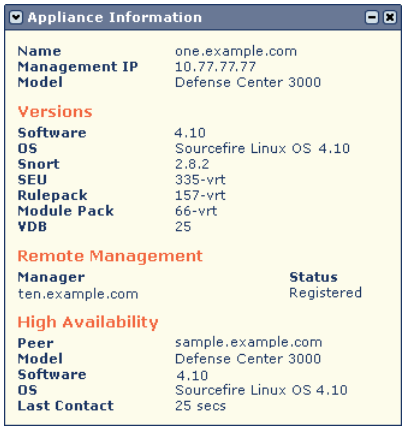
IMPORTANT! The dashboard widgets you can view depend on the type of appliance you are using and on your user role. For more information, see Understanding Widget Availability on page 63.

Understanding the Appliance Information Widget

Platform: DC/MDC/3D Sensor

Component: None

The Appliance Information widget provides a snapshot of the appliance.



The widget provides:

- the name, management interface IP address, and model of the appliance
- the versions of the Sourcefire 3D System software, operating system, Snort, SEU, rule pack, module pack, and vulnerability database (VDB) 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 66.

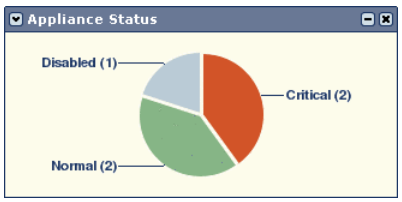
Keysight, EX-1004
Page 68

Understanding the Appliance Status Widget

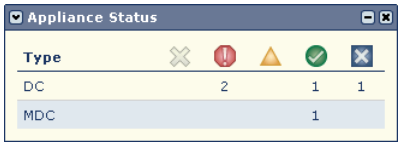
Platform: DC/MDC/3D Sensor

Component: None

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 sensors, you must manually apply a health policy or their status appears as Disabled.



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



The preferences also control how often the widget updates. For more information, see Understanding Widget Preferences on page 66.

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 1945.

Understanding the Compliance Events Widget

Platform: DC/MDC

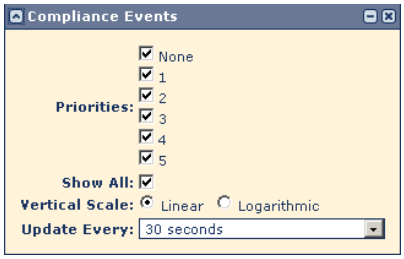
Component: RNA

The Compliance Events widget shows the average events per second by priority, over the dashboard time range.



Keysight, EX-1004
Page 69

You can configure the widget to display compliance 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 compliance events, regardless of priority. The preferences also control how often the widget updates. For more information, see Understanding Widget Preferences on page 66.

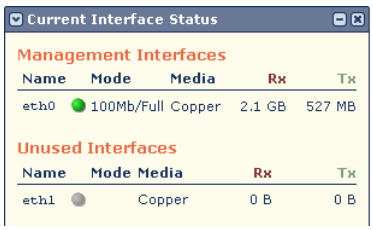
You can click a graph to view compliance events of a specific priority, or click the **All** graph to view all compliance events. In either case, the events are constrained by the dashboard time range; accessing compliance events via the dashboard changes the events (or global) time window for the appliance. For more information on compliance events, see Viewing Compliance Events on page 1272.

Understanding the Current Interface Status Widget

Platform: DC/MDC/3D Sensor

Component: None

The Current Interface Status widget shows the status of the network interfaces for the appliance, grouped by type: management, inline, passive, and unused. Note that only 3D Sensors have interface types other than the management interface.



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)

Keysight, EX-1004
Page 70

- 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

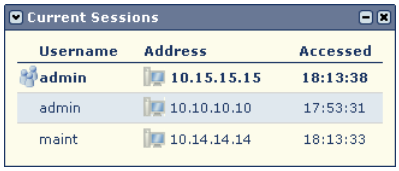
The widget preferences control how often the widget updates. For more information, see Understanding Widget Preferences on page 66.

Understanding the Current Sessions Widget

Platform: DC/MDC/3D Sensor

Component: None

The Current Sessions widget shows which users are logged into the appliance, the IP address of 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 is rendered in bold type.



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 1675
- click the host icon ( ) next to any IP address to view the host profile for that computer; see Using Host Profiles on page 944 (Defense Center with RNA 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 1969

The widget preferences control how often the widget updates. For more information, see Understanding Widget Preferences on page 66.

Understanding the Custom Analysis Widget

Platform: DC/MDC/Standalone Sensor

Component: None

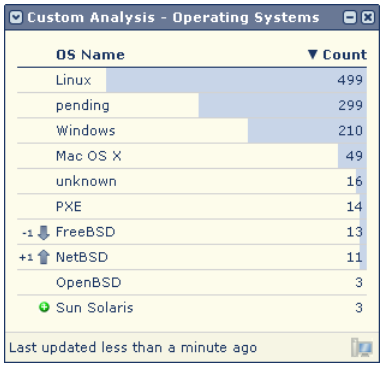
The Custom Analysis widget is a highly customizable widget that allows you to display detailed information on the events collected and generated by the Sourcefire 3D System.

Keysight, EX-1004
Page 71

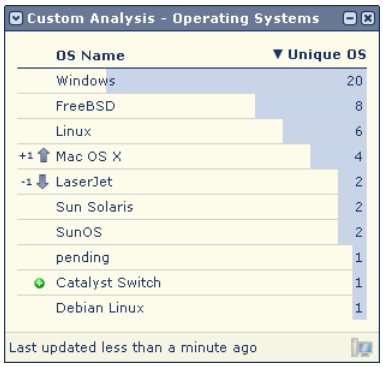
The Custom Analysis widget is delivered with several presets, which are groups of configurations that are predefined by Sourcefire. The presets serve as examples and can provide quick access to information about your deployment. You can use these presets or you can create a custom configuration.

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, if you are using Sourcefire RNA as part of your deployment, you can configure the Custom Analysis widget to display which operating systems are running on the hosts in your organization by configuring the widget to display **OS** data from the **RNA Hosts** table. Aggregating this data by **Count** tells you how many hosts are running each operating system.



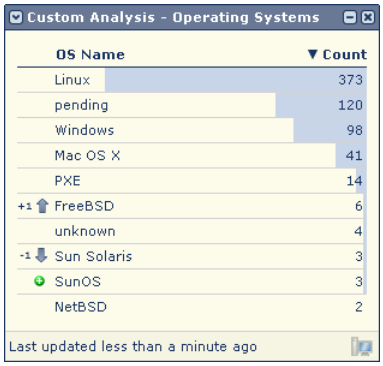
On the other hand, aggregating by **Unique OS** tells you how many unique versions of each operating system are running on the same hosts (for example, how many unique versions of Linux, Microsoft Windows, Mac OS X, and so on).



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 (operating systems

Keysight, EX-1004
Page 72

aggregated by **Count**) using the **Local Systems** search tells you how many hosts within one hop of your 3D Sensors are running each operating system.



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 upwards-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 additions or movement from the most recent results:

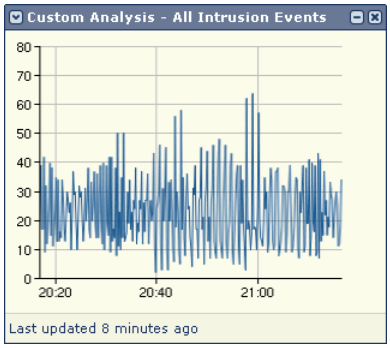
- The new event icon (●) signifies that the event is new to the results.
- The up-arrow icon (↑) indicates that the event has moved up in the standings since the last time the widget updated. A number indicating how many places the event has moved up appears next to the icon.
- The down-arrow icon (↓)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

Keysight, EX-1004
Page 73

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.

IMPORTANT! Depending on how they are configured, Custom Analysis widgets can 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 74
- Viewing Associated Events from the Custom Analysis Widget on page 80
- Custom Analysis Widget Limitations on page 81

Configuring the Custom Analysis Widget

Platform: DC/MDC/Standalone Sensor

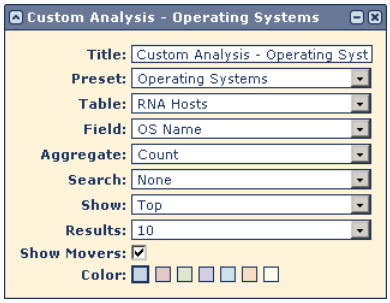
Component: None

As with all widgets, the Custom Analysis widget has preferences that determines its behavior. To configure a Custom Analysis widget, show the preferences as described in Understanding Widget Preferences on page 66.

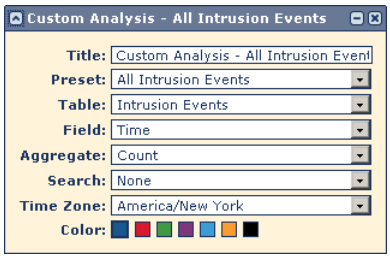
A different set of preferences appears depending on whether you configure the widget to show relative occurrences of events (that is, a bar graph), or you configure the widget to show a graph over time (that is, a line graph).

Keysight, EX-1004
Page 74

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	the title of the widget. If you do not specify a title, the appliance uses the configured event type as the widget title.
Preset	the preset for the widget. The Custom Analysis widget is delivered with several 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 you can create a custom configuration. For a detailed list of presets, see the Custom Analysis Widget Presets table on page 77.

Keysight, EX-1004
Page 75

Custom Analysis Widget Preferences (Continued)

Use this preference...	To control...
Table	the table of events which contains the event data the widget displays.
Field	the specific field of the event type you want to display. **TIP!** To display a graph over time, select **Time**.
Aggregate	the aggregation method for the widget. The aggregation method configures how the widget groups the data it displays. For most event types, the default aggregation criterion is Count.
Search	the saved search you want to use to further constrain the data that the widget displays. You do not have to specify a search, although some presets use predefined searches.
Show	whether you want to display the most frequently occurring events (**Top**) or the least frequently occurring events (**Bottom**).
Results	the number of results rows you want to display. You can display from 10 to 25 result rows, in increments of five.
Show Movers	whether you want to display the icons that indicate additions or movement from the most recent results.
Time Zone	which time zone you want to use to display results. The time zone appears whenever you select a time-based field.
Color	the color of the bars in the widget background that show the relative number of occurrences of each result.

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

Keysight, EX-1004
Page 76

each preset. (The predefined dashboards on the Master Defense Center and 3D Sensor do not include Custom Analysis widgets.)

Custom Analysis Widget Presets

Preset	Description	Predefined Dashboards	Requires
All Intrusion Events	Displays a graph of the total number of intrusion events on your monitored network over the dashboard time range.	Default Dashboard Detailed Dashboard	DC/MDC/ Standalone Sensor + IPS
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	DC/MDC/ Standalone Sensor + IPS
Client Applications	Displays the most active client applications on your monitored network, by application type.	Detailed Dashboard	DC + RNA
Dropped Intrusion Events	Displays counts for the most frequently occurring intrusion events, by classification, where the packet was dropped.	Default Dashboard	DC/MDC/ Standalone Sensor + IPS
Flows by Initiator IP	Displays the most active hosts on your monitored network, based on the number of flows where the host initiated the session.	Flow Summary	DC + RNA
Flows by Port	Displays the most active ports on your monitored network, based on the number of detected flows.	Flow Summary	DC + RNA
Flows by Responder IP	Displays the most active hosts on your monitored network, based on the number of flows where the host was the responder in the session.	Flow Summary	DC + RNA
Flows by Service	Displays the most active services on your monitored network, based on the number of detected flows.	Flow Summary	DC + RNA

Keysight, EX-1004
Page 77

Custom Analysis Widget Presets (Continued)

Preset	Description	Predefined Dashboards	Requires
Flows over Time	Displays a graph of the total number of flows on your monitored network, over the dashboard time range.	Flow Summary	DC + RNA
Intrusion Events Requiring Analysis	Displays a count of intrusion event requiring analysis, based on event classification.	Detailed Dashboard	DC/MDC + IPS + RNA
Intrusion Events by Hour	Displays the most active hours of the day, based on frequency of intrusion events.	none	DC/MDC/ Standalone Sensor + IPS
Intrusion Events to High Criticality Hosts	Displays the most frequently occurring types of intrusion events, based on the number of intrusion events occurring on high criticality hosts.	Detailed Dashboard	DC/MDC + IPS + RNA
Operating Systems	Displays the most common operating system, based on the number of hosts running each operating system within your network.	Detailed Dashboard	DC + RNA
Services	Displays the most common RNA service vendors, based on the number of hosts on the network running services made by that vendor.	Detailed Dashboard	DC + RNA
Top Attackers	Displays the most active hosts on your monitored network, based on the number of intrusion events where the host was the attacking host in the flow that caused the event.	Default Dashboard	DC/MDC/ Standalone Sensor + IPS
Top Targets	Displays the most active hosts on your monitored network, based on the number of intrusion events where the host was the targeted host in the flow that caused the event.	Default Dashboard	DC/MDC/ Standalone Sensor + IPS

Keysight, EX-1004
Page 78

Custom Analysis Widget Presets (Continued)

Preset	Description	Predefined Dashboards	Requires
Traffic by Initiator IP	Displays the most active hosts on your monitored network, based on the number of kilobytes per second of data transmitted by the hosts.	Detailed Dashboard Flow Summary	DC + RNA
Traffic by Initiator User	Displays the most active RUA users on your monitored network, based on the total number of kilobytes of data received by the hosts where those users are logged in.	Detailed Dashboard	DC + RNA + RUA
Traffic by Port	Displays the most active responder ports on your monitored network, based on the number of kilobytes per second of data transmitted via the port.	Flow Summary	DC + RNA
Traffic by Responder IP	Displays the most active hosts on your monitored network, based on the number of kilobytes per second of data received by the hosts.	Detailed Dashboard Flow Summary	DC + RNA
Traffic by Service	Displays the most active services on your monitored network, based on the number of kilobytes per second of data transmitted by the service.	Detailed Dashboard Flow Summary	DC + RNA
Traffic over Time	Displays a graph of the total kilobytes of data transmitted on your monitored network over the dashboard time range.	Detailed Dashboard Flow Summary	DC + RNA

Keysight, EX-1004
Page 79

Custom Analysis Widget Presets (Continued)

Preset	Description	Predefined Dashboards	Requires
Unique Intrusion Events by Destination IP	Displays the most active targeted hosts, based on the number of unique intrusion events per targeted host.	none	DC/MDC/ Standalone Sensor + IPS
Unique Intrusion Events by Impact	Displays the number of unique intrusion event types associated with each impact flag level.	none	DC/MDC + IPS + RNA
White List Violations	Displays the hosts with the most white list violations, by violation count	Detailed Dashboard	DC + RNA

Viewing Associated Events from the Custom Analysis Widget

Platform: DC/MDC/Standalone Sensor

Component: None

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 dashbaord, 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 47 and Specifying Time Constraints in Searches on page 1556.

Keysight, EX-1004
Page 80

To view associated events from the Custom Analysis Widget:

Access: Any except Restricted or Database

▶ 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 flow 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:

- Viewing Audit Records on page 1969
- Viewing Intrusion Events on page 274
- Viewing RNA Network Discovery and Host Input Events on page 1012
- Viewing Hosts on page 1020
- Viewing Host Attributes on page 1034
- Viewing Services on page 1041
- Viewing Client Applications on page 1051
- Viewing Sourcefire Vulnerabilities on page 1057
- Viewing Third-Party Vulnerabilities on page 1065
- Viewing Flow Data on page 1093
- Viewing RUA Users on page 1490
- Viewing RUA Events on page 1501
- Viewing Compliance Events on page 1272
- Viewing White List Events on page 1197
- Viewing White List Violations on page 1204
- Viewing Health Events on page 1958
- Viewing the SEU Import Log on page 1859
- Working with Active Scan Results on page 1443
- Understanding Custom Tables on page 1562

Custom Analysis Widget Limitations

Platform: DC/MDC/Standalone Sensor

Component: None

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

Keysight, EX-1004
Page 81

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, Intrusion Event Analysts cannot view RNA 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, the Master Defense Center does not store flow data. If your dashboard includes a Custom Analysis widget that displays data that 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.

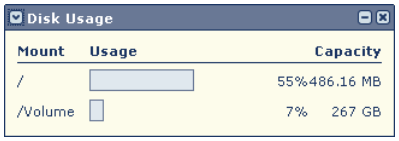
You enable or disable the Custom Analysis widget from the Dashboard settings in your system policy. For more information, see Configuring Dashboard Settings on page 1746.

Understanding the Disk Usage Widget

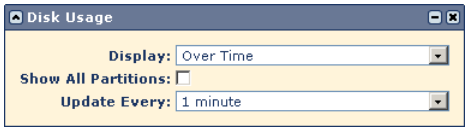
Platform: DC/MDC/3D Sensor

Component: None

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.



You can configure the widget to display just the 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 66.

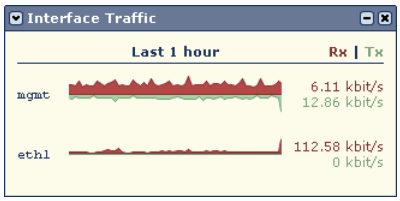
Keysight, EX-1004
Page 82

Understanding the Interface Traffic Widget

Platform: DC/MDC/3D Sensor

Component: None

The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's interfaces over the dashboard time range. Note that only 3D Sensors have interfaces other than the management interface.



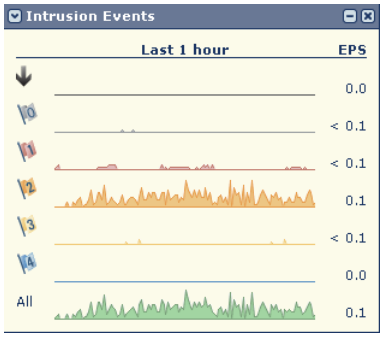
The widget preferences control how often the widget updates. On 3D Sensors, the preferences also control whether the widget displays the traffic rate for unused interfaces (by default, the widget only displays the traffic rate for interfaces that belong to an interface set). For more information, see Understanding Widget Preferences on page 66.

Understanding the Intrusion Events Widget

Platform: DC/MDC/Standalone Sensor

Component: IPS

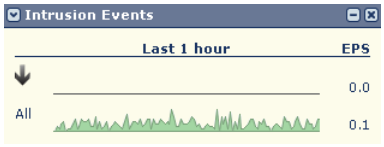
The Intrusion Events widget shows the rate of intrusion events that occurred over the dashboard time range. On the Defense Center and Master Defense Center, this includes statistics on intrusion events with dropped packets and different impacts.



On the 3D Sensor, the widget can display statistics for dropped intrusion events, all intrusion events, or both. Note that for managed 3D Sensors, you must enable

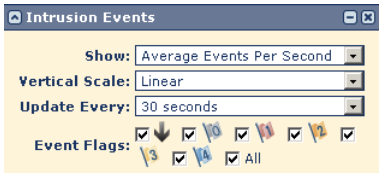
Keysight, EX-1004
Page 83

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 and Master Defense Center, but not the 3D Sensor, you can configure the widget to display intrusion events with dropped packets and different impacts by modifying the widget preferences. You can display dropped events on Defense Centers, Master Defense Centers, and 3D Sensors. The following graphic shows the Defense Center version of the widget preferences.

Note that IPS does not drop packets in a passive deployment regardless of the rule state or the inline drop behavior of the intrusion policy. For more information, see Setting Rule States on page 435 and Setting Drop Behavior in an Inline Deployment on page 358.



In the widget preferences, you can:

- on a Defense Center or Master Defense Center, select one or more **Event Flags** check boxes to display separate graphs for events with 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 307 for more information.
- select **Show** to choose **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 66.

On the Intrusion Events widget, you can:

- on a Defense Center or Master Defense Center, click a graph corresponding to dropped packets or a specific impact to view intrusion events of that type
- click the graph corresponding to dropped events to view 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 274.

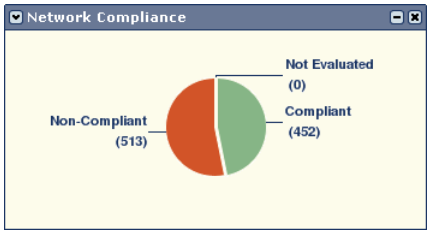
Keysight, EX-1004
Page 84

Understanding the Network Compliance Widget

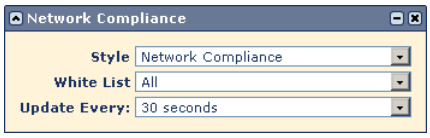
Platform: DC

Component: RNA

The Network Compliance widget summarizes your hosts' compliance with the compliance white lists you configured (see Using RNA as a Compliance Tool on page 1156). 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 compliance policies.



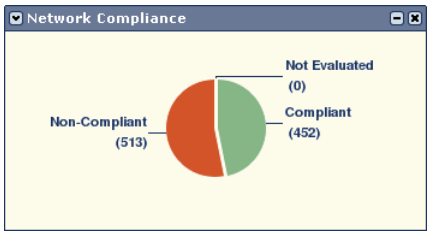
You can configure the widget to display network compliance either for all white lists, or for a specific white list, by modifying the widget preferences.



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 compliance policy.

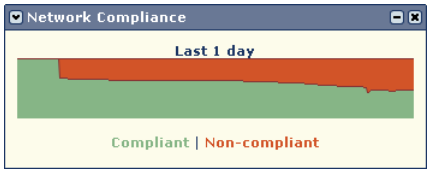
You can also use the widget preferences to specify which of three different styles you want to use to display network compliance.

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 1204.

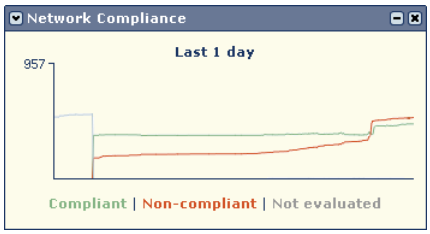


Keysight, EX-1004
Page 85

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 66.

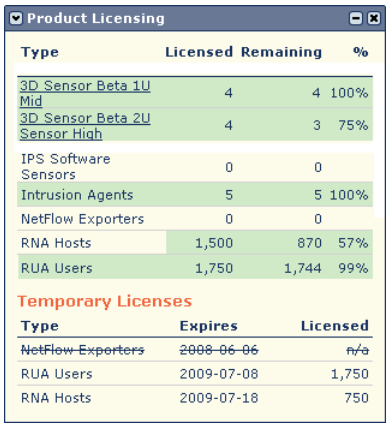
Keysight, EX-1004
Page 86

Understanding the Product Licensing Widget

Platform: DC

Component: None

The Product Licensing widget shows the sensor 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.



The top section of the widget displays all of the sensor and feature licenses installed on the Defense Center, including temporary licenses, while the Temporary Licenses section displays only temporary and expired licenses. For example, if you have two feature licenses for RNA 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 an RNA Hosts feature license with 1500 licensed hosts, while the Temporary Licenses section displays an RNA 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 66.

You can click any of the license types to go to the License page of the System Settings and add or delete feature licenses. For more information, see Understanding Licenses on page 1780.

Keysight, EX-1004
Page 87

Understanding the Product Updates Widget

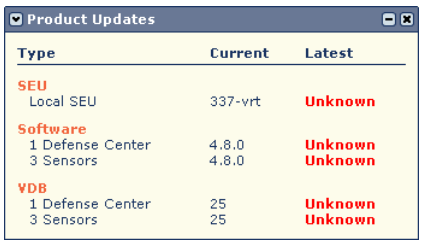
Platform: DC/MDC/3D Sensor

Component: None

The Product Updates widget provides you with a summary of the software (Sourcefire 3D System software, SEU, and VDB) currently installed on the appliance as well as information on available updates that you have downloaded, but not yet installed, for that software.

Note that the widget displays **Unknown** as the latest version of the software unless you have configured a scheduled task to download, push, or install software updates; the widget uses scheduled tasks to determine the latest version. For more information, see Scheduling Tasks on page 1699.

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 sensors. Note that you cannot update the VDB on a sensor or a Master Defense Center.



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 66.

On the Product Updates widget, you can:

- manually update an appliance by clicking the current version of the Sourcefire 3D System software, SEU, or VDB.
  - to update the system software or VDB, see Updating System Software on page 1831.
  - to update the SEU, see Importing SEUs and Rule Files on page 1848.
- create a scheduled task to download the latest version of the Sourcefire 3D System software, SEU, or VDB by clicking either the latest version or the **Unknown** link in the Latest column; see Scheduling Tasks on page 1699.

Keysight, EX-1004
Page 88

Understanding the RSS Feed Widget

Platform: DC/MDC/3D Sensor

Component: None

The RSS Feed widget adds an RSS feed to a dashboard. By default, the widget shows a feed of Sourcefire company news.



You can also configure the widget to display a preconfigured feed of Sourcefire security news, 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

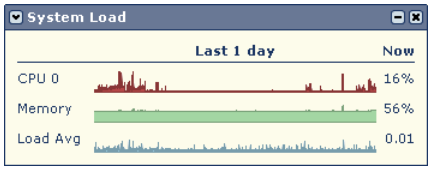
Keysight, EX-1004
Page 89

Understanding the System Load Widget

Platform: DC/MDC/3D Sensor

Component: None

The System Load widget shows the CPU usage (for each CPU), memory (RAM) usage, and system load (also called the load average, measured by the number of processes waiting to execute) on the appliance, both currently and over the dashboard time range.



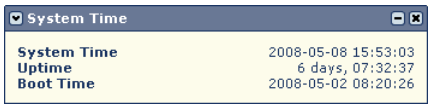
You can configure the widget to show or hide the load average by modifying the widget preferences. The preferences also control how often the widget updates. For more information, see Understanding Widget Preferences on page 66.

Understanding the System Time Widget

Platform: DC/MDC/3D Sensor

Component: None

The System Time widget shows the local system time, uptime, and boot time for the appliance.



You can configure the widget to hide the boot time by modifying the widget preferences. The preferences also control how often the widget synchronizes with the appliance's clock. For more information, see Understanding Widget Preferences on page 66.

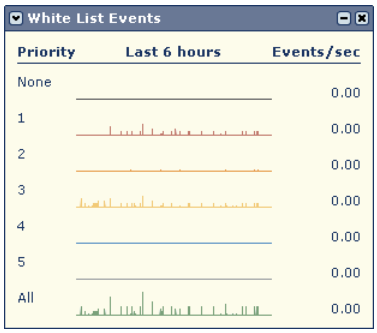
Keysight, EX-1004
Page 90

Understanding the White List Events Widget

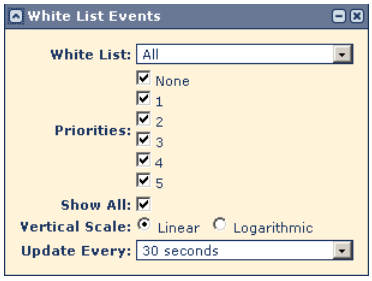
Platform: DC/MDC

Component: RNA

The White List Events widget shows the average events per second by priority, over the dashboard time range.



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 66.

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 1197.

Keysight, EX-1004
Page 91

Working with Dashboards

Platform: DC/MDC/3D Sensor

Component: None

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

You manage dashboards on the Dashboard List page (see Viewing Dashboards on page 94). You can create, view, modify, export, and delete dashboards.



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 Admin 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 52.

For more information on working with dashboards, see:

- Creating a Custom Dashboard on page 92
- Viewing Dashboards on page 94
- Modifying Dashboards on page 96
- Deleting a Dashboard on page 102
- Exporting a Dashboard on page 2000

Creating a Custom Dashboard

Platform: DC/MDC/3D Sensor

Component: None

When you create a new dashboard, you can choose to base it on any pre-existing dashboard, including the Sourcefire default dashboard, or on any user-defined dashboard. This makes a copy of the pre-existing 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 pre-existing 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

Keysight, EX-1004
Page 92

refreshed. This can 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, 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 and 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 3D Sensor or Master Defense Center may display some invalid, disabled widgets. For more information, see Importing and Exporting Objects on page 1997.

To create a new dashboard:

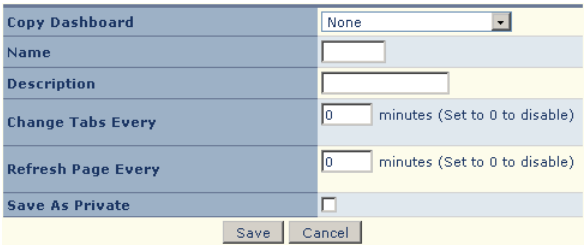
Access: Any except Restricted or Database

1. Select **Analysis & Reporting > Dashboards**.

   If you have a default dashboard defined, it appears. If you do not have a default dashboard defined, the Dashboard List page appears.

2. In either case, click **New Dashboard**.

   The **New Dashboard** page appears.



Keysight, EX-1004
Page 93

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 pre-existing dashboard, by rearranging and deleting widgets). For more information, see Modifying Dashboards on page 96.

Viewing Dashboards

Platform: DC/MDC/3D Sensor

Component: None

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

Keysight, EX-1004
Page 94

Dashboard List page, where you can choose a dashboard to view. To view the details of all available dashboards, click **Dashboards** from the Dashboard toolbar.

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 51 and Specifying Your Default Dashboard on page 52.

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 the 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 the 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.

IMPORTANT! Although your session normally logs you out after 3.5 hours of inactivity, this will not happen while you are viewing a dashboard, unless the dashboard is paused.

Keysight, EX-1004
Page 95

To view a dashboard:

Access: Any except Restricted or Database

- Select **Analysis & Reporting > 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 **Dashboards** menu on the toolbar.
  - If you do not have a default dashboard defined, the Dashboard List page appears. Click **View** next to the dashboard you want to view.

  The dashboard you selected appears.

To change the dashboard time range:

Access: Any except Restricted or Database

- 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: Any except Restricted or Database

- On the time range control, click the pause icon ( ).

  The dashboard is paused until you unpause it.

To unpause the dashboard:

Access: Any except Restricted or Database

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

  The dashboard is unpaused.

Modifying Dashboards

Platform: DC/MDC/3D Sensor

Component: None

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.

Keysight, EX-1004
Page 96

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.

IMPORTANT! Any user, 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.

For more information, see the following sections

- Changing Dashboard Properties on page 97
- Adding Tabs on page 98
- Deleting Tabs on page 98
- Renaming Tabs on page 99
- Adding Widgets on page 99
- Rearranging Widgets on page 101
- Minimizing and Maximizing Widgets on page 101
- Deleting Widgets on page 101

Changing Dashboard Properties

Platform: DC/MDC/3D Sensor

Component: None

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: Any except Restricted or Database

1. Select **Analysis & Reporting > Dashboards**.

   If you have a default dashboard defined, it appears; continue with the next step.

   If you do not have a default dashboard defined, the Dashboard List page appears; skip to step 3.

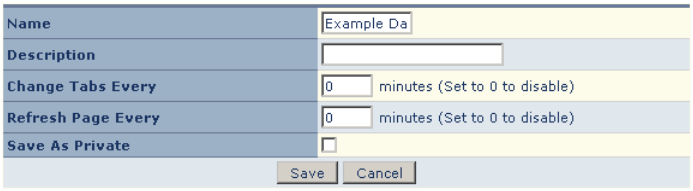
2. On the toolbar, click **Dashboards**.

   The Dashboard List page appears.

Keysight, EX-1004
Page 97

3. Click **Edit** next to the dashboard whose properties you want to change.

   The Edit Dashboard page appears. See Creating a Custom Dashboard on page 92 for information on the various configurations you can change.



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

   The dashboard is changed.

Adding Tabs

Platform: DC/MDC/3D Sensor

Component: None

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

To add a tab to a dashboard:

Access: Any except Restricted or Database

1. View the dashboard where you want to add a tab.

   For more information, see Viewing Dashboards on page 94.

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 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 99.

   The new tab is added.



   You can now add widgets to the new tab. For more information, see Adding Widgets on page 99.

Deleting Tabs

Platform: DC/MDC/3D Sensor

Component: None

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.

Keysight, EX-1004
Page 98

To delete a tab from a dashboard:

Access: Any except Restricted or Database

1. View the dashboard where you want to delete a tab.
   For more information, see Viewing Dashboards on page 94.
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

Platform: DC/MDC/3D Sensor

Component: None

Use the following procedure to rename a dashboard tab.

To rename a tab:

Access: Any except Restricted or Database

1. View the dashboard where you want to rename a tab.
   For more information, see Viewing Dashboards on page 94.
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 and click **OK**.
   The tab is renamed.

Adding Widgets

Platform: DC/MDC/3D Sensor

Component: None

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 left-most 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 101.

Keysight, EX-1004
Page 99

To add a widget to a dashboard:

Access: Any except Restricted or Database

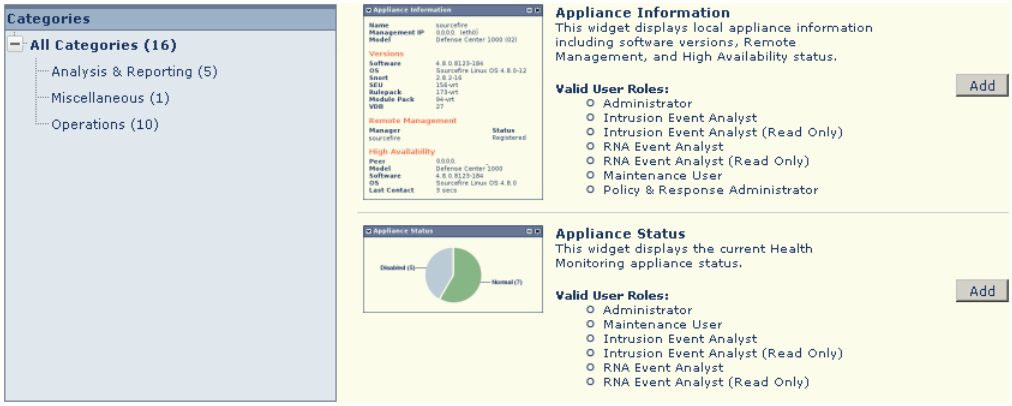
1. View the dashboard where you want to add a widget.

   For more information, see Viewing Dashboards on page 94.

2. Select the tab where you want to add the widget.

3. Click **Add Widgets**.

   The Add Widgets page appears.



   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.



Keysight, EX-1004
Page 100

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

Platform: DC/MDC/3D Sensor

Component: None

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: Any except Restricted or Database

- Click the title bar of the widget you want to move, then drag it to its new location.

Minimizing and Maximizing Widgets

Platform: DC/MDC/3D Sensor

Component: None

You can minimize widgets to simplify your view, then maximize them when you want to see them again.

To minimize a widget:

Access: Any except Restricted or Database

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

To maximize a widget:

Access: Any except Restricted or Database

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

Deleting Widgets

Platform: DC/MDC/3D Sensor

Component: None

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

Keysight, EX-1004
Page 101

To delete a widget:

Access: Any except Restricted or Database

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

Platform: DC/MDC/3D Sensor

Component: None

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 52.

To delete a dashboard:

Access: Any except Restricted or Database

1. Select **Analysis & Reporting > Dashboards**.

   If you have a default dashboard defined, it appears; continue with the next step.

   If you do not have a default dashboard defined, the Dashboard List page appears; skip to step 3.
2. On the toolbar, click **Dashboards**.

   The Dashboard List page appears.
3. Click **Delete** next to the dashboard you want to delete.
4. Confirm that you want to delete the dashboard.

   The dashboard is deleted.

Keysight, EX-1004
Page 102

Chapter 3
Using the Defense Center

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 sensors that are a part of 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 sensors, you can:

- configure policies for all your sensors from a single location, making it easier to change configurations
- install various types of software updates on sensors
- push health policies to your managed sensors and monitor their health status from the Defense Center

The Defense Center aggregates and correlates intrusion events, network discovery information, and sensor performance data, allowing you to monitor the information that your sensors are reporting in relation to one another, and to assess the overall activity occurring on your network.

IMPORTANT! Some of the components in the Sourcefire 3D System (such as the Virtual 3D Sensors, Intrusion Agents, and Crossbeam-based software sensors) do not provide a web interface that you can use to view events or manage policies. You must use a Defense Center if your deployment includes any of these products.

Keysight, EX-1004
Page 103

See the following sections for more information about using the Defense Center to manage your sensors:

- Management Concepts on page 104 describes some of the features and limitations involved with managing your sensors with a Defense Center.
- Working in NAT Environments on page 117 describes the principles of setting up the management of your sensors in Network Address Translation environments.
- Working with Sensors on page 118 describes how to establish and disable connections between sensors and your Defense Center. It also explains how to add, delete, and change the state of managed sensors, and how to reset management of a sensor.
- Managing Sensor Groups on page 137 describes how to create sensor groups as well as how to add and remove sensors from groups.
- Editing a Managed Sensor's System Settings on page 139 describes the sensor attributes you can edit and explains how to edit them.
- Managing a Clustered Pair on page 146 describes how to create a clustered pair of 3D Sensors and how to remove 3D Sensors from a cluster.
- Configuring High Availability on page 151 describes how to set up two Defense Centers as a high availability pair to help ensure continuity of operations.

Management Concepts

You can use a Defense Center to manage nearly every aspect of a sensor's behavior. Use only a single Defense Center to manage your sensor unless you are using a second Defense Center as a 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.

- Understanding the Benefits of Managing Your Sensors on page 105
- What Can Be Managed by a Defense Center? on page 106
- Understanding Software Sensors on page 110
- Beyond Policies and Events on page 116
- Using Redundant Defense Centers on page 117

Keysight, EX-1004
Page 104

Understanding the Benefits of Managing Your Sensors

There are several benefits to using a Defense Center to manage your sensors.

- Use the Defense Center as a central point of management.

  The Defense Center' s web interface can accomplish nearly any task that can be done from its managed sensor's local web interface, and can apply that action to one or many sensors managed by the Defense Center.

  As an example, you can create an intrusion policy on the Defense Center and apply it to all your managed 3D Sensors with IPS. Similarly, you can create and apply RNA appliance and detection polices to managed 3D Sensors with RNA. This eliminates the laborious task of having to replicate and apply potentially thousands of intrusion rules and detection policies.

  You can also create and apply system policies to your managed sensors. A system policy controls several appliance-level settings such as the login banner and the access control list. Because most of the sensors in your deployment are likely to have similar settings in the system policy, you can create the policy on the Defense Center and push it to the appropriate sensors instead of replicating it locally.
- Intrusion events and RNA events are automatically sent to the Defense Center.

  You can view the events from a single web interface instead of having to log into each sensor's interface to view the events there. You can also generate reports based on events from multiple sensors.
- If your Defense Center manages sensors with IPS and RNA, and those sensors view the same network traffic, then the Defense Center can correlate the intrusion events it receives with the information about hosts that RNA provides. The Defense Center can then assign impact flags to each intrusion event. The impact flag indicates how likely it is that an intrusion attempt will affect its target.
- You can use your Defense Center to configure external authentication through an Lightweight Directory Access Protocol (LDAP) or Remote Authentication Dial In User Service (RADIUS) server. You can use user information from an external server to authenticate users on your Sourcefire 3D System appliances. By pushing a system policy with configured authentication objects to your sensor, you push the external authentication object to the sensor. External authentication cannot be managed on the sensor, so you must use the Defense Center to manage it.
- The Defense Center includes a health monitoring featuthat you can use to check the status of critical functionality across your Sourcefire 3D System deployment. You can take advantage of health monitoring by applying health policies to each of your managed sensors and then reviewing the health data that they send back to the Defense Center. You can also apply a health policy to the Defense Center to monitor its health.

Keysight, EX-1004
Page 105

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 the following devices:

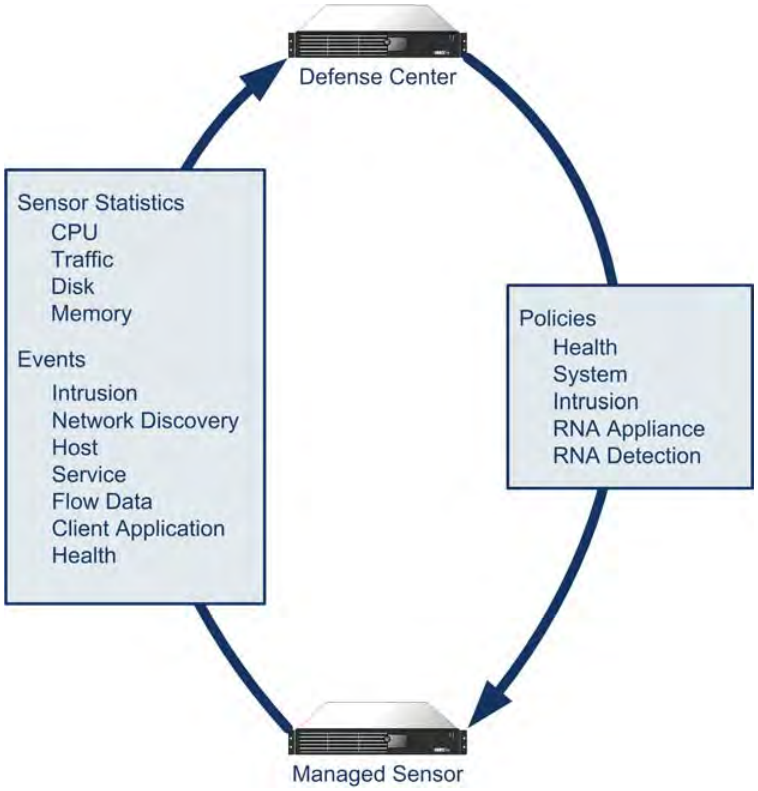
- Sourcefire 3D Sensors
- 3D Sensor Software for Crossbeam Systems X-Series
- Intrusion Agents on various platforms

IMPORTANT! Sourcefire recommends that you manage no more than three 3D Sensors with the DC500 model Defense Center. You can also use a DC500 to manage Sourcefire 3D Sensor software on approved platforms, as well as intrusion agents and RNA software on approved platforms. For details on DC500 database limitations see Database Event Limits on page 1748.

When you manage a sensor (or a software sensor), information is transmitted between the Defense Center and the sensor over a secure, SSL-encrypted TCP tunnel.

Keysight, EX-1004
Page 106

The following illustration lists what is transmitted between a Sourcefire Defense Center and its managed sensors. Note that the types of events and policies that are sent between the appliances are based on the sensor type.

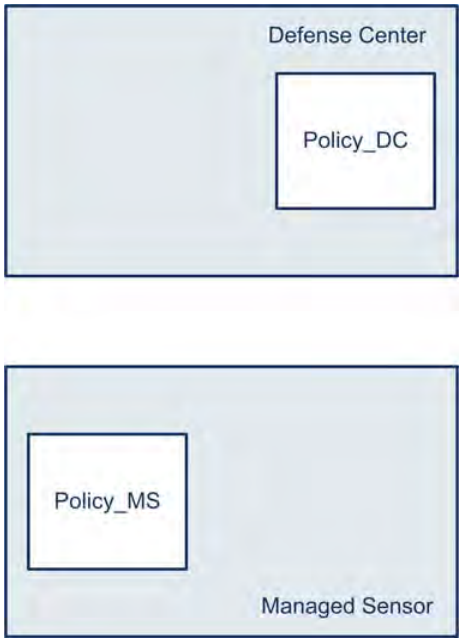


If you apply a policy on a sensor before you begin managing it with a Defense Center, you can see a read-only version of the policy on the Defense Center's web interface.

Similarly, after you set up communications with a Defense Center and apply policies from the Defense Center to your sensor, you can see a read-only version of the running policies on the sensor's web interface. The following graphics

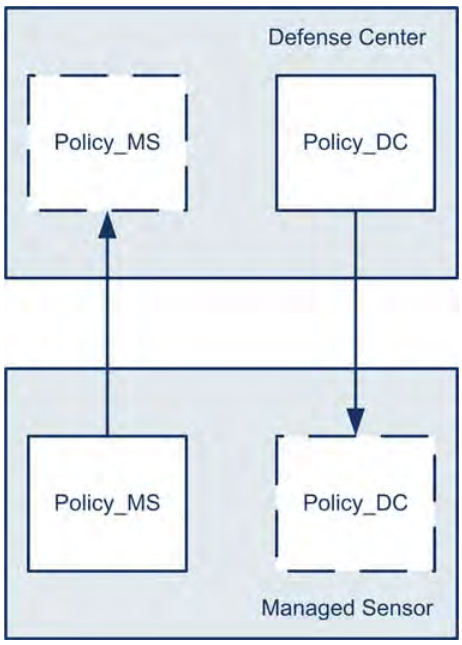
Keysight, EX-1004
Page 107

illustrate this process. First, before you set up sensor management, each appliance has its own policies:



Keysight, EX-1004
Page 108

Then, after communications are set up, read-only versions of running policies (represented by the dotted lines) are available:



The appliance where you originally create a policy is the policy's "owner" and is identified that way if you view the policy on a different appliance. For example, the following graphic shows the Detection Engine page on a 3D Sensor with IPS. The Sample Intrusion Policy that is currently applied to the sensor's two detection engines was created on the Defense Center (pine.example.com).

Name	DE Type	Resources	Interface Set	Set Type	Policy			
DE1	IS	1	passive interfaces	passive	Sample Policy (Remotely Authored by pine.example.com)	Edit	Delete	Variables
DE2	IS	1	passive interfaces	passive	Sample Policy (Remotely Authored by pine.example.com)	Edit	Delete	Variables

If you want to edit a policy, you must do it on the appliance where the policy was created.

TIP! After you set up management with a Defense Center, Sourcefire recommends that you use **only** the Defense Center's web interface to view events and manage policies for your managed sensors.

Keysight, EX-1004
Page 109

The following user-created data and configurations are retained locally on the sensor and are **not** shared with the Defense Center:

- user accounts
- user preferences
- bookmarks
- saved searches
- custom workflows
- report profiles
- audit events
- syslog messages
- reviewed status for intrusion events (IPS only)
- contents of the clipboard (IPS only)
- incidents (IPS only)

If you create custom fingerprints on the Defense Center, they are automatically shared with managed 3D Sensors with RNA.

Also note that operations you perform on data on one appliance are not transmitted to other appliances. For example, if you delete an intrusion event from the Defense Center, the event remains on the sensor that discovered it. Similarly, deleting an intrusion event from a sensor does not delete it from the Defense Center.

Understanding Software Sensors

Several of the sensors you can manage with a Defense Center are software-based sensors. A *software-based sensor* is a software-only installation of Sourcefire 3D System sensor software. The following Sourcefire 3D System sensors are software-based:

- Intrusion Agents for various platforms - for more information, see Managing Intrusion Agents on page 111
- 3D5800, 3D3800, and 3D9800 sensors - for more information, see Managing 3Dx800 Sensors on page 112.
- RNA for Red Hat Linux - for more information, see Managing RNA Software for Red Hat Linux on page 114
- 3D Sensor Software with RNA for Crossbeam X-Series - for more information, see Managing 3D Sensor Software with RNA for Crossbeam on page 115
- 3D Sensor Software with IPS for Crossbeam X-Series - for more information, see Managing 3D Sensor Software with IPS for Crossbeam on page 116

Software-based sensors do not have a user interface on the sensor; they can only be managed from a Defense Center. In addition, some of the functionality in the

Keysight, EX-1004
Page 110

Defense Center interface cannot be used with software-based sensors. For some software-based sensors, certain aspects of functionality are managed through the operating system or other features on the appliance.

Managing Intrusion Agents

The Sourcefire Intrusion Agent transmits events generated by open source Snort sensor installations to the Sourcefire Defense Center. These events can then be viewed along with data from 3D Sensors with IPS so you can easily analyze all the intrusion information gathered on your network.

The Defense Center cannot apply intrusion policies to the Intrusion Agent. You must tune your Snort rules and options manually on the computer where the Intrusion Agent resides. Also, high availability is not supported on Intrusion Agents.

IMPORTANT! When using Intrusion Agents registered to Defense Centers configured for high availability and managed by a Master Defense Center, register all Intrusion Agents to the primary Defense Center.

Keysight, EX-1004
Page 111

See the Supported Features for Intrusion Agents table for more information.

Supported Features for Intrusion Agents

Supported through Defense Center	Supported through CLI and .conf files	Not Supported
• Intrusion event collection and management • Licensing • Reports generated on the Defense Center	• Process management • Registration of remote manager • Rules tuning	• Detection engine management • Event storage on sensor • Health policy apply • High availability synchronization • Host Statistics • Interface set management • Intrusion policy apply • Network interface management • Network settings • Performance Statistics • Remote backup and restore • Remote reports • Sensor information management (System Settings) • SEU updates • Software updates • System policy apply • Time settings

Managing 3Dx800 Sensors

Although Version 4.10 is not supported on the Sourcefire 3D3800, 3D5800, and 3D9800 sensor models (referred to as the 3Dx800 sensors), you can manage these sensors with a Defense Center if those sensor models are using Version 4.9.x. However, because these models do not have a web interface and because configuration and event data cannot be stored on the sensors, certain features cannot be used with these sensors.

Keysight, EX-1004
Page 112

Features Supported through the Defense Center

The following features are supported on all 3Dx800 sensors through the Defense Center:

- detection engine management
- health policy apply
- high availability synchronization
- host statistics
- interface set management
- intrusion policy apply (no OPSEC support)
- intrusion event collection and management
- licensing
- performance statistics (may be underreported because of multiple detection resources)
- process management
- reports generated on the Defense Center
- sensor information management (System Settings)
- SEU updates
- software updates
- system policy apply
- time settings

The following features are supported on only the 3D3800 and 3D5800 sensor models through the Defense Center:

- compliance policy apply
- RNA and compliance event collection and management
- RNA detection policy apply
- VDB updates

Features Supported through the CLI

The following features are supported through the sensor's command line interface:

- network interface management
- network settings
- registering the remote manager

Keysight, EX-1004
Page 113

Unsupported Features

The following features are unsupported on 3Dx800 sensors:

- custom fingerprinting
- event storage on sensor
- remote backup and restore

Managing RNA Software for Red Hat Linux

Although RNA for Red Hat Linux is not supported on Version 4.10, you can manage Version 4.9.x of these software sensors with a Defense Center. However, not all of the features function in the same manner.

Features Supported through the Defense Center

The following features are supported on all RNA for Red Hat Linux sensors through the Defense Center:

- compliance policy apply
- detection engine management
- high availability synchronization
- host statistics
- interface set management
- licensing
- performance statistics
- reports generated on the Defense Center
- RNA and compliance event collection and management
- RNA detection policy apply
- sensor information management (system settings)
- VDB updates

Features Supported through the CLI

The following features are supported through the sensor's command line interface:

- network interface management
- network settings
- process management
- registering the remote manager
- time settings

Keysight, EX-1004
Page 114

Unsupported Features

The following features are unsupported on RNA for Red Hat Linux sensors:

- custom fingerprinting
- event storage on sensor
- health policy apply
- remote backup and restore
- remote reports
- applying a system policy

Managing 3D Sensor Software with RNA for Crossbeam

3D Sensor Software with RNA for Crossbeam provides many of the features found on 3D Sensors with RNA. However, not all of the features function in the same manner. See the Supported Features for RNA on Crossbeam table for more information.

Supported Features for RNA on Crossbeam

Supported through Defense Center	Supported through Crossbeam X-Series CLI	Not Supported
• Compliance policy apply • Detection engine management • High availability synchronization • Host Statistics • Interface set management • Licensing • Performance Statistics • Reports generated on the Defense Center • RNA detection policy apply • RNA and compliance event collection and management • Sensor information management (in System Settings) • Software updates • VDB updates	• Backup and restore • Network interface management • Network settings • Process management • Registration of remote manager • Time settings	• Custom fingerprinting • Event storage on sensor • Health policy apply • Remote backup and restore • Remote reports • System policy apply • Subnet detection

Keysight, EX-1004
Page 115

Managing 3D Sensor Software with IPS for Crossbeam

3D Sensor Software with IPS for Crossbeam provides many of the features found on 3D Sensors with IPS. However, because the Crossbeam sensors do not have a user interface and because configuration and event data cannot be stored on the sensors, certain features cannot be used with this software. See the Supported Features for IPS on Crossbeam table for more information.

Supported Features for IPS on Crossbeam

Supported through Defense Center	Supported through Crossbeam X-Series CLI	Not Supported
• Detection engine management • High availability synchronization • Host Statistics • Interface set management • Intrusion policy apply • Intrusion event collection and management • Licensing • Performance Statistics • Reports generated on the Defense Center • SEU updates • Sensor information management (in System Settings) • Software updates	• Backup and restore • Network interface management • Network settings • Process management • Registration of remote manager • Time settings	• Custom fingerprinting • Event storage on sensor • Health policy apply • Remote backup and restore • Remote reports • System policy apply

Beyond Policies and Events

In addition to applying policies to sensors and receiving events from them, you can also perform other sensor-related tasks on the Defense Center.

Backing Up a Sensor

If you are storing event data on your sensor 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 sensor. See Performing Sensor Backup with the Defense Center on page 1991 for more information.

Keysight, EX-1004
Page 116

Running Remote Reports

You can create a report profile on the Defense Center and run it remotely using the data on a managed sensor. This is particularly useful if you want to generate a report for the audit events on a managed sensor. Audit events are stored locally and are not sent to the Defense Center, but you can design a report on the Defense Center, select a managed sensor, and run the report. If you set up the report so that it is automatically emailed to you, you do not even need a user account on the sensor to read the resulting report. See Working with Event Reports on page 1531 for more information.

Updating Sensors

From time to time, Sourcefire releases updates to the Sourcefire 3D System, including:

- Security Enhancement Updates (SEUs), which can contain new and updated intrusion rules, as well as new and updated preprocessors and protocol decoders
- vulnerability database updates
- software patches and updates

You can use the Defense Center to install an update on the sensors it manages.

Using Redundant Defense Centers

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 151 or more information.

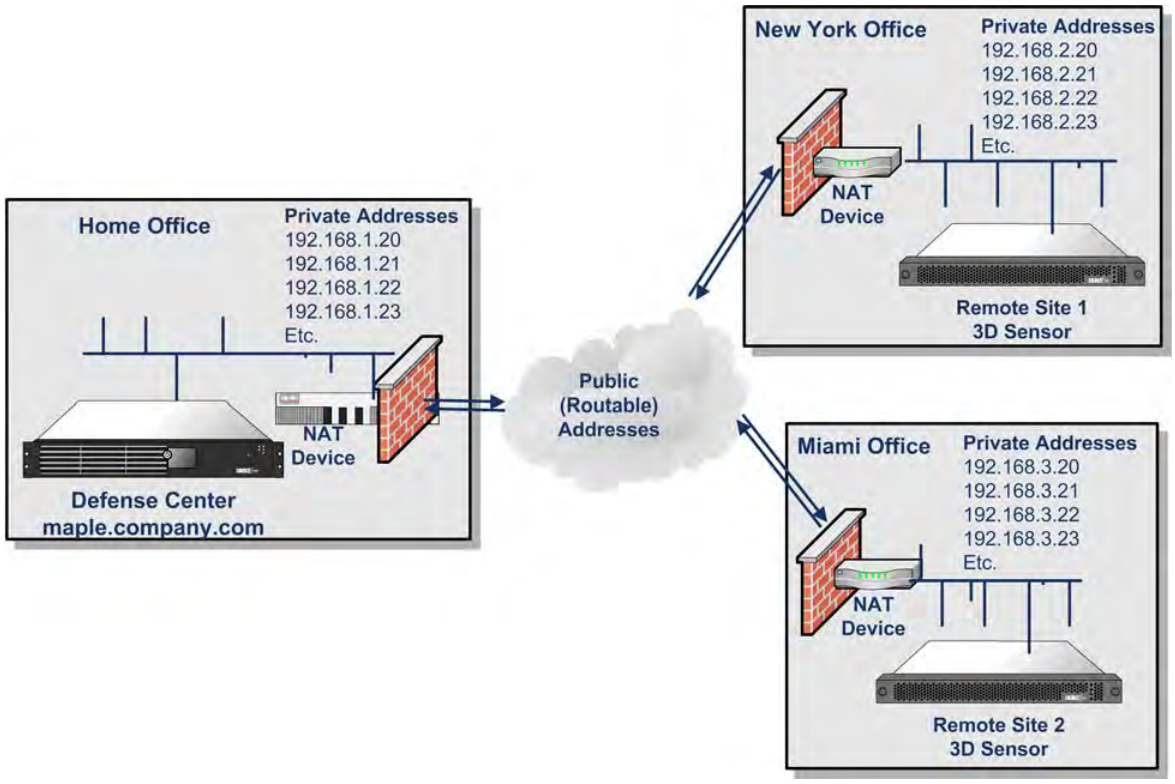
Working in NAT Environments

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 an appliance, you establish connections between appliances and register the appliances with one another. If you establish that communication in an environment without NAT, the two required pieces of common information during registration are the registration key and the unique IP address or the fully qualified domain name of the host. If you establish that communication in an environment with NAT, the two required pieces of common information during registration are the registration key and the unique NAT ID.

Keysight, EX-1004
Page 117

In the following example diagram, when you set up the remote office 3D Sensors connections to the home office, use the Defense Center's fully qualified domain name as its host name. You can use thesame registration key, you can use `snort` when adding either sensor, because the registration key does not have to be unique. However, you must use a **unique** NAT ID when adding the New York 3D Sensor to the Defense Center, and then use a different **unique** NAT ID when adding the Miami 3D Sensor. Each NAT ID has to be unique among all NAT IDs used to register sensors on the Defense Center.



Working with Sensors

When you manage a sensor, you set up a two-way, SSL-encrypted communication channel between the Defense Center and the sensor. The Defense Center uses this channel to send information (in the form of policies) to the sensor about how you want to analyze your network traffic. As the sensor evaluates the traffic, it generates events and sends them to the Defense Center

Keysight, EX-1004
Page 118

using the same channel. You can create the following policies on your Defense Center and apply them to managed sensors:

- health policies
- system policies
- RUA policies
- RNA detection policies
- intrusion policies

There are several steps to managing a sensor with a Defense Center:

1. Begin by setting up a communications channel between the two appliances.

   This is a two-step process, with procedures that you need to perform on each side of the communications channel. See Adding Sensors to the Defense Center on page 122 for more information. (Deleting Sensors on page 126 explains how to remove a sensor from the Defense Center.)

2. Create the appropriate policies on the Defense Center and apply them to the sensor or to the appropriate detection engines on the sensor.
   - IPS detection engines require an intrusion policy that determines which types of attacks 3D Sensor with IPS detect. See Creating Intrusion Policies on page 331 for more information.
   - RNA detection engines require an RNA detection policy, which controls the networks that 3D Sensors with RNA monitor. See What is an RNA Detection Policy? on page 889 for more information.
   - You can also create and apply system policies, which control certain appliance-level features on your sensors. Note that the system policy applied to the Defense Center controls the types of RNA events that are logged to the database. See Managing System Policies on page 1732 for more information.
   - You can create and apply health policies that allow you to monitor the processes and status of your sensors. See Configuring Health Policies on page 1895 for more information.

3. Confirm that you are receiving the events generated by your sensors. See Viewing Intrusion Event Statistics on page 269 and Viewing RNA Event Statistics on page 996 for more information.

Note that the procedure for managing a 3Dx800 sensor differs from the procedure for managing other sensors. See Managing a 3Dx800 Sensor on page 131 for more information.

The process for setting up communications between the Defense Center and other products such as the Crossbeam-based software sensors, RNA Software for Red Hat Linux, and the Intrusion Agents are slightly different. Refer to the configuration guides for those products for more information.

Many sensor management tasks are performed on the Sensors page and are described in Understanding the Sensors Page on page 120.

Keysight, EX-1004
Page 119

Understanding the Sensors Page

The Sensors page (**Operations > Sensors**) provides you with a range of information and options that you can use to manage your sensors (including software-based sensors), intrusion agents, and sensor groups.

The following sections describe some of the features on the Sensors page.

Virtual Sensor Count



When you manage Virtual 3D Sensors from the Defense Center, the field for a **Virtual Sensor** count appears above the sensor list on the Sensors page. For details about Virtual 3D Sensors, see the *Virtual Defense Center and 3D Sensor Installation Guide*.

Sort-by Drop-Down List

Use this drop-down list to sort the Sensors page according to your needs. You can sort by:

- Group (that is, sensor group; see Managing Sensor Groups on page 137)
- Model (that is, the sensor model)

Sensor List

The first column lists the hostname, sensor type, sensor model, and software version for each sensor. You can click the folder icon next to the name of the category to expand and contract the list of sensors. If you use a clustered pair of sensors, they are designated in the sensor list by a peer icon.



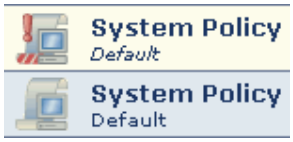
When you move your cursor over the peer icon, you can see which sensors are paired and if you configured the sensor as a master or a slave.

Health Policy

The next column lists the health policy for the sensor, if one has been applied. You can click the name of the health policy to view a read-only version of the policy. See Editing Health Policies on page 1931 for information about modifying an existing health policy.

Keysight, EX-1004
Page 120

System Policy



The next column lists the currently applied system policy. The policy name and the icon for the system policy in the top row highlight a special feature of the Sensors page. If a policy has a different icon and its name is in italics, that indicates the policy was modified after it was applied to the sensor. The icon and the name of the policy in the bottom row indicate that the version applied to the sensor is up to date. Note that this is the case for any policy that you create and apply from the Defense Center.

As with the health policy, you can click the name of the system policy to view a read-only version. See Managing System Policies on page 1732 for more information.

Status Icons



The status icons indicate the state of a sensor. The green check mark icon indicates that the sensor and the Defense Center are communicating properly. The red exclamation point icon indicates that the Defense Center has not received communications from the sensor in the last three minutes. If you hover your cursor over the icon, a pop-up window indicates the amount of time (in hours, minutes, and seconds) since the last contact. If the Defense Center has not received a communication from a sensor within the last two minutes, it sends a two-byte heartbeat packet to establish contact and ensure that the communications channel is still running. If your network is constrained in bandwidth, you can contact technical support to change the default time interval.

Edit and Delete Icons



Click the Edit icon next to a sensor if you want to change the sensor's current system settings. The system settings include the storage settings for the sensor, the time, the remote management configuration, and access to the processes for stopping and restarting the sensor or its software. See Editing a Managed Sensor's System Settings on page 139 for more information.

If you sort your Sensors page by sensor group, you can click the Edit icon next to the name of a sensor group to modify the list of sensors that belong to the group. See Editing Sensor Groups on page 138 for more information.

Keysight, EX-1004
Page 121

Click the Delete icon next to a sensor if you no longer want to manage the sensor with the Defense Center. See Deleting Sensors on page 126 for more information.

If you sort your Sensors page by sensor group, you can click the Delete icon next to the name of a sensor group to remove the sensor group from the Defense Center. See Deleting Sensor Groups on page 139 for more information.

Adding Sensors to the Defense Center

When you manage a sensor, you set up a two-way, SSL-encrypted communication channel between the Defense Center and the sensor. The Defense Center uses this channel to send information about how you want to analyze your network traffic (in the form of policies) to the sensor. As the sensor evaluates the traffic, it generates events and sends them to the Defense Center using the same channel. You can create the following policies on your Defense Center and apply them to managed sensors:

- system policies, which control appliance-level configurations such as database limits, DNS cache settings, and custom login banners
- RNA detection policies, which control RNA data-gathering behavior and determine which networks are monitored which detection engines
- intrusion policies, which control how protocol decoders and preprocessors are configured and which intrusion rules are enabled
- health policies, which monitor the health of your managed sensors

Note that before you add sensors to a Defense Center, you must make sure that the network settings are configured correctly on the sensor. This is usually completed as part of the installation process, but you can refer to Configuring Network Settings on page 1800 for details.

You can also add Intrusion Agents to the Defense Center. For more information, see Adding Intrusion Agents on page 136 and the *Sourcefire Intrusion Agent Configuration Guide*.

IMPORTANT! If you registered a Defense Center and 3D Sensor using IPv4 and want to convert them to IPv6, you must delete and re-register the sensor.

Keysight, EX-1004
Page 122

To add a sensor, you need:

- the sensor's IP address or hostname (in the connection context "hostname" is the fully qualified domain name or the name that resolves through the local DNS to a valid IP address)
- the Defense Center's IP address or hostname
- to decide if you want to store the events generated by the sensor only on the Defense Center, or on both the Defense Center and the sensor

TIP! Set up the managed appliance first.

You must begin the procedure for setting up the management relationship between a Defense Center and a sensor on the sensor.

Three fields are provided for setting up communications between appliances:

- **Management Host** - for the hostname or IP address.
- **Registration Key** - for registration key.
- **Unique NAT ID** - for a unique alphanumeric ID. Refer to Working in NAT Environments on page 117 for more information.

Valid combinations include:

- **Management Host** and **Registration Key** used on both appliances
- **Registration Key** and **Unique NAT ID** used on the 3D Sensor with **Host**, **Registration Key,** and **Unique NAT ID** used on the Defense Center.
- **Management Host**, **Registration Key,** and **Unique NAT ID** used on the 3D Sensor with **Registration Key** and **Unique NAT ID** used on the Defense Center.

IMPORTANT! The **Management Host** or **Host** field (hostname or IP address) must be used on at least one of the appliances.

To add a sensor to a Defense Center:

Access: Admin

1. Log into the web interface of the sensor you want to add.
2. Select **Operations > System Settings**.

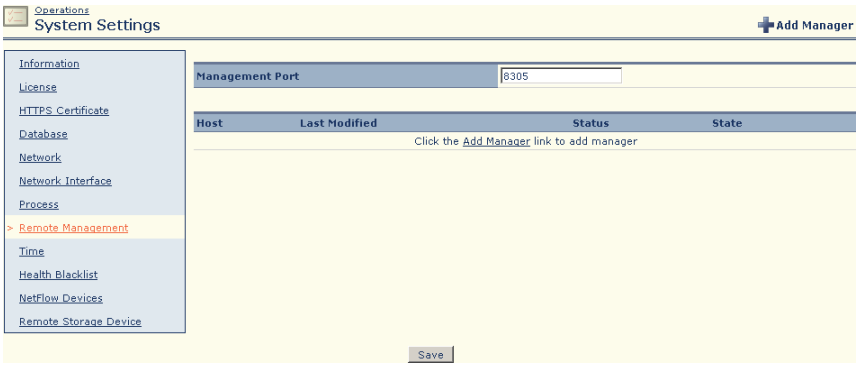
   The Information page appears.

Keysight, EX-1004
Page 123

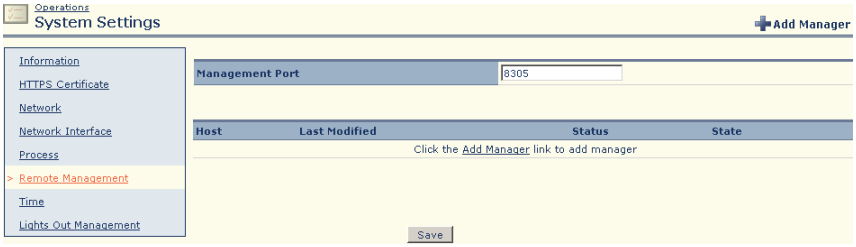
3. Click **Remote Management**.

   The Remote Management page appears.

   The figure below shows a Series 2 sensor version of the page.

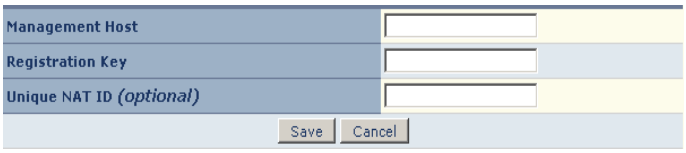


   The figure below shows an 8000 Series sensor version of the page.



4. Click **Add Manager**.

   The Add Remote Management page appears.



5. In the **Management Host** field, type the IP address or the host name of the Defense Center that you want to use to manage the sensor.

WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

TIP! You can leave the **Management 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.

Keysight, EX-1004
Page 124

6. In the **Registration Key** field, type the one-time use registration key that you want to use to set up a communications channel between the sensor and the Defense Center.
7. Optionally, in the **Unique NAT ID** field, type a **unique** alphanumeric ID that you want to use to identify the sensor.
8. Click **Save**.

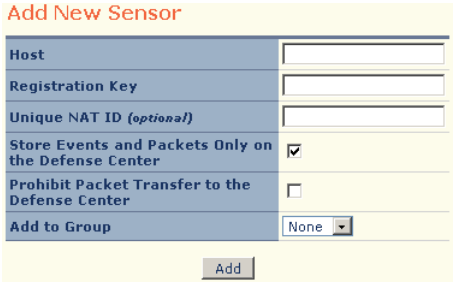
   After the sensor confirms communication with the Defense Center, the Pending Registration status appears.



9. Log into the Defense Center's web interface using a user account with Admin access, and select **Operations > Sensors**.

   The Sensors page appears.
10. Click **New Sensor**.

    The Add New Sensor page appears.
11. Type the IP address or the hostname of the sensor you want to add in the **Host** field.



WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

12. In the **Registration Key** field, enter the same registration key that you used in step 6.
13. If you used a NAT ID in step 7, enter the same ID in the **Unique NAT ID (optional)** field.

Keysight, EX-1004
Page 125

14. You can store data on both the Defense Center and the sensor by clearing the **Store Events and Packets Only on the Defense Center** check box. On 8000 Series sensors, you may store events on the sensor, but you must view events from the managing Defense Center.

    By default, data is stored only on the Defense Center and not on the sensor.

    IMPORTANT! Software-based sensors such as the 3D Sensor Software for Crossbeam cannot store data locally. You **must** store events on the Defense Center. For more information on supported functionality for software-based sensors, see Understanding Software Sensors on page 110.

15. You can prevent packet data from leaving a sensor by enabling the **Prohibit Packet Transfer to the Defense Center** check box.

    IMPORTANT! If you elect to prohibit sending packets and you do not store events on the 3D Sensor, packet data is not retained. Packet data is often important for forensic analysis.

16. To add the sensor to a group, select the group from the **Add to Group** list.

    For more information about groups, see Managing Sensor Groups on page 137.

17. Click **Add**.

    The sensor is added to the Defense Center. It can take up to two minutes for the Defense Center to verify the sensor's heartbeat and establish communication. You can view the sensor's status on the sensor's page (**Operations > Sensors**).

    IMPORTANT! In some high availability deployments where network address translation is used, you may need to use the Add Manager feature a second time to add the secondary Defense Center. Contact technical support for more information.

Deleting Sensors

If you no longer want to manage a sensor, you can delete it from the Defense Center. Deleting a sensor severs all communication between the Defense Center and the sensor. To manage the sensor again at a later date, you must re-add it to the Defense Center. To keep the sensor from trying to reconnect to the Defense Center, you should also delete the manager on the sensor.

If the Defense Center can no longer communicate with a detection engine on a managed sensor (for example, if the sensor is down or the network interface card is damaged), you should delete the managed sensor from the Defense Center

Keysight, EX-1004
Page 126

and then re-add it rather than try to delete the non-communicative detection engine.

IMPORTANT! If you delete a sensor 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 re-adding it. This interval ensures that the high availability pair re-synchronizes 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 sensor to both Defense Centers.

To delete a sensor from the Defense Center:

Access: Admin

1. Log into the Defense Center web interface and select **Operations > Sensors**.

   The Sensors page appears.

2. Click **Delete** next to the sensor you want to delete.

   Communication between the sensor and the Defense Center is discontinued and the sensor is deleted from the Sensors page.

3. Using a user account with Admin access, log into the web interface of the sensor you want to delete.

4. Select **Operations > System Settings**.

   The Information page appears.

5. Click **Remote Management**.

   The Remote Management page appears.

6. Click **Delete** next to the Defense Center where you want to reset management.

   The manager is removed. If the sensor has a system policy that causes it to receive time from the Defense Center via NTP, the sensor reverts to local time management.

Resetting Management of a Sensor

If communications fail between the Defense Center and one of your sensors, you can reset management of the sensor. If you want to manage a sensor with a different Defense Center, you must also reset management before adding the sensor to another Defense Center. You must first delete the manager on the

Keysight, EX-1004
Page 127

sensor and delete the sensor on the Defense Center. You can then re-add the manager on the sensor and then add the sensor to a Defense Center.

TIP! To temporarily disable communications between appliances without having to reset management, you can disable the manager on the sensor. For more information, see Managing Communication on a Managed Sensor on page 144.

The procedures for resetting management on the 3Dx800 sensors and on Crossbeam-based software sensors differ from the procedure for other sensors. For more information on resetting management on a 3Dx800 sensor, see Resetting Communications on the 3Dx800 on page 134. For more information on resetting management on a Crossbeam-based software sensor, see the *Sourcefire 3D Sensor Software for X-Series Installation Guide.*

To reset management:

Access: Admin

1. Log into the web interface of the Defense Center where you want to reset communications.
2. Select **Operations > Sensors**.

   The Sensors page appears.
3. Click **Delete** next to the sensor you want to delete.

   Communication between the sensor and the Defense Center is discontinued and the sensor is deleted from the Sensors page.

   If your sensor is no longer communicating with the Defense Center, you can delete the management on the sensor. If you attempt to delete management on the sensor while it is communicating with the Defense Center you will receive an error similar to:

   ```
   Delete failed. You must delete the appliance from its manager,
   maple.example.com.
   ```

To delete management on the sensor:

Access: Admin

1. Log into the web interface of the sensor where you want to reset communications.
2. Select **Operations > System Settings**.

   The Information page appears.
3. Click **Remote Management**.

   The Remote Management page appears.
4. Click **Delete** next to the Defense Center where you want to reset management.

   The manager is removed.

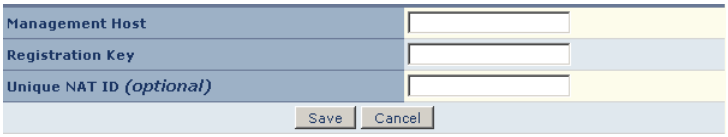
Keysight, EX-1004
Page 128

To re-add the sensor to the Defense Center:

Access: Admin

1. Log into the web interface of the sensor where you want to reset communications and click **Add Manager**.

   The Add Remote Management page appears.



2. In the **Management Host** field, type the IP address or the host name of the Defense Center that you want to use to manage the sensor.

   WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

   TIP! You can leave the Management 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.

3. In the **Registration Key** field, type the one-time use registration key that you want to use to set up a communications channel between the sensor and the Defense Center.
4. Optionally, in the **Unique NAT ID** field, type a **unique** ID that you want to use to identify the sensor.
5. Click **Save**.

   After the sensor confirms communication with the Defense Center, the Pending Registration status appears.



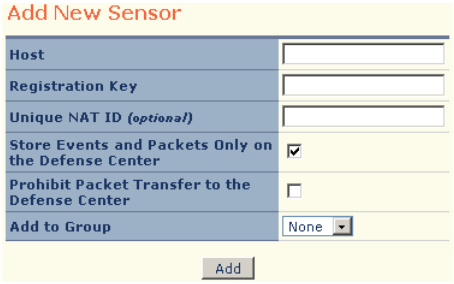
6. Log into the Defense Center's web interface using a user account with Admin access, and select **Operations > Sensors**.

   The Sensors page appears.

Keysight, EX-1004
Page 129

7. Click **New Sensor**.

   The Add New Sensor page appears.



8. Type the IP address or the hostname of the sensor you want to add in the **Host** field.

   WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

9. In the **Registration Key** field, type the same one-time use registration key that you used in step 3.
10. If you used a unique NAT ID in step 4, type the same value in the **Unique NAT ID** field.
11. You can store data on both the Defense Center and the sensor by clearing the **Store Events and Packets Only on the Defense Center** check box.

    By default, data is stored only on the Defense Center and not on the sensor.
12. You can prevent packet data from leaving a sensor by checking the **Prohibit Packet Transfer to the Defense Center** check box.

    IMPORTANT! If you elect to prohibit sending packets and you do not store events on the 3D Sensor, packet data is not retained. Packet data is often important for forensic analysis.

13. To add the sensor to a group, select the group from the **Add to Group** list.

    For more information about groups, see Managing Sensor Groups on page 137.

Keysight, EX-1004
Page 130

14. Click **Add**.

    The sensor is added to the Defense Center. It can take up to two minutes for the Defense Center to verify the sensor's heartbeat and establish communication. You can view the sensor's status on the Sensors page (**Operations > Sensors**).

IMPORTANT! In some high availability deployments where network address translation is used, you may need to use the Add Manager feature a second time to add the secondary Defense Center. Contact technical support for more information.

Managing a 3Dx800 Sensor

Because the Sourcefire 3D Sensor 3800, 3D Sensor 5800, and 3D Sensor 9800 (usually called the 3Dx800 sensors) do not have their own web interfaces, you must add them to a Defense Center as managed sensors so that you can perform procedures such as:

- creating and applying intrusion and RNA detection policies
- viewing events
- generating reports
- uploading and installing software updates

The following sections explain how to manage 3Dx800 sensors with a Defense Center:

- Managing 3Dx800 Sensors with a Defense Center on page 131
- Deleting a 3Dx800 Sensor from the Defense Center on page 133
- Resetting Communications on the 3Dx800 on page 134

Managing 3Dx800 Sensors with a Defense Center

Setting up communications between a 3Dx800 sensor and a Defense Center is a two-step process that involves setting up the sensor and then adding the sensor to the Defense Center.

This procedure assumes that you have completed the setup steps described in the sensor's *Installation Guide*.

To manage a 3Dx800 sensor with a Defense Center:

Access: Admin

1. Log into the 3D Sensor using the admin account.

   The CLI prompt appears.

   ```
   sensor.domain [admin]
   ```


Keysight, EX-1004
Page 131

2. Enter the following at the CLI prompt:

   ```
   [admin] configure sensor
   ```

3. Use the following command to determine whether remote management is already enabled:

   ```
   [admin:sensor] show management
   ```

   If management is already enabled, the sensor may be managed by another Defense Center. See Resetting Communications on the 3Dx800 on page 134 for information about deleting the sensor from the other Defense Center and preparing it for new management.

4. Use one of the following commands to enable management on the 3D Sensor:

   - If you are deploying your sensor in a network that **does not** use network address translation, enter the following command:

     ```
     [admin:sensor] set management enable ip_address reg_key
     ```

     where *ip_address* is the IP address of the Defense Center and *reg_key* is a unique single-use alphanumeric registration key. The IP address and registration key pair must uniquely identify the communications channel between the sensor and the Defense Center.

   - If you are deploying your sensor in a network that **does** use network address translation, enter the following command:

     ```
     [admin:sensor] set management enable NONE reg_key nat_id
     ```

     where NONE is a placeholder for the unresolvable IP address of the Defense Center, *reg_key* is a unique single-use alphanumeric registration key, and *nat_id* is a unique alphanumeric string. The NAT ID together with the registration key must uniquely identify the communications channel between the sensor and the Defense Center.

   In either case, a message appears indicating that remote management is enabled.

5. If you changed the management port on the Defense Center, you must change it on the 3Dx800 also:

   ```
   [admin:sensor] set management port port_number
   ```

   where *port_number* is the same port number you used on the Defense Center.

6. Use the following command to exit the CLI and return to the login prompt:

   ```
   [admin:sensor] exit
   ```

7. Using a user account with Admin access, log into the web interface of the Defense Center where you want to add the sensor.

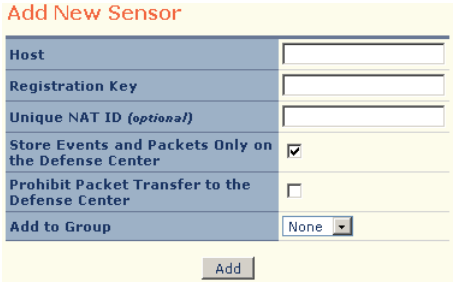
8. Select **Operations > Sensors**.

   The Sensors page appears.

Keysight, EX-1004
Page 132

9. Click **New Sensor**.

   The Add New Sensor page appears.



10. In the **Host** field, type the IP address or the hostname of the sensor you want to add.
11. In the **Registration Key** field, type the same one-time use registration key that you used on the sensor.
12. If you used a NAT ID in step 4, type the same value in the **Unique NAT ID** field.

IMPORTANT! Because 3Dx800 sensors do not have any local storage for events, make sure the **Store Events and Packets Only on the Defense Center** check box is selected.

---

13. You can prevent packet data from leaving a sensor by checking the **Prohibit Packet Transfer to the Defense Center** check box.

IMPORTANT! If you prohibit sending packets to the Defense Center, packet data, which is often important for forensic analysis, is not retained anywhere.

---

14. To add the sensor to a group, select the name of the group from the **Add to Group** list.

    For more information about groups, see Managing Sensor Groups on page 137.

15. Click **Add**.

    The 3Dx800 is added to the Defense Center.

    It can take up to two minutes for the Defense Center to verify the sensor's heartbeat and establish communication.

Deleting a 3Dx800 Sensor from the Defense Center

If you want to delete a 3Dx800 sensor from a Defense Center (for example, to manage it with a different Defense Center), you must complete a two-step process to disable remote management and then delete it from the Defense Center.

Keysight, EX-1004
Page 133

To delete a 3Dx800 sensor from a Defense Center:

Access: Admin

1. Log into the web interface of the Defense Center where you want to delete the sensor.
2. Select **Operations > Sensors**.

   The Sensors page appears.
3. Click **Delete** next to the sensor you want to delete.

   The sensor is deleted.
4. On the sensor, access the command prompt and use the admin account to log in.

   The CLI prompt appears.

   ```
   sensor.domain [admin]
   ```
5. Enter the following at the CLI prompt:

   ```
   [admin] configure sensor
   ```
6. Enter the following command to disable remote management:

   ```
   [admin:sensor] set management disable
   ```

   A message appears indicating that remote management is disabled.
7. Enter the following command to exit the CLI and return to the login prompt:

   ```
   [admin:sensor] exit
   ```

IMPORTANT! To add the sensor to either the same or a different Defense Center, you must re-enable remote management and then add the sensor to the Defense Center. For more information, see the next section, Resetting Communications on the 3Dx800.

Resetting Communications on the 3Dx800

If communication fails between a 3Dx800 sensor and the Defense Center that manages it, you can manually reset communications on the sensor.

To reset communications between the sensor and the Defense Center:

Access: Admin

1. Log into the web interface of the Defense Center that manages the sensor.
2. Select **Operations > Sensors**.

   The Sensors page appears.
3. Click **Delete** next to the sensor that is no longer communicating with the Defense Center.

   The sensor is deleted.

Keysight, EX-1004
Page 134

4. On the sensor, access the command prompt and use the admin account to log in.

   The CLI prompt appears.

   ```
   sensor.domain [admin]
   ```

5. Enter the following at the CLI prompt:

   ```
   [admin] configure sensor
   ```

6. Enter the following command to disable remote management:

   ```
   [admin:sensor] set management disable
   ```

   Remote management is disabled.

7. Use one of the following commands to enable remote management.

   - If your sensor is in a network that **does not** use network address translation, enter the following command:

     ```
     [admin:sensor] set management enable ip_address reg_key
     ```

     where *ip_address* is the IP address of the Defense Center and *reg_key* is a unique single-use alphanumeric registration key. The IP address and registration key pair must uniquely identify the communications channel between the sensor and the Defense Center.

   - If your sensor is in a network that **does** use network address translation, enter the following command:

     ```
     [admin:sensor] set management enable NONE reg_key nat_id
     ```

     where NONE is a placeholder for the unresolvable IP address of the Defense Center, *reg_key* is a unique single-use alphanumeric registration key, and *nat_id* is a unique alphanumeric string. The NAT ID together with the registration key must uniquely identify the communications channel between the sensor and the Defense Center.

   In either case, remote management is enabled again.

8. Enter the following command to exit the CLI and return to the login prompt:

   ```
   [admin:sensor] exit
   ```

9. On the Defense Center's Sensors page, re-add the sensor by clicking **New Sensor**.

   The Sensors page appears.

10. In the **Host** field, type the IP address or hostname of the sensor and make sure the **Store Events and Packets Only on the Defense Center** check box is selected.

11. Click **Add**.

    Communications are restarted and the sensor is re-added to the Defense Center.

Keysight, EX-1004
Page 135

Adding Intrusion Agents

The Add Agent page allows you to add an Intrusion Agent.

IMPORTANT! When using Intrusion Agents registered to Defense Centers configured for high availability and managed by a Master Defense Center, register all Intrusion Agents to the primary Defense Center.

To add an Intrusion Agent:

Access: Admin

1. Access the Defense Center web interface and select **Operations > Sensors**.

   The Managed Sensors page appears.

2. Click **New Agent**.

   The Agent Administration page appears.

3. In the **Name Of Agent** field, type an identifying name for the agent.

   This is the name that the Defense Center uses to identify the Intrusion Agent. It will appear on the event summary, event view pages, and reports.

4. In the **Hostname** or **IP Address** field, type the Intrusion Agent's host name (if DNS resolution is enabled on the Defense Center) or IP address.

   WARNING! If your Intrusion Agent sensor resides behind a NAT device, enter the IP address granted by the NAT device; that is, you should the IP address that the Defense Center will "see" when the Intrusion Agent attempts to communicate with it.

5. Click **Add Agent**.

   The Intrusion Agent is added and the page reloads, displaying a link that allows you to download authentication credentials.

6. Click **Download Auth Credentials** and save them for later use on the Intrusion Agent.

   To download authentication credentials, see Sensor Attributes - Intrusion Agent Page on page 136.

For information on the requirements for the intrusion agent side of the connection, see the *Sourcefire Intrusion Agent Configuration Guide*.

Sensor Attributes - Intrusion Agent Page

The Sensor Attributes page for Intrusion Agents allows you to view basic information about the Intrusion Agent and allows you to download authentication credentials. During configuration, you copy this file to the Intrusion Agent

Keysight, EX-1004
Page 136

appliance to allow the Intrusion Agent to authenticate with the Defense Center. Authentication credentials are unique to each Intrusion Agent appliance and Defense Center and cannot be copied from one appliance to another.

To download authentication credentials from the Sensor Attributes page:

Access: Admin

1. Access the Defense Center web interface and select **Operations > Sensors**.

   The Managed Sensors page appears.

2. Click **Edit** next to the Intrusion Agent.

   The System Settings page for the Intrusion Agent appears.

3. Click **Download Credential File**.

   You are prompted to download the credentials to your local computer.

For more information about copying the credentials, see the *Sourcefire Intrusion Agent Configuration Guide*.

Managing Sensor Groups

The Defense Center allows you to group sensors so that you can easily apply policies and install updates on multiple sensors.

See the following sections for more information:

- Creating Sensor Groups on page 137 explains how to create a sensor group on the Defense Center.
- Editing Sensor Groups on page 138 explains how to modify the list of sensors in a sensor group.
- Deleting Sensor Groups on page 139 explains how to delete a sensor group.

Creating Sensor Groups

Grouping managed sensors allows you to configure multiple sensors with a single system or health policy, and update multiple sensors with new software updates at the same time.

IMPORTANT! For information about Defense Center groups, see Managing Appliance Groups on page 191.

Keysight, EX-1004
Page 137

To create a sensor group and add sensors to it:

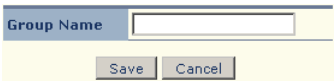
Access: Admin

1. On the Defense Center, select **Operations > Sensors**.

   The Sensors page appears.

2. Click **Create New Sensor Group**.

   The Create Sensor Group page appears.

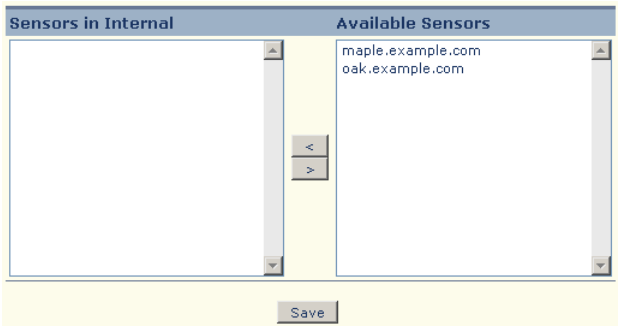


3. In the **Group Name** field, type the name of the group you want to create.
4. Click **Save**.

   The group is added.

5. To add sensors to the group, return to the Sensors page (**Operations > Sensors**) and click **Edit** next to the name of the sensor group.

   The Sensor Group Edit page appears.



6. Select the IP addresses or hostnames of the sensors you want to add from the Available Sensors list and click the arrow to move them into sensor group.
7. Click **Save**.

   The sensors are added to the group.

Editing Sensor Groups

You can change the set of sensors that reside in any sensor group. You must remove a sensor from its current group before you can add it to a new group.

Moving a sensor to a new group does not change its policy to the policy previously applied to the group. To change the sensor's policy, you must apply a new policy to the sensor or sensor group. See Applying an Intrusion Policy on page 347 for details.

Keysight, EX-1004
Page 138

To edit a sensor group:

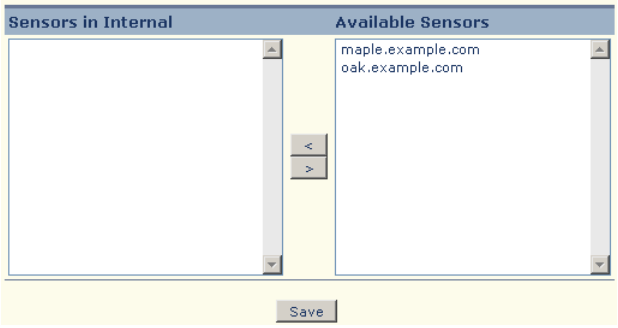
Access: Admin

1. On the Defense Center, select **Operations > Sensors**.

   The Sensors page appears.

2. Click **Edit** next to the sensor group you want to edit.

   The Sensor Group Edit page appears.



3. Select the sensor you want to move and click the arrow to add or remove it from the group.
   - To add a sensor to the group, select it from the **Available Sensors** list and click the arrow pointing toward the group you are editing.
   - To remove a sensor from a group, select it from the list in the group you are editing and click the arrow pointing to the **Available Sensors** list.
4. Click **Done**.

Deleting Sensor Groups

If you delete a group that contains sensors, the sensors are moved to Ungrouped on the Sensors page. They are not deleted from the Defense Center.

To delete a sensor group:

Access: Admin

1. Select **Operations > Sensors**.

   The Sensors page appears.

2. Click **Delete** next to the group you want to delete.

Editing a Managed Sensor's System Settings

Each sensor has a number of system settings. On an unmanaged sensor you can use the sensor's web interface to modify the settings as needed. When you

Keysight, EX-1004
Page 139

manage one or more sensors with a Defense Center, you can modify their system settings through the Defense Center's web interface.

IMPORTANT! You cannot edit the network settings or add a license file to a sensor through the Defense Center's web interface. You must perform those tasks on the sensor's web interface (generally before you begin to manage the sensor with the Defense Center). See Configuring System Settings on page 1775 for more information about system settings.

To edit the system settings for a managed sensor:

Access: Admin

1. On the Defense Center, select **Operations > Sensors**.

   The Sensors page appears.

2. Click **Edit** next to the name of the sensor where you want to edit the system settings.

   The Appliance page appears and includes a list of links on the left side of the page that you can use to navigate between pages.

3. From the System Settings page, you can:

   - view detailed information about the sensor. For more information, see Viewing a Sensor's Information Page on page 141.
   - modify the default settings for each network interface on the managed sensor. For more information, see Editing Network Interface Configurations on page 1803.

     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.

   - reboot or restart the processes on the managed sensor. For more information, see Stopping and Restarting a Managed Sensor on page 143.

Keysight, EX-1004
Page 140

- manage communications between the sensor and the Defense Center. For more information, see Managing Communication on a Managed Sensor on page 144.
- manage time settings on the managed sensor. For more information, see Setting the Time on a Managed Sensor on page 145.
- blacklist individual health policy modules on the managed sensor. For more information, see Blacklisting a Health Policy Module on page 1938.

Viewing a Sensor's Information Page

The Information page for a managed sensor includes the fields described in the Sensor Information table.

IMPORTANT! When you view the Information page for a managed Defense Center from the Master Defense Center's web interface, the fields are slightly different. See Editing Settings for a Managed Defense Center on page 185.

Sensor Information

Field	Description
Name	The assigned name for the managed sensor. Note that is the name of the sensor in the Defense Center web interface, not the hostname.
Product Model	The model name for the managed sensor.
Serial Number	The serial number of the chassis of the managed sensor.
Software Version	The version of the software currently installed on the managed sensor.
Store Events Only on Defense Center	Enable this check box to store event data on the Defense Center, but not the managed sensor. Clear this check box to store event data on both appliances.
Prohibit Packet Transfer to the Defense Center	Enable this check box to prevent the managed sensor from sending packet data with the events. Clear this check box to allow packet data to be stored on the DC with events.
Operating System	The operating system currently running on the managed sensor.

Keysight, EX-1004
Page 141

Sensor Information (Continued)

Field	Description
Operating System Version	The version of the operating system currently running on the managed sensor.
VDB Version	The version level of the vulnerability database currently loaded on the managed sensor.
IPv4 Address	The IPv4 address of the managed sensor.
IPv6 Address	The IPv6 address of the managed sensor.
Current Policies	The appliance-level policies currently applied to the managed sensor. If a policy has been updated since it was last applied, the name of the policy appears in italics. • The name of the current system policy is listed under **System**. • The name of the current health policy is listed under **Health**, if you applied one from the Defense Center that manages the sensor.
Status	An icon showing the current status of the managed sensor. If you hover your cursor over the icon, a pop-up message indicates how long it has been (in hours, minutes, and seconds) since the sensor communicated with the Defense Center. You can click **Refresh** to update the Status icon and its accompanying pop-up message.
Model Number	The model number for the sensor. This number can be important for troubleshooting.
Current Group	The sensor group that the sensor belongs to, if any. See Creating Sensor Groups on page 137 for more information.

To edit a managed sensor's settings:

Access: Admin

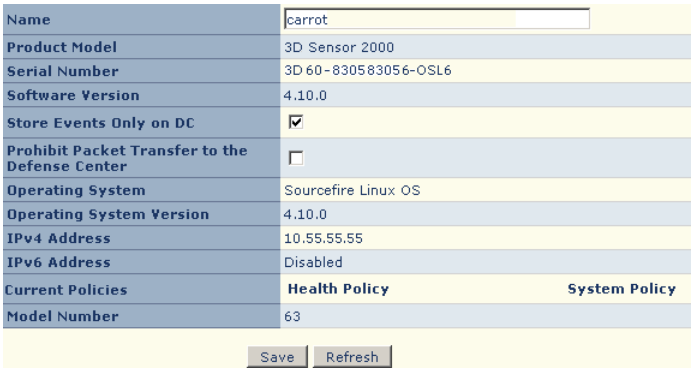
1. Select **Operations > Sensors**.

   The Sensors page appears.

Keysight, EX-1004
Page 142

2. Click **Edit** next to the name of the sensor whose system settings you want to edit.

   The Information page for that sensor appears. See the Sensor Information table on page 141 for a description of each field.



3. Change the sensor's attributes as needed.

   You can edit the following:

   - the sensor's hostname
   - where events generated by the sensor are stored
   - the group in which the sensor resides

   WARNING! Sensor host names must be made up of a combination of alphanumeric characters and should not be made up of numeric characters only.

4. Click **Save**.

   The updated sensor attributes are saved.

Stopping and Restarting a Managed Sensor

For 3D Sensors, you can reboot or restart the processes on a managed sensor using the Defense Center's web interface.

You must use the command line interface (CLI) to manage processes on Crossbeam-based software sensors, RNA Software for Red Hat Linux, and Intrusion Agents.

To shut down or restart a managed sensor:

Access: Admin

1. Select **Operations > Sensors**.

   The Sensors page appears.

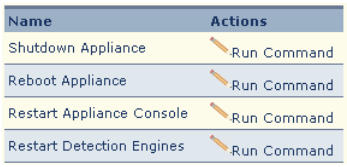
Keysight, EX-1004
Page 143

2. Click **Edit** next to the name of the sensor that you want to restart.

   The Information page for that sensor appears.

3. Click **Process** in the list to the left of the page.

   The Process page appears for your managed sensor.



4. Specify what command you want to perform:
   - If you want to shut down the sensor, click **Run Command** next to **Shutdown Appliance**.
   - If you want to reboot the sensor, click **Run Command** next to **Reboot Appliance**.
   - If you want to restart the software processes on the sensor, click **Run Command** next to **Restart Appliance Console**.
   - If you want to restart the Snort and RNA processes, click **Run Command** next to **Restart Detection Engines**.

WARNING! 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, you must press the power button on the appliance.

Managing Communication on a Managed Sensor

For most 3D Sensors, you can manage communications between a managed sensor and the Defense Center managing it using the Defense Center's web interface.

You must use the command line interface (CLI) to manage communication on 3Dx800 sensors, Crossbeam-based software sensors, RNA Software for Red Hat Linux, and Intrusion Agents.

To disable communications between the Defense Center and the sensor:

Access: Admin

1. Select **Operations > Sensors**.

   The Sensors page appears.

2. Click **Edit** next to the name of the sensor that you want to manage.

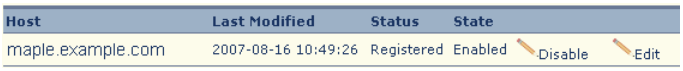
   The Information page for that sensor appears.

Keysight, EX-1004
Page 144

3. Click **Remote Management** in the list to the left of the page.

   The Remote Management page appears.

4. Click **Disable** next to the name of the sensor.



   Communications between the two appliances are interrupted.

> **TIP!** To enable communications between the two appliances again, click **Enable**.

For information about editing the remote management communications from a sensor see Configuring Remote Access to the Defense Center on page 1810.

Setting the Time on a Managed Sensor

If your managed sensor is receiving its time from an NTP server, which is the recommended setting for a managed sensor and its Defense Center, then you cannot change the time manually. See the NTP Status table on page 1814 for a description of the values you are likely to see for a sensor that is synchronized with an NTP server. However, if the system policy applied to the managed sensor allows you to set the time manually, then you can change it as part of the system settings.

For 3D Sensors, you can manage time settings on a managed sensor using the Defense Center's web interface.

You must use the command line interface (CLI) to manage time settings on Crossbeam-based software sensors and RNA Software for Red Hat Linux. You cannot manage time settings on Intrusion Agents.

To set the time for a managed sensor:

Access: Admin

1. Select **Operations > Sensors**.

   The Sensors page appears.

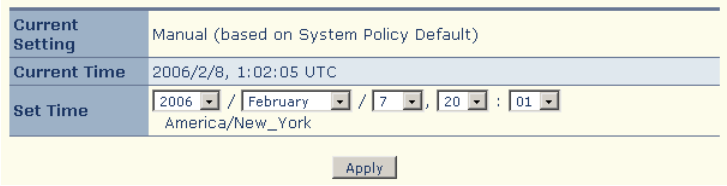
2. Click **Edit** next to the name of the sensor where you want to set the time.

   The Information page for that sensor appears.

Keysight, EX-1004
Page 145

3. Click **Time** in the list to the left of the page.

   The Time page appears showing the current time.



4. From the **Set Time** drop-down lists, select the following:
   - year
   - month
   - day
   - hour
   - minute
5. Click **Apply**.

   The time is updated.
6. 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.
7. Select your time zone and click **Save** and, after the time zone setting is saved, click **Close** to close the pop-up window.

IMPORTANT! Changing the time zone with this option is equivalent to changing the time zone using the Time Zone Settings option in the user preferences. In other words, this time zone option changes the time setting your user account uses on the Defense Center web interface. This setting does not affect the time zone setting on the managed sensor.

Managing a Clustered Pair

Platform: DC/3D9900/3D8250

Component: None

You can increase the amount of traffic inspected on a network segment by connectiong two fiber-based 3D9900 or two 3D8250 sensors in a clustered pair configuration. When you establish a clustered pair, you combine the resources of both sensors into a single, shared configuration.

To use the Defense Center to manage a clustered pair, both sensors in the pair must be running the same version of the Sourcefire 3D System. If the sensors in the clustered pair are no longer running the same version (for example, if an

Keysight, EX-1004
Page 146

upgrade on one of the sensors fails), the sensors no longer function as a clustered pair.

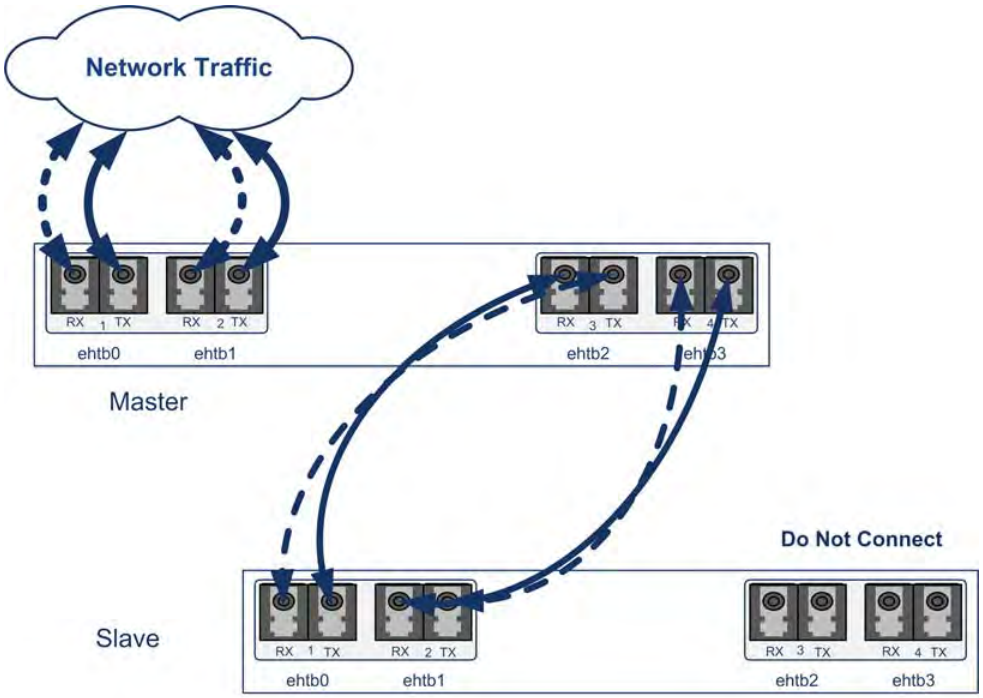
One sensor is designated as the *master* and the other is the *slave*. Connect the master to the network segment you wish to analyze in the same way you would configure a single sensor (either passive, inline, or inline with failopen). Connect the slave to the master as indicated by the clustering diagram and the Cluster Interconnect table. Use a Defense Center to establish the clustered pair relationship between the two sensors and manage their joint resources.

For information on the detection engines, interface set, and data from a clustered pair, see:

- Using Clustered 3D Sensors on page 250
- Identifying Detection Engines and Interface Sets on a Clustered Pair on page 251

The Defense Center manages the clustered pair, and local management is blocked on the shared portion of the clustered pair.

The following diagram shows interfaces on the master and slave 3D9900 sensors.



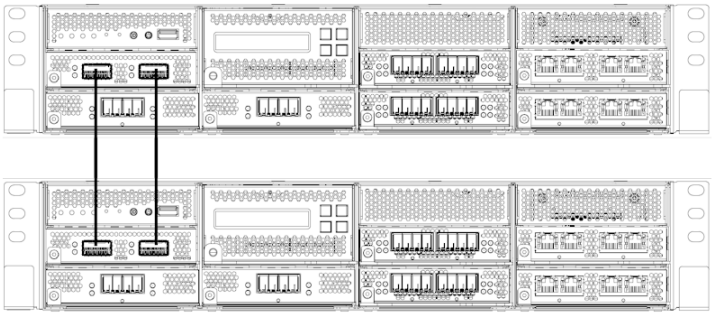
Keysight, EX-1004
Page 147

For information about the connections between the master and slave 3D9900 sensors, see the Cluster Interconnect table.

Cluster Interconnect

Master Interface	Slave Interface
ethb2 RX	ethb0 TX
ethb2 TX	ethb0 RX
ethb3 RX	ethb1 TX
ethb3 TX	ethb1 RX

The following diagram shows how to connect the interfaces joining the master and slave 3D8250 sensors.



After the relationship between the two sensors is established, the devices act like two separate sensors with a single, shared detection configuration. Local administration on the 3D Sensors is limited; the Defense Center manages detection configuration identically for the clustered pair.

If the master fails, traffic is handled according to the configuration of the master (either passive, inline, or inline with failopen). Although the slave continues to function, no traffic is passed to the slave. Two health alerts are generated, the first indicating that the master is not sending a heartbeat, and the second indicating that the slave is not receiving traffic.

If the slave fails, the master continues sense traffic and generate alerts. Although the master continues to function, the master alone may not have the capacity to handle the traffic, and some traffic may be dropped. One health alert is generated indicating that the slave is not sending a heartbeart.

After you establish the relationship, you cannot change which sensor is the master or slave unless you break and reestablish the relationship using the Defense Center.

Keysight, EX-1004
Page 148

For more information, see:

- Establishing a Clustered Pair on page 149
- Separating a Clustered Pair on page 151

Establishing a Clustered Pair

Platform: DC/3D9900/ 3D8250

Component: None

You can group two fiber-based 3D9900 or two 8000 Series sensors in a clustered pair to increase throughput. Before you begin, you must:

- decide which unit will be the master
- have SEU 2.8.6 or later loaded on your sensors and Defense Center
- cable the units properly prior to designating the master/slave relationship

Connect the 3D9900 master's ethb0 and ethb1 pair to the network. Connect the 3D9900 master's ethb2 and ethb3 pair to the 3D9900 slave's ethb0 and ethb1 pair as shown in the Cluster Interconnect table. Note that you **cannot** connect the 3D9900 slave's ethb2 and ethb3 pair when you establish the clustered pair.

On 8000 Series sensors, connect the left port on the master sensor to the left port on the slave sensor, and right port on the master sensor to the right port on the slave sensor.

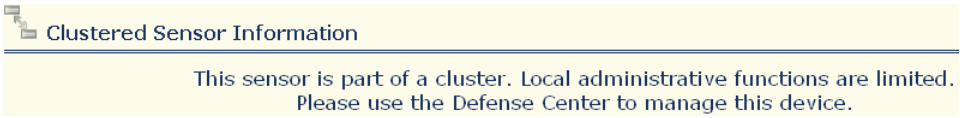
For more information about cabling, see the *Sourcefire 3D Sensor Installation Guide*.

After you establish the master/slave relationship, the detection engines and interface set are combined on the two sensors.

IMPORTANT! If you apply an RNA detection policy to the RNA detection engines on two different sensors and then establish clustering with those two sensors, you **must** edit and reapply your detection policy after you establish clustering.

There is one detection engine and interface set shared over the paired sensors. They are managed from the Defense Center, instead of the sensors.

If you attempt to manage the combined detection engines and interface set on the paired sensors, the following message appears.



Keysight, EX-1004
Page 149

To establish a clustered pair:

Access: Admin

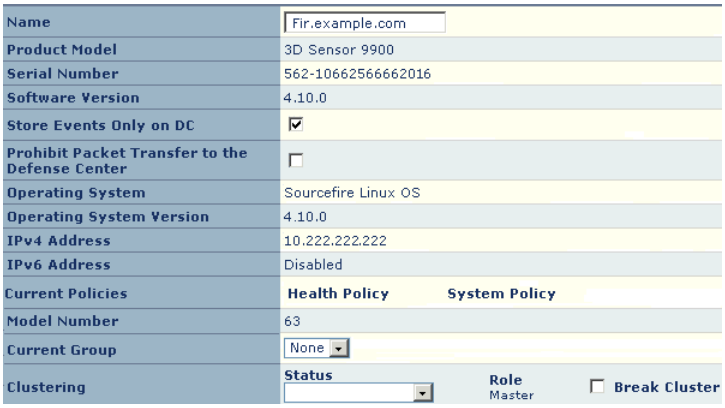
1. Select **Operations > Sensors** on your Defense Center.

   The Sensor page appears.

2. Click **Edit** next to the sensor that you cabled for master operation.

   TIP! If you edit a sensor that is not cabled as the master, you cannot perform the next series of steps.

   The System Settings page appears and there is a **Clustering** field at the bottom.



3. In the **Clustering** field, under status, select the sensor you want to form a cluster with. For example, if the other member of your pair is birch.example.com, select **Clustered with birch.example.com.** Note that you must form a cluster with a similar sensor: 3D9900 with 3D9900, or 3D8250 with 3D8250.

   Clustering is established and a confirmation message appears.

4. Review the confirmation message and confirm the correct the master/slave pair.

   IMPORTANT! While system verifies the cabling configuration, the sensing traffic is interrupted. If the system determines that the cabling is correct, it removes detection configurations (interface sets, detection engines) from the slave.

5. Click **OK** to confirm the master/slave pair.

Keysight, EX-1004
Page 150

6. After clustering is established, verify that the **Clustering** field changes to indicate the correct state.
   - On the master, the field reads: **Status Clustered** *sensor_name*, where *sensor_name* is the name of the sensor you designated as the slave in step 3 and **Role Master**.
   - On the slave, the field reads: **Status Clustered** and **Role Slave**

   3D9900 clustering is established. Use the managing Defense Center to establish the cluster's detection configurations for the interface set and detection engines.

Separating a Clustered Pair

Platform: DC/3D9900/ 3D8250

Component: None

If you no longer need to use the two sensors as a clustered pair, use the Defense Center to break the cluster.

To separate a clustered pair:

Access: Admin

1. Select **Operations > Sensors** on your Defense Center.

   The Sensor page appears.

2. Click **Edit** next to the sensor that you designated as the master sensor when you connected the pair's cables.

   The System Settings page appears with the **Clustering** field at the bottom.

3. Select **Break Cluster** in the **Clustering** field.



4. Click **Save**.
5. Review the confirmation message. Note the master/slave pair and click **OK** to confirm the master/slave that you want to separate the clustered pair.

   The sensors separate and the confirmation message disappears.

Configuring High Availability

To ensure the continuity of operations, the high availability feature allows you to designate redundant Defense Centers to manage 3D Sensors. Event data streams from managed sensors to both Defense Centers and certain configuration elements are maintained on both Defense Centers. If one Defense Center fails, you can monitor your network for intrusion events, RNA events, RUA

Keysight, EX-1004
Page 151

events, and compliance events without interruption using the second Defense Center.

See the following sections for more information about setting up high availability:

- Using High Availability on page 152 list the items that are and are not duplicated when you implement high availability.
- Guidelines for Implementing High Availability on page 157 outlines some guidelines you must follow if you want to implement high availability.
- Setting Up High Availability on page 158 explains how to specify primary and secondary Defense Centers.
- Monitoring the High Availability Status on page 160 explains how to check the status of your linked Defense Centers.
- Disabling High Availability and Unregistering Sensors on page 162 explains how to permanently remove the link between linked Defense Centers.
- Pausing Communication between Paired Defense Centers on page 163 explains how to pause communications between linked Defense Centers.
- Restarting Communication between Paired Defense Centers on page 163 explains how to restart communications between linked Defense Centers.

Using High Availability

Platform: DC

Component: None

The DC750/1500/3500, DC1000 and DC3000 models of the Defense Center support high availability configurations. The DC500 model of the Defense Center and the Virtual Defense Center **do not** support high availability.

Recommendations and Prerequisites

Prior to establishing a High Availability pair, please note the following recommendations and prerequisites:

- Sourcefire **strongly** recommends that both Defense Centers in an HA pair be the same model. Do **not** attempt to set up high availability between a Defense Center 1000 and a Defense Center 3000 or Defense Center 750/1500/3500.
- Sourcefire recommends that you change configurations **only** on the primary Defense Center and that you keep your secondary Defense Center as a backup.
- Both Defense Centers must have an `admin` account and each `admin` account must use the same password.

Keysight, EX-1004
Page 152

- No duplicate account names can exist between the two Defense Centers. If an identical account name exists on both Defense Centers, remove the duplicate user before establishing a high availability pair.
- If your deployment includes intrusion agents and you are also using a Master Defense Center to manage your linked Defense Centers, register all intrusion agents to the primary Defense Center.

For more information on:

- sensor attributes and user information shared in a high availability pair, see Sensor Configurations and User Information on page 153
- health and system policies shared in a high availability pair, see Health and System Policies on page 154
- feature license operation in a high availability pair, see Feature Licenses on page 155
- details of high availability pair operation, see Understanding High Availability on page 156

Sensor Configurations and User Information

Platform: DC

Component: None

Defense Centers in a high availability pair (also called an HA pair) share the following sensor attributes and user information:

- user account attributes and authentication configurations
- custom dashboards
- authentication objects for Sourcefire 3D System user accounts
- custom workflows
- custom tables
- sensor attributes, such as the sensor's host name, where events generated by the sensor are stored, and the group in which the sensor resides
- intrusion, RNA, and RUA detection engines
- intrusion policies and their associated rule states
- local rules
- custom intrusion rule classifications
- variable values and user-defined variables
- RNA detection policies
- RNA custom service detectors
- activated custom fingerprints
- host attributes
- traffic profiles

Keysight, EX-1004
Page 153

- RNA user feedback, including notes and host criticality; the deletion of hosts, services, and networks from the network map; and the deactivation or modification of vulnerabilities
- compliance policies and their associated rules
- compliance white lists
- SEU

To avoid launching duplicate responses and remediations when compliance policies are violated, Defense Centers do not share the associations between the policies and their responses and remediations.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 compliance policies. If the primary Defense Center fails, you should quickly associate your compliance policies with the appropriate responses and remediations on the secondary Defense Center to maintain continuity of operations. For more information, see Creating Compliance Policies on page 1262 and Configuring Remediations on page 1296.

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.

Health and System Policies

Platform: DC

Component: None

Health and system policies for Defense Centers and 3D Sensors are shared in high availability pairs. Allow enough time to ensure that 3D Sensor information about health policies, modules, blacklists, is synchronized on a newly activated Defense Center.

TIP! If you employ an HA paired Defense Center as a NTP server, the NTP function does not automatically switch. However, you can synchronize time with multiple alternative NTP servers. For 3D Sensors, you can point to one Defense Center as your first NTP server and the other Defense Center as your second NTP server. For more information, see Synchronizing Time on page 1768.

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.

Keysight, EX-1004
Page 154

Defense Centers in an HA 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

Feature Licenses

Platform: DC

Component: IPS/RNA/RUA/NetFlow

Defense Centers in an HA pair do **not** share RNA, RUA, and NetFlow licenses:

- Both Defense Centers must have RNA host licenses if you want to manage 3D Sensors with RNA with the high availability pair.
- While NetFlow data and devices are shared, the two Defense Centers must have enough NetFlow licenses to merge the list of devices on each, if you want to use NetFlow data to supplement the data gathered by your 3D Sensors with RNA.

  TIP! Both Defense Centers in a high-availability pair must have NetFlow licenses for at least the number of NetFlow-enabled devices you are using. If one Defense Center does not have a NetFlow license, it will not receive data from your NetFlow-enabled devices.

- While RUA LDAP authentication objects are shared, both Defense Centers must have RUA licenses if you want to manage 3D Sensors with RUA with the high availability pair.

  IMPORTANT! An RUA Agent can only connect to one Defense Center at a time. In an high-availability environment, if the primary Defense Center fails, you must make sure that your RUA Agents can communicate with the secondary Defense Center. For more information, see Configuring an RUA Agent on an Active Directory Server on page 1483.

Keysight, EX-1004
Page 155

Understanding High Availability

Platform: DC

Component: None

Although Defense Centers in high availability mode are named "primary" and "secondary," you can make policy or other changes to either Defense Center. 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.) However, during this ten-minute window, policies may appear incorrectly on the other Defense Center.

For example, if you create a policy on your primary Defense Center and apply it to a sensor that is also managed by your secondary Defense Center, the sensor could contact the secondary Defense Center before the Defense Centers contact each other. Because the sensor 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.

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, see Guidelines for Implementing High Availability on page 157.

Keysight, EX-1004
Page 156

Guidelines for Implementing High Availability

Platform: DC

Component: None

To take advantage of high availability, you must follow these guidelines:

- 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 sensors, and so on before you set up high availability.

  TIP! 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 sensors with it. To make sure the secondary Defense Center is in its original state, use the Restore CD to remove changed settings. Note that this also deletes event and configuration data from the Defense Center.

  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.

- By default, the Defense Centers use port 8305/tcp for communications. You can change the port as described in Configuring the Communication Channel on page 1807.
- Both Defense Centers must be running the same software version.
- Both Defense Centers must be running the same SEU version.
- The Defense Center software version must be the same or newer than the software version of managed 3D Sensors.
- All RNA software sensors managed by Defense Centers in high availability mode must be the same software version.
- If you use a Master Defense Center to manage a high-availability pair of Defense Centers, use this sequence to establish communications between the three of them: First, set up remote management between each Defense Center and the Master Defense Center as detailed in Adding and Deleting Defense Centers on page 174, then set up high availability as detailed in Setting Up High Availability on page 158.
- 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 sensors they share. That is, the primary Defense Center must be able to contact the secondary Defense Center at

Keysight, EX-1004
Page 157

the IP address on the secondary Defense Center's own management interface, and vice versa. In addition, either each Defense Center must be able to contact the sensors it manages or the sensors must be able to contact the Defense Center.

Setting Up High Availability

Platform: DC

Component: None

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 the Management Virtual Network on page 1809.

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 1768.

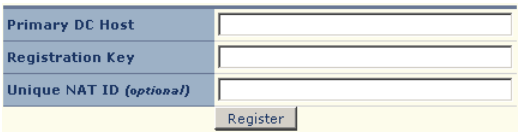
To set up high availability for two Defense Centers:

Access: Admin

1. Log into the Defense Center that you want to designate as the secondary Defense Center.
2. Select **Operations > Configuration > High Availability**.

   The High Availability page appears.
3. Click the secondary Defense Center option.

   The Secondary Defense Center Setup page appears.



Keysight, EX-1004
Page 158

4. 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.

   TIP! 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.

5. Type a one-time-use registration key in the **Registration Key** text box
6. 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 117 for more information.
7. Click **Register**.

   A success message appears, and the Peer Manager page appears, showing the current state of the secondary Defense Center.
8. Using an account with Admin access, log into the Defense Center that you want to designate as the primary.
9. Select **Operations > Configuration > High Availability**.

   The High Availability page appears.
10. Click the primary Defense Center option.

    The Primary Defense Center Setup page appears.



11. 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.

12. Type the same one-time-use registration key in the **Registration Key** text box you used in step 5.
13. If you used a unique NAT ID on the secondary Defense Center, type the same registration ID that you used in step 6 in the **Unique NAT ID** text box.

Keysight, EX-1004
Page 159

14. Click **Register**.

    A success message appears, and the Peer Manager page appears, showing the current state of the primary Defense Center.

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 the High Availability Status on page 160.

If one of the Defense Centers in the HA pair must be reimaged, disable the high availability link first. After you reimage the Defense Center, re-establish the HA 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.

TIP! To add an existing high availability pair of Defense Centers to a Master Defense Center, add the primary Defense Center and the secondary Defense Center is automatically added. For information about adding a Defense Center to a Master Defense Center, see Adding a Master Defense Center on page 174.

Monitoring the High Availability Status

Platform: DC

Component: None

Once 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
- time the Defense Centers last synchronized

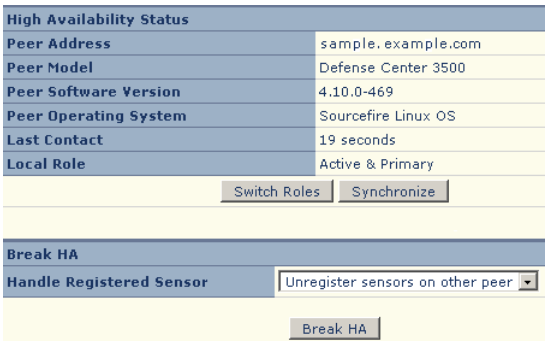
Keysight, EX-1004
Page 160

To check high availability status:

Access: Admin

1. Log into one of the Defense Centers that you linked using high availability.
2. Select **Operations > Configuration > High Availability**.

   The High Availability page appears.



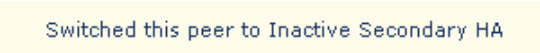
3. Under High Availability Status, you can view the following information about the other Defense Center in the high availability pair:
   - the IP address
   - the model name
   - the software version
   - the operating system
   - the length of time since the last contact between the two Defense Centers
   - the local role (Active & Primary or Inactive & Secondary)
   - the option to switch roles between the two Defense Centers
4. 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 sensor 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 sensor back. This interval ensures that the high availability pair re-synchronizes first. If you do not wait five minutes, it may take more than one synchronization cycle to add the sensor to both Defense Centers.

Keysight, EX-1004
Page 161

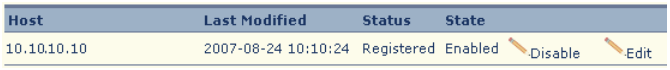
5. Click **Switch Roles** to change the local role from Active & Primary to Inactive & Secondary, or from Inactive & Secondary to Active & Primary.

   The roles are switched between the two peers.

6. Click **Peer Manager** in the toolbar.

   The Peer Manager page appears.

   You can view the following information:

   - the IP address of the other Defense Center in the HA pair
   - the status, registered or unregistered, of the communications link
   - the state, enabled or disabled, of the HA pair

   For information about editing the remote management communications between the two appliances, see Editing the Management Virtual Network on page 1809.

Disabling High Availability and Unregistering Sensors

Platform: DC

Component: None

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 HA pair.
2. Select **Operations > Configuration > High Availability**.

   The High Availability page appears.
3. Select one of the following options from the **Handle Registered Sensors** drop-down list:
   - To control all the managed sensors with the Defense Center where you are accessing this page, select **Unregister sensors on the other peer**.
   - To control all the managed sensors with the other Defense Center, select **Unregister sensors on this peer**.
   - To stop managing the sensors altogether, select **Unregister sensors on both peers**.

Keysight, EX-1004
Page 162

4. Click **Break HA**.

   After you answer the prompt **Do you really want to Break High Availability?** by selecting **OK**, high availability is disabled and any managed sensors 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 158.

Pausing Communication between Paired Defense Centers

Platform: DC

Component: None

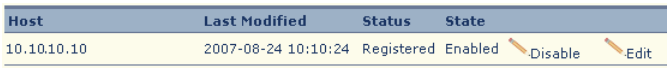
If you want to temporarily disable high availability, you can disable the communications channel between the Defense Centers.

To disable the communications channel for a high availability pair:

Access: Admin

1. Click **Peer Manager**.

   The Peer Manager page appears.



2. Click **Disable** to disable the communications channel between the two Defense Centers.

   For information about editing the remote management communications between the two appliances, see Editing the Management Virtual Network on page 1809.

Restarting Communication between Paired Defense Centers

Platform: DC

Component: None

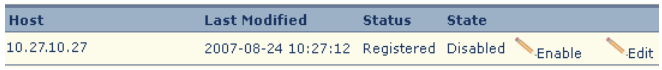
If you temporarily disabled high availability, you can enable the communications channel between the Defense Centers to restart high availability.

To enable the communications channel for a high availability pair:

Access: Admin

1. Click **Peer Manager**.

   The Peer Manager page appears.



Keysight, EX-1004
Page 163

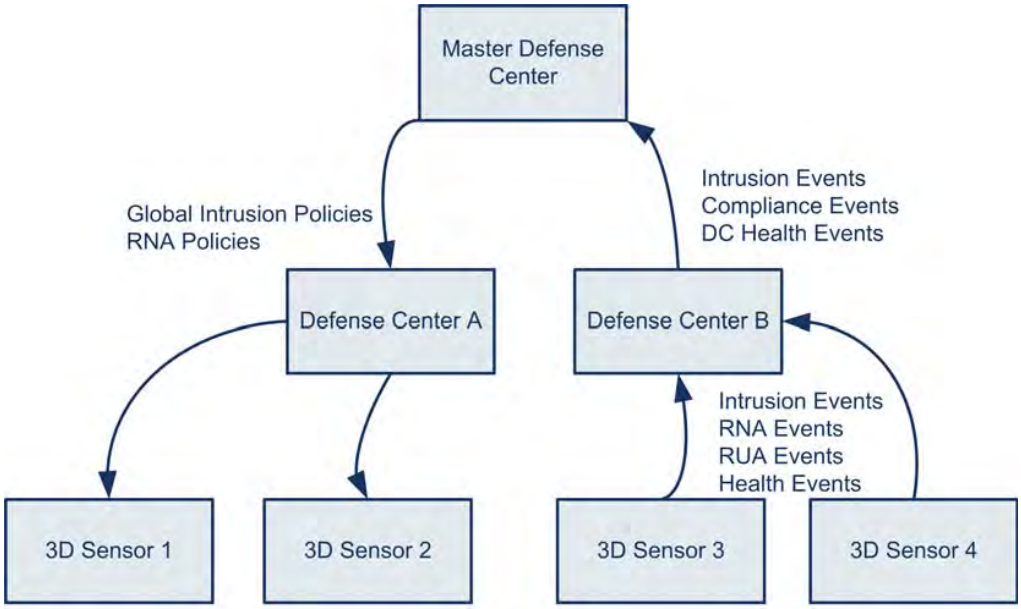
2. Click **Enable** to disable the communications channel between the two Defense Centers.

   For information about editing the remote management communications between the two appliances, see Editing the Management Virtual Network on page 1809.

Keysight, EX-1004
Page 164

Chapter 4
Using the Master Defense Center

The Sourcefire Master Defense Center is a key component in the Sourcefire 3D System. You can use the Master Defense Center to aggregate and analyze intrusion events, compliance events, and white list events from up to ten Defense Centers within your Sourcefire 3D System deployment.



You can use the Master Defense Center to build and dispatch global detection and intrusion policies. When you apply intrusion policies from a Master Defense

Keysight, EX-1004
Page 165

Center, the Sourcefire 3D System checks the SEU on the managing Defense Center. If it finds an older SEU, it updates the managing Defense Center's SEU. The Master Defense Center can also aggregate events related to the health of managed Defense Centers. In this way, you can view the current status of the Defense Centers across your enterprise from a web interface.

IMPORTANT! The Product Compatibility section of the release notes for each version describes which versions of the Defense Center you can manage with a Master Defense Center.

The following sections explain more about using a Master Defense Center in your Sourcefire 3D System deployment:

- Understanding Event Aggregation on page 166 explains which types of events you can send from your Master Defense Centers to your Master Defense Center.
- Understanding Global Policy Management on page 170 explains which policies you can send from your Master Defense Center to 3D Sensors and Defense Centers.
- Adding and Deleting Defense Centers on page 174 explains how to configure a Defense Center to communicate with a Master Defense Center.
- Editing Settings for a Managed Defense Center on page 185 explains how to change some of the settings for a Defense Center from the Master Defense Center's web interface.
- Managing Appliance Groups on page 191 explains how to use appliance groups to aid in managing 3D Sensors and Defense Centers.

Understanding Event Aggregation

Platform: MDC

Component: IPS/RNA

A Master Defense Center can aggregate intrusion events and compliance events (including white list events) from up to ten Defense Centers. You can configure a Defense Center to send intrusion events based on their flag. You can also choose whether to include the packet data collected with the intrusion events.

The settings on the Filter Configuration page determine which events are forwarded from the Defense Center to the Master Defense Center. You can set up a different configuration for each Defense Center, although most deployments will use the same configuration across the enterprise.

Keysight, EX-1004
Page 166

See the following sections for more information:

- Aggregating Intrusion Events on page 167
- Aggregating Compliance Events on page 168
- Limitations on Event Aggregation on page 168

Aggregating Intrusion Events

Platform: MDC

Component: IPS

An intrusion event is generated by IPS when it analyzes network traffic and finds one or more packets that violate the currently applied intrusion policy. Packet decoders, preprocessors, and intrusion rules are all able to generate intrusion events.

When you use the Filter Configuration page to specify which events are forwarded to the Master Defense Center, you can choose one of the following options:

- **Do Not Send** - Intrusion events are not forwarded to the Master Defense Center.
- **Events Only** - The intrusion events specified in the Flags section are forwarded to the Master Defense Center; however, any packets captured for the event are not sent.
- **Events and Packet Data** - The intrusion events specified in the Flags section, along with any related packets, are forwarded to the Master Defense Center.

You can use the Flags section of the Filter Configuration page to forward only the intrusion events that are important to your analysis. For example, you may want to limit the intrusion events on the Master Defense Center to only those with the greatest impact, that is, the red impact flag. If your 3D Sensors are deployed inline and you are using intrusion rules set to Drop and Generate Events, you may also want to send intrusion events with the black down arrow.

You can also use flag settings to reduce the number of intrusion events that are sent to the Master Defense Center in deployments where large numbers of intrusion events are being generated from your 3D Sensors. For example, you can greatly reduce the number of events sent from a Defense Center by excluding events with the blue or gray impact flags.

IMPORTANT! You must deploy both RNA and IPS on your network to generate intrusion events with meaningful impact flags. If you do not deploy 3D Sensors with RNA on your network, then intrusion events are limited to gray impact flags to indicate unknown impact.

Keysight, EX-1004
Page 167

Aggregating Compliance Events

Platform: MDC

Component: RNA

A compliance event is generated by a Defense Center when the conditions for a compliance rule in an active compliance policy are met. The conditions that can trigger a compliance rule include intrusion events, RNA events, flow data, and anomalous network traffic.

When you use the Filter Configuration page to specify which events are forwarded to the Master Defense Center, you can choose to send or not send compliance events. See the following sections for more information:

- Adding a Defense Center on page 178
- Editing the Event Filter Configuration on page 187

Limitations on Event Aggregation

Platform: MDC

Component: IPS/RNA

The Master Defense Center is a powerful tool for analyzing the potential malicious activity across your enterprise's network. However, there are certain limitations that you should take into consideration when you design your Master Defense Center deployment. The Master Defense Center and Defense Center

Keysight, EX-1004
Page 168

Functional Comparison table compares and contrasts Defense Center and Master Defense Center functional areas.

Master Defense Center and Defense Center Functional Comparison

Function	Master Defense Center	Defense Center
License provisions	N/A	provides sensor license, and NetFlow, RNA and RUA feature licenses
3D Sensor configuration	allows you to configure detection engines	allows you to configure detection engines, interface sets, network interfaces.
Analysis and reporting search	allows you to search for intrusion events, compliance events, white list events, SEU import log, audit log, health events.	allows you search for intrusion events, RNA events, hosts, host attributes, services, client applications, flow data, vulnerabilities, compliance events, white list events, white list violations, remediation status, SEU import log, audit log, health events, scan results, users, and RUA events.
Network scans	does not allow Nmap scans.	allows you to run Nmap scans and view results.
Global policies	allows you to build intrusion policies and to distribute them through connected Defense Centers to their managed 3D Sensors throughout the enterprise	policies are normally downloaded only to their managed 3D Sensors
Event consolidation	allows for collection of events from up to ten Defense Centers	events are collected only from managed 3D Sensors

Data Generated by RNA

The Master Defense Center cannot aggregate RNA events or flow data generated by RNA and forwarded to a Defense Center. In addition, the Master Defense Center does not build a network map or host data for the hosts on your network. However, because you can forward compliance events and white list events from your managed Defense Centers to your Master Defense Center, you can gain insight into RNA-detected activity across your enterprise. To take advantage of this, on your Defense Centers you need to build compliance rules and policies that are triggered by the RNA events that interest you and forward the resulting compliance events to the Master Defense Center.

Keysight, EX-1004
Page 169

Event Rate

The event rate limit for the Master Defense Center is the same rate limit on Defense Centers. This means that if your Defense Centers are accepting events from their 3D Sensors up to the rate limit, you must adjust the event filter on the Master Defense Center so that only the most important events are forwarded from the Defense Centers. For example, in cases where the intrusion event rate is high, you might want to adjust the filter to send only intrusion events with red impact flags. You can also limit the amount of data transferred between a Defense Center and its Master Defense Center by sending only intrusion event data, and not sending the packet data.

Intrusion Agents

Intrusion events generated by intrusion agents are not forwarded to the Master Defense Center.

Understanding Global Policy Management

Platform: MDC

Component: IPS/RNA

You can use the Master Defense Center to generate global intrusion policies and coordinate them with potential vulnerabilities detected by RNA policies. Global intrusion policies are beneficial in rapid response scenarios and during enterprise-wide intrusion policy updates. The Master Defense Center sends the policy through a Defense Center to a 3D Sensor's detection engine. Master Defense Center generated policies are not accessible on an intermediate Defense Center, however if a newer SEU resides on the Master Defense Center than on a Defense Center in the path, then the downstream SEU is updated. This ensures that a global intrusion policies utilize the latest SEU.

RNA compares the data it collects and analyzes with its vulnerability database to determine the potential vulnerabilities on the detected host. You can build, apply edit, delete and export RNA on a Master Defense Center. Existing RNA policies are available for viewing so that you can determine:

- RNA policy name and description
- Detection policy settings such as update interval, if banners and HTTP URLs are captured, if client application are being detected, and so on.
- Which networks and ports are monitored by the RNA policy
- If NetFlow is used to generate host information, which networks and NetFlow-enabled devices are monitored by NetFlow.

For information on creating and applying as well as deleting RNA policies, see What is an RNA Detection Policy? on page 889.

Keysight, EX-1004
Page 170

You can also import and export compliance policies and rules, custom service decoders, as well as intrusion, system, and health policies. For information on import and export functions, see Importing and Exporting Objects on page 1997.

Managing Global Intrusion Policies

Platform: MDC

Component: IPS/RNA

Refer to the following sections for information about managing intrusion policies:

- Creating an Intrusion Policy on page 338 explains how to create an intrusion policy.
- Editing an Intrusion Policy on page 341 explains how to modify existing intrusion policies.
- Applying an Intrusion Policy on page 347 explains how to apply a new or updated intrusion policy to the appropriate IPS detection engines.
- Defining IP Addresses and Ports for Your Network on page 398 provides the syntax used to specify IP addresses and port numbers within the variables and rules in your policy.
- Managing Variables on page 371 explains how to create and manage variables that you can use within intrusion policies.
- Managing Rules in an Intrusion Policy on page 403 explains how to enable and disable intrusion rules within an intrusion policy. This section also explains how to configure rules so that they drop malicious packets in an inline deployment.
- Importing SEUs and Rule Files on page 1848 explains how to download and import Security Enhancement Updates (SEUs) that contain new intrusion rules. Note that SEUs can also contain new and updated decoders and preprocessors.

Using RNA Detection Policies on a Master Defense Center

Platform: MDC

Component: RNA

You can create, edit, delete, export, and apply RNA detection policies from a Master Defense Center. Refer to the following, for information on the following RNA detection policy functions:

- Creating RNA Detection Policies on page 900
- Applying an RNA Detection Policy on page 919
- Editing an RNA Detection Policy on page 921
- Deleting an RNA Detection Policy on page 921

Keysight, EX-1004
Page 171

Using Health Policies on a Master Defense Center

Platform: MDC

Component: None

You can edit, delete, and apply default health policies to the Master Defense Center and to connected Defense Centers. For information about health policies see the following:

- Understanding Health Monitoring on page 1891
- Configuring Health Policies on page 1895
- Using the Health Monitor Blacklist on page 1935
- Configuring Health Monitor Alerts on page 1940
- Using the Health Monitor on page 1945
- Using Appliance Health Monitors on page 1947
- Working with Health Events on page 1957

See Health Policies on page 173 to distinguish the health policy modules that are useful on a Master Defense Center or Defense Center from those that are not, and for brief descriptions of those modules that are used.

Using System Policies on a Master Defense Center

Platform: MDC

Component: None

System policies allow you to manage the following functions on your Defense Centers or Master Defense Center:

- access configuration
- authentication profiles (Defense Center only)
- database limits
- DNS cache settings
- the mail relay host and a notification address for database prune messages
- language selection (English or Japanese)
- login banner
- the kinds and amount of RNA data stored in the database (Defense Center only)
- time synchronization settings

See Managing System Policies on page 1732 for information about system policy usage.

Keysight, EX-1004
Page 172

Master Defense Center Policy Management Limitations

Platform: MDC

Component: IPS/RNA/RUA

There are several types of policies including intrusion policies, RNA detection policies, RUA detection policies, and health policies. The Defense Center and Master Defense Center do not handle these policies in the same manner.

Intrusion Policies

You can create, edit, delete, export, and apply intrusion detection and prevention policies from a Master Defense Center. The Sourcefire 3D System bases intrusion policies on SEUs residing on the appliance where the policy is built. When you apply an intrusion policy to a 3D Sensor's detection engines from a Master Defense Center, the Sourcefire 3D System checks for any older SEUs on Defense Center(s) managing those detection engines. If it finds SEUs older than those on the Master Defense Center, they are updated. Therefore, a warning message with a check box appears. After you acknowledge the message by clicking its check box, the **Apply** button activates.

You can apply one or more custom intrusion policies filtered to monitor VLAN or subnetwork traffic on the network monitored by the detection engine where you apply the policy.

TIP! Before applying a filtered policy, you must apply a non-filtered policy to the detection engine from the same Defense Center or Master Defense Center. You cannot apply a non-filtered policy from a Defense Center then add filters to it from a managing Master Defense Center.

RNA Detection Policies

RNA analysis and reporting functions such as using the network map, listing RNA hosts and events, and listing client applications and vulnerabilities are performed on Defense Centers and not on Master Defense Centers. However, if your deployment includes RNA, you can view host profiles from event views by clicking the host profile icon () next to an IP address.

RUA Detection Policies

There are currently no Real-Time User Awareness functions on a Master Defense Center. RUA functions are available only on properly licensed Defense Centers.

Health Policies

The Master Defense Center monitors its health and the health of connected Defense Centers. Master Defense Centers apply health policies only to Master

Keysight, EX-1004
Page 173

Defense Centers and Defense Centers. Default 3D Sensor, Default IPS, Default IPS (3Dx800 only), and Default RNA Health Policies are not used on the Master Defense Center.

Currently, only the generic Default Health Policy is available for editing and application to appliances. For a listing of the health policy modules that apply to Defense Centers, see the Enabled Defense Center Health Modules - Default Health Policy table on page 1899. For a listing of the health policy modules that apply to Master Defense Centers, see the Enabled MDC Health Modules - Default Health Policy table on page 1900. Policies that are not applicable are implicitly disabled when there is an attempt to apply them to a Defense Center or an Master Defense Center. For details about editing appropriate health policies, see Editing Health Policies on page 1931.

System Policies

System policies are applied only to Master Defense Centers and Defense Centers from a Master Defense Center.

Adding and Deleting Defense Centers

Platform: MDC

Component: None

When you manage a Defense Center with your Master Defense Center, you set up a two-way, SSL-encrypted communication channel between the appliances. The Defense Center uses this channel to send events to the Master Defense Center. As the Defense Center receives events from its sensors, it evaluates which events, based on filter configuration, it should send to the Master Defense Center using the same channel.

- Adding a Defense Center on page 178
- Deleting a Defense Center on page 181
- Resetting Management of a Defense Center on page 182

Adding a Master Defense Center

Platform: MDC

Component: None

You can add a Master Defense Center connection to your Defense Center, however before you do, you must make sure that the network settings are configured correctly on both appliances. This is usually completed as part of the

Keysight, EX-1004
Page 174

installation process, but you can see Configuring Network Settings on page 1800 for details.

TIP! To add an existing high availability pair of Defense Centers to a Master Defense Center, add the primary Defense Center and the secondary Defense Center is automatically added.

Three fields are provided for setting up communications between appliances:

- **Management Host** or **Host**- for the hostname or IP address.
- **Registration Key** - registration key
- **Unique NAT ID (optional)** - for a unique alphanumeric ID. See Working in NAT Environments on page 117 for more information.

TIP! Set up the managed appliance first. At a Defense Center, add the remote management then at the managing Master Defense Center, add the Defense Center.

Valid combinations include:

- **Management Host** or **Host** and **Registration Key** used on both appliances
- **Registration Key** and **Unique NAT ID** used on the Defense Center with **Host**, **Registration Key**, and **Unique NAT ID** used on the Master Defense Center
- **Management Host**, **Registration Key**, and **Unique NAT ID** used on the Defense Center with **Registration Key** and **Unique NAT ID** used on the Master Defense Center

IMPORTANT! The **Management Host** or **Host** field (hostname or IP address) must be used on at least one of the appliance.

To add a Master Defense Center, you need to determine which events on the Defense Center you want to forward to the Master Defense Center.

To add a Master Defense Center to a Defense Center:

Access: Admin

1. Log into the web interface of the Defense Center you want to add.
2. Select **Operations > System Settings**.

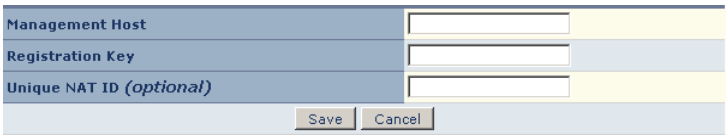
   The Information page appears.
3. Click **Remote Management**.

   The Remote Management page appears.

Keysight, EX-1004
Page 175

4. Click **Add Manager**.

   The Add Remote Management page appears.



5. In the **Management Host** field, type the IP address or the host name of the Master Defense Center that you want to use to manage the Defense Center.

   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 **Management 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. In the **Registration Key** field, type the one-time use registration key that you want to use to set up a communications channel between the Master Defense Center and the Defense Center.

7. Optionally, in the **Unique NAT ID** field, type a **unique** alphanumeric NAT ID that you want to use to identify the Defense Center.

8. Click **Save**.

   After the Defense Center confirms communication with the Master Defense Center, the Pending Registration status appears.



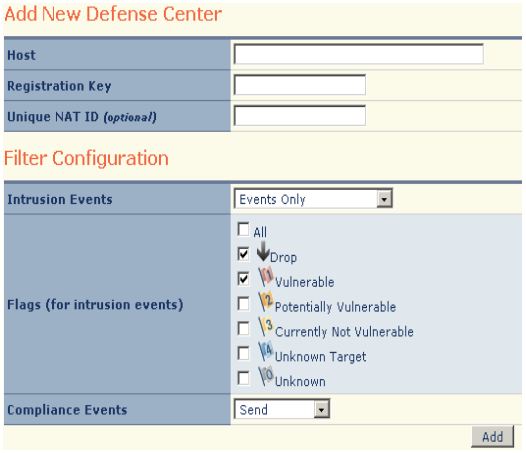
9. Log into the Master Defense Center's web interface using a user account with Admin access, and select **Operations > Appliances**.

   The Defense Centers page appears.

Keysight, EX-1004
Page 176

10. Click **New Defense Center**.

    The New Defense Center page appears.



11. Type the IP address or the hostname of the Defense Center you want to add in the **Host** field.

    TIP! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses. You can leave the **Host** field empty if the host does not have a routable address. In that case, use both the **Registration Key** and the **Unique NAT ID** fields

12. In the **Registration Key** field, type the same one-time use registration key that you used in step 6.

13. If you used an unique NAT ID in step 6, type the same value in the **Unique NAT ID (optional)** field.

14. Under Filter Configuration, identify the types of events you want to forward from the Defense Center to the Master Defense Center.

    Note that if you select intrusion events, you can send events or events and packet data. You can also filter which intrusion events are forwarded based on their impact flag or inline result. If you chose to send compliance events to the Master Defense Center, white list events are also sent. See Editing the Event Filter Configuration on page 187 for more information.

    IMPORTANT! You must select at least one type of flag if you want to send intrusion events.

Keysight, EX-1004
Page 177

15. Click **Add**.

    The Defense Center is added to the Master Defense Center. It can take up to two minutes for the Defense Center to establish communication with the Master Defense Center. You can view the status on the Defense Centers page (**Operations > Appliances**).

16. After communications between the two appliances are established, continue with the next section, Adding a Defense Center.

Adding a Defense Center

Platform: DC/MDC

Component: None

Before you add a Defense Center to a Master Defense Center, you must make sure that the network settings are configured correctly on both appliances. This is usually completed as part of the installation process. For more information see Configuring Network Settings on page 1800.

IMPORTANT! If you registered a Master Defense Center and Defense Center using IPv4 and want to convert them to IPv6, you must delete and re-register the Defense Center.

Three fields are provided for setting up communications between appliances:

- **Management Host** or **Host**- for the hostname or IP address.
- **Registration Key** - one-time use registration key
- **Unique NAT ID (optional)** - for a unique alphanumeric ID. See Working in NAT Environments on page 117 for more information.

TIP! Set up the managed appliance first. At a Defense Center, add the remote management, then at the managing Master Defense Center add the Defense Center.

Keysight, EX-1004
Page 178

Valid combinations include:

- **Management Host** or **Host** and **Registration Key** used on both appliances
- **Registration Key** and **Unique NAT ID** used on the Defense Center with **Host**, **Registration Key**, and **Unique NAT ID** used on the Master Defense Center
- **Management Host**, **Registration Key**, and **Unique NAT ID** used on the Defense Center with **Registration Key** and **Unique NAT ID** used on the Master Defense Center

IMPORTANT! The **Management Host** or **Host** field (hostname or IP address) must be used on at least one of the appliance.

To add a Defense Center, you need to predetermine which events on the Defense Center you want to forward to the Master Defense Center.

To add a Defense Center to a Master Defense Center:

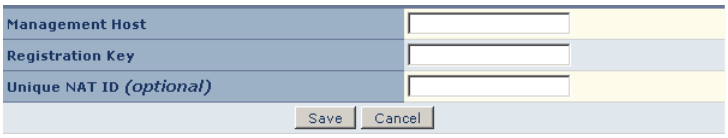
Access: Admin

1. Using a user account with Admin access, log into the web interface of the Defense Center you want to add.
2. Select **Operations > System Settings**.

   The Information page appears.
3. Click **Remote Management**.

   The Remote Management page appears.
4. Click **Add Manager**.

   The Add Remote Management page appears.



5. In the **Management Host** field, type the IP address or the host name of the Master Defense Center that you want to use to manage the Defense Center.

   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 **Management 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. In the **Registration Key** field, type the one-time use registration key that you want to use to set up a communications channel between the Master Defense Center and the Defense Center.

Keysight, EX-1004
Page 179

7. Optionally, in the **Unique NAT ID** field, type a **unique** alphanumeric NAT ID that you want to use to identify the Defense Center.
8. Click **Save**.

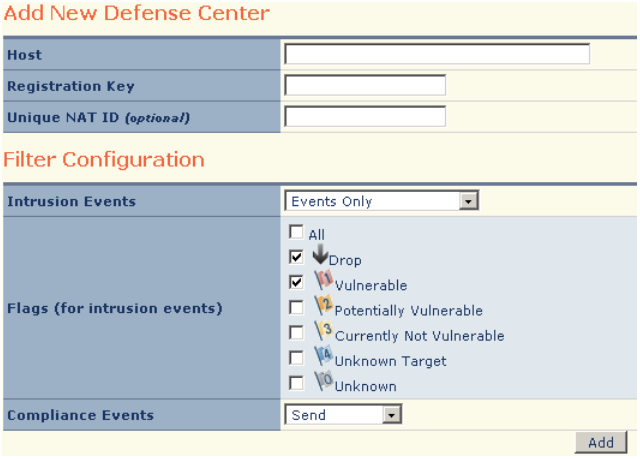
   After the Defense Center confirms communication with the Master Defense Center, the Pending Registration status appears.



9. Log into the Master Defense Center's web interface using a user account with Admin access, and select **Operations > Appliances**.

   The Defense Centers page appears.
10. Click **New Defense Center**.

    The New Defense Center page appears.



11. Type the IP address or the hostname of the Defense Center you want to add in the **Host** field.

WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

12. In the **Registration Key** field, type the same one-time use registration key that you used in step 6.
13. If you used a NAT ID in step 7, type the same value in the **Unique NAT ID (optional)** field.

Keysight, EX-1004
Page 180

14. Under Filter Configuration, identify the types of events you want to forward from the Defense Center to the Master Defense Center.

    Note that if you select intrusion events, you can send events or events and packet data. You can also filter which intrusion events are forwarded based on their impact flag or inline result. If you chose to send compliance events to the Master Defense Center, white list events are also sent. See Editing the Event Filter Configuration on page 187 for more information.

    IMPORTANT! You must select at least one type of flag or inline result if you want to send intrusion events.

15. Click **Add**.

    The Defense Center is added to the Master Defense Center. It can take up to two minutes for the Defense Center to establish communication with the Master Defense Center. You can view the status on the Defense Centers page (**Operations > Appliances**).

Deleting a Defense Center

Platform: DC/MDC

Component: None

If you no longer want to manage a Defense Center, you can delete it from the Master Defense Center. Deleting a Defense Center severs all communication between the Defense Center and the Master Defense Center. To manage the Defense Center again at a later date, you must re-add it to the Master Defense Center. To keep the Defense Center from trying to reconnect to the Master Defense Center, you should also delete the manager on the Defense Center.

To delete a Defense Center from the Master Defense Center:

Access: Admin

1. Log into the Master Defense Center web interface, and select **Operations > Appliances**.

   The Defense Centers page appears.

2. Click **Delete** next to the Defense Center you want to delete.

   Communication between the Master Defense Center and the Defense Center is discontinued and the Defense Center is deleted from the Defense Centers page.

3. Log into the web interface of the Defense Center you want to delete.

4. Select **Operations > System Settings**.

   The Information page appears.

Keysight, EX-1004
Page 181

5. Click **Remote Management**.

   The Remote Management page appears.

6. Click **Delete** next to the Master Defense Center that was managing the Defense Center.

   The manager is removed.

Resetting Management of a Defense Center

Platform: DC/MDC

Component: None

If communications fail between the Master Defense Center and one of your Defense Centers, you can reset management of the Defense Center. If you want to manage a Defense Center with a different Master Defense Center, you must also reset management before adding the Defense Center to the another Master Defense Center. To do this, you must first delete the manager on the Defense Center and delete the Defense Center on the Master Defense Center. You can then re-add the Master Defense Center on the Defense Center and then add the Defense Center to a Master Defense Center.

To reset management from a Master Defense Center:

Access: Admin

1. Log into the web interface of the Master Defense Center where you want to reset communications.
2. Select **Operations > Appliances**.

   The Defense Centers page appears.

3. Click **Delete** next to the Defense Center you want to delete.

   Communication between the Defense Center and the Master Defense Center is discontinued and the Defense Center is deleted from the Defense Centers page.

To delete management on the Defense Center:

Access: Admin

1. Log into the web interface of the Defense Center where you want to reset communications.
2. Select **Operations > System Settings**.

   The Information page appears.

3. Click **Remote Management**.

   The Remote Management page appears.

Keysight, EX-1004
Page 182

4. Click **Delete** next to the Master Defense Center where you want to reset management.

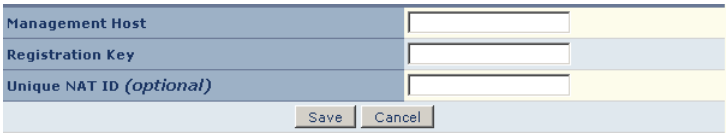
   The manager is removed.

To re-add the Defense Center to the Master Defense Center:

Access: Admin

1. Log into the web interface of the Defense Center where you want to reset communications and click **Add Manager**.

   The Remote Management page appears.



2. In the **Management Host** field, type the IP address or the host name of the Master Defense Center that you want to use to manage the Defense Center.

WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

TIP! You can leave the **Management 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

3. In the **Registration Key** field, type the one-time use registration key that you want to use to set up a communications channel between the Defense Center and the Master Defense Center.
4. Optionally, in the **Unique NAT ID** field, type a **unique** alphanumeric NAT ID that you want to use to identify the Defense Center. See Working in NAT Environments on page 117 for more information.
5. Click **Save**.

   After the Defense Center confirms communication with the Master Defense Center, the Pending Registration status appears.
6. Log into the Master Defense Center's web interface and select **Operations > Appliances**.

   The Defense Centers page appears.
7. Click **New Defense Center**.

   The Add New Defense Center page appears.

Keysight, EX-1004
Page 183

8. Type the IP address or the hostname of the Defense Center you want to add in the **Host** field.

   WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

9. In the **Registration Key** field, type the same one-time use registration key that you used in step 3.
10. If you used an alphanumeric NAT ID in step 4, type the same value in the **Unique NAT ID (optional)** field.
11. To add the Defense Center to a group, select the group from the **Add to Group** list.

    For more information about Defense Center groups, see Managing Appliance Groups on page 191.
12. Click **Add**.

    The Defense Center is added to the Master Defense Center. It can take up to two minutes for the Master Defense Center to verify communication with the Defense Center. You can view the Defense Center's status on the Defense Centers page (**Operations > Appliances**).

Using the Appliances Page

Platform: MDC

Component: None

The Appliances page (**Operations > Appliances**) provides you with a range of information and options that you can use to manage your Defense Centers. The following sections describe the features on the Appliances page.

Sort-by Drop-Down List

Use this drop-down list to sort the Appliances page according to your needs. You can sort by:

- Group, which sorts by Appliance group (see Managing Appliance Groups on page 191)

  TIP! High availability Defense Center pairs are automatically listed as an appliance group. An HA pair is listed as a group named with the name of the active Defense Center.

Keysight, EX-1004
Page 184

- Manager, which sorts by the Defense Center then the 3D Sensor connected to it.
- Model, which sorts by appliance model number, that is, the Defense Center 1000 and the Defense Center 3000, 3D Sensor 2100, and so on.

Status Icons



The status icons indicate the state of a Defense Center. The green check mark icon indicates that the Master Defense Center and the Defense Center are communicating properly. The red exclamation point icon indicates that the Master Defense Center has not received communications from the Defense Center in the last three minutes. If you hover your cursor over the icon, a pop-up window indicates the amount of time (in hours, minutes, and seconds) since the last contact. If the Master Defense Center has not received a communication from a Defense Center within the last two minutes, it sends a two-byte heartbeat packet to establish contact and ensure that the communications channel is still running. If your network is constrained in bandwidth, you can contact technical support to change the default time interval.

Edit and Delete Icons



Click the Edit icon next to a sensor if you want to change the Defense Center's current system settings. The system settings include the filter configuration for the Defense Center, the remote management configuration, the health blacklist settings, and the high availability settings. See Editing Settings for a Managed Defense Center on page 185 for more information.

Click the Delete icon next to a Defense Center if you no longer want to manage the Defense Center with the Master Defense Center. See Deleting a Defense Center on page 181 for more information.

Editing Settings for a Managed Defense Center

Platform: MDC

Component: None

After you configure management of a Defense Center by a Master Defense Center, you can use the Master Defense Center web interface to view and edit the configuration of the Defense Center. See the following sections for more information.

- Viewing the Defense Center Information Page on page 186
- Editing the Event Filter Configuration on page 187

Keysight, EX-1004
Page 185

- Editing or Disabling Remote Management Communications on page 189
- Managing the Health Blacklist on page 189
- Managing High Availability Defense Centers on page 190

Viewing the Defense Center Information Page

Platform: MDC

Component: None

To access the system settings information page for a managed Defense Center, select **Appliances** from the **Operations** menu, then click Edit next to the Defense Center. The Information page for a managed Defense Center includes the fields described in the Defense Center Information table.

Defense Center Information

Field	Description
Name	The assigned name for the Defense Center. Note that this is the name of the Defense Center in the Master Defense Center web interface, not the hostname.
Product Model	The model name for the managed Defense Center.
Serial Number	The serial number for the chassis of the Defense Center.
Software Version	The version of the software currently installed on the managed Defense Center.
Operating System	The operating system currently running on the managed Defense Center.
Operating System Version	The version of the operating system currently running on the managed Defense Center.
VDB Version	The Vulnerability Database version on the managed Defense Center.
IP Address	The IP address of the managed Defense Center.
Status	An icon showing the current status of the managed Defense Center. If you hover your cursor over the icon, a pop-up message indicates how long it has been (in hours, minutes, and seconds) since the Defense Center communicated with the Master Defense Center. You can click **Refresh** to update the Status icon and its accompanying pop-up message.

Keysight, EX-1004
Page 186

Defense Center Information (Continued)

Field	Description
Model Number	The model number for the Defense Center. This number can be important for troubleshooting.
Current Group	The group that the Defense Center belongs to, if any.

To edit a managed Defense Center's settings:

Access: Admin

1. Change the Defense Center's attributes as needed.

   You can edit the following:

   - the name of the Defense Center
   - the group in which the Defense Center resides

   WARNING! The name must be made up of a combination of alphanumeric characters and should not be made up of numeric characters only.

2. Click **Save**.

   The updated Defense Center attributes are saved.

Editing the Event Filter Configuration

Platform: MDC

Component: None

The settings on the Filter Configuration page control which events are sent from the Defense Center to the Master Defense Center that manages it. Your options are to send intrusion events, intrusion events and related packet data, and compliance events.

If you want to send intrusion events (with or without packet data), you can also specify which intrusion events are sent based on their impact flag or inline result. See the Impact Flags table on page 307 for an explanation of what each impact flag means. Note that you must deploy both RNA and IPS as part of your Sourcefire 3D System deployment to generate meaningful impact flags.

TIP! If you set up the 3D Sensor so it does not send packet data to the intermediate Defense Center, then packet data is not forwarded to the Master Defense Center.

Keysight, EX-1004
Page 187

To modify the event filter configuration:

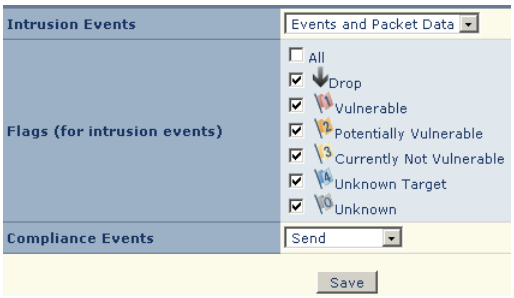
Access: Admin

1. On the Master Defense Center's web interface, select **Operations > Appliances**.

   The Appliances page appears.

2. Next to the Defense Center whose filter configuration you want to change, click **Edit**.

   The Filter Configuration page appears.



3. In the Intrusion Events area, use the drop-down list to indicate whether you want to forward intrusion events to the Master Defense Center. The options are **Do Not Send**, **Events Only**, and **Events and Packet Data**.

4. If you indicated that you want to send intrusion events, then you must specify which events you want to send based on their impact flag or inline result. The Flags options are:

   - All
   - Black down arrow (or Drop)
   - Red (or Vulnerable)
   - Orange (or Potentially Vulnerable)
   - Yellow (or Currently Not Vulnerable)
   - Blue (or Unknown Target)
   - Gray (or Unknown)

   TIP! If you select All, then all the options are immediately selected. If you want to send intrusion events to the Master Defense Center, then you must select at least one impact flag or the drop inline result option.

5. In the Compliance Events area, use the drop-down list to indicate whether you want to forward compliance events to the Master Defense Center. The options are **Do Not Send** and **Send**.

Keysight, EX-1004
Page 188

6. Click **Save**.

   Your settings are saved and the Defense Center begins forwarding the events you specified to the Master Defense Center that manages it.

Editing or Disabling Remote Management Communications

Platform: MDC

Component: None

You can manage communications between a managed Defense Center and its Master Defense Center using the Master Defense Center's web interface. For example, if a Defense Center is no longer responding, you can temporarily disable communications between the Defense Center and its Master Defense Center.

IMPORTANT! Master Defense Centers do not currently use a Management Virtual Network. You cannot edit the **Management Virtual Network** field of a Master Defense Center. The field is filled with 0.0.0.0/24 to indicate that the Management Virtual Network is disabled on a Master Defense Center.

To disable communications between the Defense Center and the Master Defense Center:

Access: Admin

- Click **Disable** next to the name of the Defense Center. Communications between the two appliances are interrupted.

  To enable communications between the two appliances again, click **Enable**.

For more information about editing the Management Virtual Network, see Editing the Management Virtual Network on page 1809.

Managing the Health Blacklist

Platform: MDC

Component: None

You can blacklist individual health policy modules on Defense Centers. You may want to do this to prevent events from the module from changing the status for the appliance to warning or critical.

For information on using the blacklisting function, see Using the Health Monitor Blacklist on page 1935.

Keysight, EX-1004
Page 189

Managing High Availability Defense Centers

Platform: MDC

Component: None

You can configure, monitor, disable, pause and restart Defense Center High Availability from a Defense Center. See the following sections for more information:

- Using Redundant Defense Centers on page 117
- Setting Up High Availability on page 158
- Monitoring the High Availability Status on page 160
- Disabling High Availability and Unregistering Sensors on page 162
- Pausing Communication between Paired Defense Centers on page 163
- Restarting Communication between Paired Defense Centers on page 163

If High Availability is configured, you can activate Defense Center High Availability from a Master Defense Center.

IMPORTANT! When using Intrusion Agents registered to Defense Centers configured for high availability and managed by a Master Defense Center, register all Intrusion Agents to the primary Defense Center.

To activate a redundant Defense Center:

Access: Admin

1. Select **Operations > Appliances**.

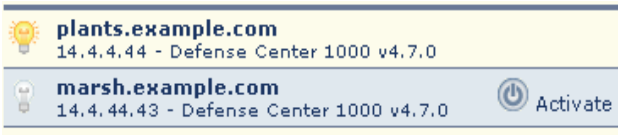
   The Appliances page appears.

2. Click **Edit** next to the appropriate Defense Center.

   The System Settings page for that Defense Center appears.

3. Click **High Availability**.

   The high availability page appears with the paired Defense Centers.



TIP! A light bulb icon shows which of the high availability paired Defense Centers is currently active.

4. Click **Activate** to activate the redundant Defense Center.

   The redundant Defense Center is activated.

Keysight, EX-1004
Page 190

Managing Appliance Groups

Platform: MDC

Component: None

The Master Defense Center allows you to group appliances so that you can easily search for events based on whether they were forwarded by one of a specific group of appliances.

TIP! High availability Defense Center pairs are automatically listed as an appliance group. An HA pair is listed as a group with the name of the active Defense Center.

See the following sections for more information:

- Creating Appliance Groups on page 191 explains how to create a Defense Center group on the Master Defense Center.
- Editing Appliance Groups on page 192 explains how to modify the list of Defense Centers in a Defense Center group.
- Deleting Appliance Groups on page 193 explains how to delete a Defense Center group.

Creating Appliance Groups

Platform: MDC

Component: None

Grouping managed appliances allows you to use the group name as a search criterion when you search for specific compliance or intrusion events.

To create an appliance group and add appliances to it:

Access: Admin

1. On the Master Defense Center, select **Operations > Appliances**.

   The Appliances page appears.
2. Click **Create New Appliance Group**.

   The Create Appliance Group page appears.
3. In the **Group Name** field, type the name of the group you want to create.
4. Click **Save**.

   The group is added.
5. To add appliances to the group, return to the Appliances page (**Operations > Appliances**) and click **Edit** next to the name of the group.

   The Appliance Group Edit page appears.

Keysight, EX-1004
Page 191

6. Select the IP addresses or hostnames of the appliances you want to add from the **Available Appliances** list and click the arrow to move them into the group.
7. Click **Save**.

   The appliances are added to the group and the Appliances page appears again.

Editing Appliance Groups

Platform: MDC

Component: None

You can change the set of appliances that reside in any appliance group.

TIP! You must remove an appliance from its current group before you can add it to a new group.

IMPORTANT! Moving an appliance to a new group does not change any of its policies or configurations.

To edit an appliance group:

Access: Admin

1. On the Master Defense Center, select **Operations > Appliances.**

   The Appliances page appears.
2. Click **Edit** next to the Appliance group you want to edit.

   The Appliance Group Edit page appears.
3. Select the appliance you want to move and click the arrow to add or remove it from the group.
   - To add an appliance to the group, select it from the **Available Appliances** list and click the arrow pointing toward the group you are editing.
   - To remove an appliance from a group, select it from the list in the group you are editing and click the arrow pointing to the **Available Appliances** list.
4. Click **Save**.

Keysight, EX-1004
Page 192

Deleting Appliance Groups

Platform: MDC

Component: None

If you delete a group that contains appliances, the appliances are moved to Ungrouped on the Appliances page. They are not deleted from the Master Defense Center.

To delete an appliance group:

Access: Admin

1. Select **Operations > Appliances**.

   The Appliances page appears.

2. Click **Delete** next to the group you want to delete.

   The appliances group is removed from the Master Defense Center.

Editing Master Defense Center System Settings

Platform: MDC

Component: None

With a few exceptions, the Master Defense Center system settings are the same as those of a Defense Center. See the following sections for information on each of the listed system settings:

IMPORTANT! NetFlow-enabled devices cannot currently be added to a Master Defense Center.

- Listing Master Defense Center Information on page 193
- Viewing a Master Defense Center License on page 194
- Configuring Network Settings on page 1800
- Shutting Down and Restarting the System on page 195
- Setting System Time on page 196
- Blacklisting Health Policies on page 196

Listing Master Defense Center Information

Platform: MDC

Component: None

For details on information listed under the Master Defense Center system settings, see Defense Center Information on page 186.

Keysight, EX-1004
Page 193

To edit a Master Defense Center's settings:

Access: Admin

1. Change the name of the Master Defense Center attributes as needed.

   WARNING! The name must be made up of a combination of alphanumeric characters and should not be made up of numeric characters only.

2. Click **Save**.

   The updated Master Defense Center attributes are saved.

Viewing a Master Defense Center License

Platform: MDC

Component: None

Unlike a Defense Center, a Master Defense Center cannot manage the licenses of Defense Centers or 3D Sensors.

To view information about the Master Defense Center license:

Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **License**.

   The License page appears.

Configuring Network Settings

Platform: MDC

Component: None

The network settings are identical to those of the Defense Center. For information on configuring the Master Defense Center network settings, see Configuring Network Settings on page 1800.

Keysight, EX-1004
Page 194

Shutting Down and Restarting the System

Platform: MDC

Component: None

You have several options for controlling the processes on your Master Defense Center. You can:

- shut down the appliance
- reboot the appliance
- restart the appliance

To shut down or restart your appliance:

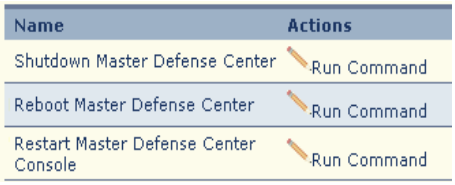
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **Process**.

   The Appliance Process page appears.



3. Specify the command you want to perform:
   - If you want to shut down the Master Defense Center, click **Run Command** next to **Shutdown Master Defense Center**.
   - If you want to reboot the system, click **Run Command** next to **Reboot Master Defense Center**.
   - If you want to restart the Defense Center, click **Run Command** next to **Restart Master Defense Center Console**. Note that restarting the Defense Center may cause deleted hosts to reappear.

Configuring Remote Management Networking

Platform: MDC

Component: None

Master Defense Centers do not currently use a Management Virtual Network. You cannot edit the **Management Virtual Network** field if the Defense Center is in the Master Defense Center operational mode. The field is filled with the address range 0.0.0.0/24 to disable the Management Virtual Network.

Keysight, EX-1004
Page 195

Setting System Time

Platform: MDC

Component: None

The system time is set and synchronized in accordance with the system policy. On the Time Synchronization page you can choose to serve time from the Master Defense Center by selecting **Enabled** in the **Serve Time via NTP** field.

TIP! Because Master Defense Centers do not currently use Management Virtual Networks, their real IP network is used to serve time.

To specify how the Master Defense Center clock is set:

Access: Admin

▶ You have two options:

- To set the time manually, select **Manually in the System Settings**.
- To receive time through NTP from a different server, select **Via NTP Server from** and, in the text box, type the IP address of the NTP server or, if DNS is enabled, type the fully qualified host and domain name.

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, you should configure your DHCP server to set the same NTP server.

For more information about setting system time, see Synchronizing Time on page 1768.

Blacklisting Health Policies

Platform: MDC

Component: None

You can blacklist health policy modules when required. The Master Defense Center supports the following health policy modules:

- Appliance Heartbeat
- CPU Usage
- Data Correlator Process
- Defense Center Status
- Disk Usage
- eStreamer Process

Keysight, EX-1004
Page 196

- Event Stream Status
- Memory Usage

For more information on blacklisting a health policy, see Blacklisting a Health Policy Module on page 1938.

Keysight, EX-1004
Page 197

Chapter 5
Using Detection Engines and Interface Sets

To give you increased flexibility in your deployment choices, the Sourcefire 3D System provides a feature called the *detection engine*. You can think of a detection engine as a collection of one or more sensing interfaces (called an *interface set*) on a 3D Sensor plus a portion of the sensor's computing resources (called a *detection resource*).

3D Sensors support three types of detection engines:

- IPS
- RNA
- RUA

TIP! You cannot use the RUA feature on Crossbeam-based software sensors. In addition, you cannot use RUA or RNA on 3D9800 sensors. However, you can combine the data from those sensors with RUA or RNA on a Defense Center.

The number of detection engines per sensor is limited by the number of detection resources that are available. Most 3D Sensor models have at least three detection resources available and can support at least three detection engines: one for IPS, one for RNA, and the third for RUA. See the Detection Resources by Model table on page 203 for more information.

Keysight, EX-1004
Page 198

The following sections describe the detection engines and interface set features and how you can use them in your Sourcefire 3D System deployment:

- Understanding Detection Engines on page 199 explains detection engines in more detail, including some of the limitations based on the sensor model. This section also describes how default detection engines are configured.
- Managing Detection Engines on page 206 explains how to create, edit, and delete detection engines.
- Using Detection Engine Groups on page 211 explains how to create and use detection engine groups.
- Using Variables within Detection Engines on page 213 explains how to use detection-engine-specific variable values to tailor your detection capabilities to more closely match your infrastructure.
- Using Interface Sets on page 222 describes how to create interface sets and how to use them with detection engines.
- Using Interface Set Groups on page 244 describes how to create and use interface sets groups.
- Inline Fail Open Interface Set Commands on page 246 explains how to force an interface set in and out of bypass mode when using an inline fiber fail open interface set.
- Using Clustered 3D Sensors on page 250 explains how to use detection engines and interface sets in a clustered sensor pairing.

Understanding Detection Engines

Platform: DC/MDC/3D Sensor

Component: None

A detection engine is the mechanism on a 3D Sensor that is responsible for analyzing the traffic on the network segment where the sensor is connected.

A detection engine consists of one of more sensing interfaces to pass traffic and a portion of the sensor's computing capacity, known as a detection resource, to power the traffic inspection. You apply a detection policy to a detection engine to specify how the detection engine should analyze network traffic.

When you view a list of detection engine, specifications for each detection engine are displayed.

Name

Lists the user-specified name of the detection engine. See Creating a Detection Engine on page 206.

DE Type

Identifies the type of detection engine: IPS, RNA, or RUA.

Keysight, EX-1004
Page 199

Depending on which components are licensed on the sensor, 3D Sensors can support three types of detection engines: IPS, RNA, and RUA.

A detection engine has two main components:

- an interface set, which can include one or more sensing interfaces
- a detection resource, which is a portion of the sensor's computing resources

For information about detection engines and detection resources, see Understanding Detection Resources and 3D Sensor Models on page 202.

Resources

Specifes the number of detection resources applied to the detection engine.

Interface Set

Lists the user-specified name of the interface set.

PEP Policy

Indicates whether PEP policy is applied to the detection engine; available on 3D9900 and later sensor models. See Using PEP to Manage Traffic on page 1510.

Set Type

Specifies the type of connection used in the interface set: failopen, inline, passive

An interface set refers to a grouping of one or more sensing interfaces on a sensor, although a sensing interface can belong to only one interface set at a time. The Sourcefire 3D System supports three types of interface sets, but the interface options available to you depend on the type of sensor and the capabilities of its sensing interfaces. The three interface types are passive, inline, and inline with fail open.

Passive: Use a passive interface set if you deployed the sensor out of band from the flow of network traffic.

Inline: Use an inline interface set if you deployed the sensor inline on your network and the sensing interfaces do **not** support automatic fail-open capabilities. Note that you can use any two of the non-fail-open interfaces on the sensor's network interface cards as part of an inline interface set. The exception is on 3D9900s, where pairs are pre-determined.

Inline with Fail-Open: Use an inline with fail open interface set if you deployed the sensor inline on your network and the sensing interfaces **do** support automatic fail-open capabilities. Note that you must use paired fail-open interfaces on the sensor's network interface cards for an inline with fail open interface set.

You can use RNA or RUA to monitor the traffic that passes through any of the three types of interface sets.

Keysight, EX-1004
Page 200

On a 3D3800 or 3D5800 sensor, if you plan to use RNA to monitor either an inline or inline with fail open interface set, you **must** either configure an IPS detection engine that uses that interface set, as well as apply an intrusion policy to that detection engine, or configure the interface set in tap mode. Otherwise, the RNA detection engine monitoring that interface set will not see any traffic. If you are monitoring the same inline interface set with both IPS and RNA or RUA, and the IPS detection engine fails for any reason, the RNA or RUA detection engine monitoring that interface set will not see any traffic until the IPS detection engine restarts. Neither RNA nor RUA are supported on the 3D9800 sensor.

See Using Interface Sets on page 222 for more information about creating and editing interface sets.

Policy

Identifies the intrusion policy or policies applied to an IPS detection engine, or the RNA detection policy applied to an RNA detection engine. See Applying an Intrusion Policy on page 347 and Creating RNA Detection Policies on page 900.

3D Sensors have different capabilities and limitations depending on whether you licensed IPS, RUA, or RNA. You can determine what the name and state of IPS and RNA policies from the following information in the policy column:

If you change an IPS and RNA policy and have not applied it to the detection engine since the change, then the icon has an exclamation point and the name is italicized.

You can click the name of an IPS policy to see details about the running policy. For more information see Viewing an Intrusion Policy Report on page 350.

If there is a network or VLAN filter applied to the IPS policy, you can click **More** or the down icon (▼) and view the type (Net for network or VLAN for virtual LAN) filter. Place your cursor over the name to display the network or VLAN range of the filter. If you want to remove the currently applied filter from the IPS policy, click the delete icon (✖) next to the filter name.

If you want to remove the currently applied IPS policy from the detection engine, click the delete icon (✖) next to the intrusion policy name. The delete icon only appears next to a non-filtered policy when there are no VLAN- or network-filtered policies applied.

Initially, the Detection Engines page does not indicate that a filtered or non-filtered intrusion policy is deleted. Select **Monitor > Task Status** to track the progress of the deletion process, which takes approximately 30 seconds.

If you want to reapply all intrusion policies, including all VLAN- and network-filtered intrusion polices, that were originally applied from the appliance, click **Reapply Intrusion Policies**.

Sensor

Identifies the sensor where the detection engine is located.

Keysight, EX-1004
Page 201

The sensor column provides the name of the sensor where the policy is applied.

If you want to edit or delete a detection engine, click **Edit** or **Delete** next to its sensor name. See Editing a Detection Engine on page 208 and Deleting a Detection Engine on page 211 for more information.

If you want to list, add, edit, reset, or delete variables associated with a detection engine's IPS or RNA policy, click **Variables**. See Using Variables within Detection Engines on page 213 for more information.

If you want to reapply a non-filtered intrusion policy and simultaneously reapply any filtered VLAN- or network-specific policy to the detection engine, click **Reapply All**, then **OK** to confirm.

For more information see Understanding Detection Resources and 3D Sensor Models on page 202

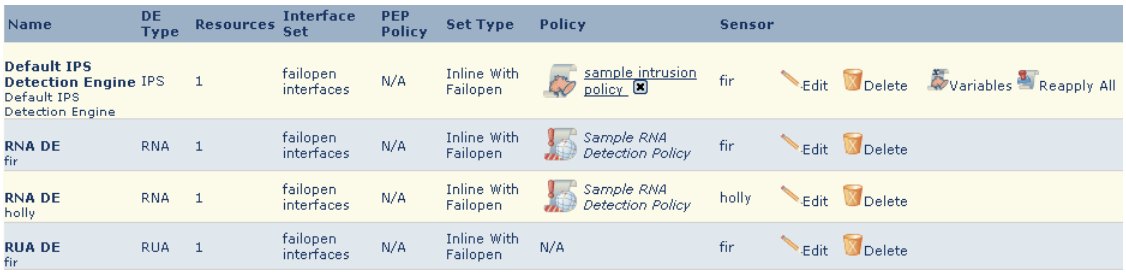
When you configure a new sensor, it has a predefined detection engine that you can choose to modify to meet your needs. See Understanding Default Detection Engines for more information.

To view the list of available detection engines:

Access: Admin

- Select **Operations > Configuration > Detection Engines**.

  The Available Detection Engines page appears. The graphic below shows the Defense Center version of the page.



Understanding Detection Resources and 3D Sensor Models

Platform: DC/Standalone Sensor

Component: None

3D Sensors with IPS can use multiple detection resources per detection engine, which allows you to use more computing resources when network traffic is high. For example, if you plan to use the 3D3500 sensor in inline mode, you could assign two detection resources to your detection engine to allow processing of more events per second.

Note that on an 8000 Series sensor, detection resources are allocated automatically using an intelligent resource algorithm which takes into

Keysight, EX-1004
Page 202

consideration such factors as aggregate link bandwidth of all interfaces in the interface set, type of detection engine, and type of interface set. Any change in configuration (such as creating a new detection engine, modifying the type of interface set, or changing the makeup of an interface set or the binding between an interface set and a detection engine) automatically causes the algorithm to reallocate resources accordingly for optimal performance. Optionally, you can manually allocate your detection resources; call Sourcefire Support for assistance.

As a best practice, use one detection resource per application per core on your appliance. Different sensor models have different numbers of detection resources available as shown in the Detection Resources by Model table.

- The **Optimal** column indicates the per-sensor total number of detection resources you should use if you want to maximize the performance of the sensor. It also indicates the maximum number of detection resources you can assign a single detection engine.
- The **Maximum** column indicates the total number of detection resources available on the sensor.
- The **Combination Restrictions** column indicates the permitted combinations of detection resources that you can allocate to detection engines on the same sensor; 3D Sensors can run combinations of IPS, RNA and RUA.

Detection Resources by Model

Model	Optimal per Sensor	Maximum per Sensor	Combination Restrictions
3D500	1	2	Maximum of one IPS and either one RNA or one RUA
3D1000	1	2	Maximum of two; can be any type
3D2000	1	2	Maximum of two; can be any type
3D2100	2	3	No restrictions
3D2500	2	4	No restrictions
3D3000	2	4	No restrictions
3D3500	2	6	No restrictions
3D3800	2	2	No restrictions
3D4500	4	8	No restrictions
3D5800	6	6	No restrictions

Keysight, EX-1004
Page 203

Detection Resources by Model (Continued)

Model	Optimal per Sensor	Maximum per Sensor	Combination Restrictions
3D6500	8	12	No restrictions
3D9800	12	12	Cannot run RUA or RNA detection engines
3D9900	7	12	No restrictions
Virtual 3D Sensor	3	3	No restrictions
Crossbeam-based software sensors	Refer to Crossbeam-Based Software Sensor Considerations on page 205		
8000 Series	Automatic; see Detection Engines and Resources on 8000 Series Sensors on page 204		

Detection Engines and Resources on 8000 Series Sensors

On an 8000 Series 3D Sensor, detection resources are allocated automatically using an intelligent resource algorithm which takes into consideration such factors as aggregate link bandwidth of all interfaces in the interface set, type of detection engine, and type of interface set.

Any change in configuration (such as creating a new detection engine, modifying the type of interface set, or making a change in an interface set or the binding between an interface set and a detection engine) automatically causes the algorithm to reallocate resources accordingly for optimal performance. Optionally, you can manually allocate your detection resources; call Sourcefire Support for assistance.

The maximum number of detection engines per sensor is two times the number of CPUs (for example, if a sensor has 24 CPUs, you can create up to 48 detection engines).

General Recommendations with Two or More Detection Resources

For improved 3D Sensor performance on sensors with optimal detection resources of two or greater, you can reduce latency by distributing your network traffic across all available interfaces on the sensor, then distribute the detection engines and detection resources across all operative interfaces on the sensor.

Keysight, EX-1004
Page 204

Note that no more than eight detection resources can be assigned to any single interface set, regardless of the number of detection engines assigned to that interface set.

Crossbeam-Based Software Sensor Considerations

Depending upon the capabilities of your X-Series and the products you are licensed to use, you have several deployment options for 3D Sensor Software. Consider how your network is configured and how you want to deploy the Sourcefire 3D System within it. As with other 3D Sensors, the maximum number of detection engines that you can create is equal to the number of available detection resources. The number of detection resource depends on the Crossbeam System hardware.

Refer to the *Sourcefire 3D Sensor Software for X-Series Installation Guide* for information on deployment scenarios, current Crossbeam System hardware and software support, and detection resources available on Crossbeam System hardware.

Understanding Default Detection Engines

Platform: DC/3D Sensor

Component: None

When you install a new 3D Sensor, you can use initial interface sets and default detection engines to quickly begin evaluating network traffic. After initial installation, you can modify interface sets and detection engines.

Initial Interface Sets

- **Inline with Fail-Open**, the default that builds paired fail-open interface sets on all 3D Sensor interfaces, less the management interface.
- **Passive** that builds a single passive interface set for all 3D Sensor interfaces, less the management interface.

Choose from these initial interface sets based on how you deployed the sensor. Select Inline with Fail-Open Mode if you cabled the sensing interfaces inline on your network as an IPS. Depending on the 3D Sensor, typically you pair adjacent interfaces; for example, a 3D2000 Sensor uses eth1 and eth2 as one inline fail-open interface set and it uses eth3 and eth4 as another inline fail-open interface set.

Select Passive Mode if the sensing interfaces are not cabled inline.

Default Detection Engines

Default detection engines are configured with the optimal (rather than maximum) number of detection resources as described in the Detection Resources by Model table on page 203. With this configuration, you can connect any of the

Keysight, EX-1004
Page 205

non-management interfaces to your network and apply the appropriate policy to the detection engine and begin analyzing your network.

Managing Detection Engines

Platform: DC/3D Sensor

Component: None

See Understanding Detection Engines on page 199 and Using Interface Sets on page 222 for more information about the capabilities of detection engines and the interface sets they depend on. The following sections explain how to create, edit, and delete detection engines.

- Creating a Detection Engine on page 206
- Editing a Detection Engine on page 208
- Deleting a Detection Engine on page 211

Creating a Detection Engine

Platform: DC/3D Sensor

Component: None

You can create a detection engine if you have an available interface set and at least one available detection resource. You can use interface sets that include multiple inline interface pairs, when they are available on your 3D Sensor.

To create a detection engine:

Access: Admin

1. Select **Operations > Configuration > Detection Engines**.

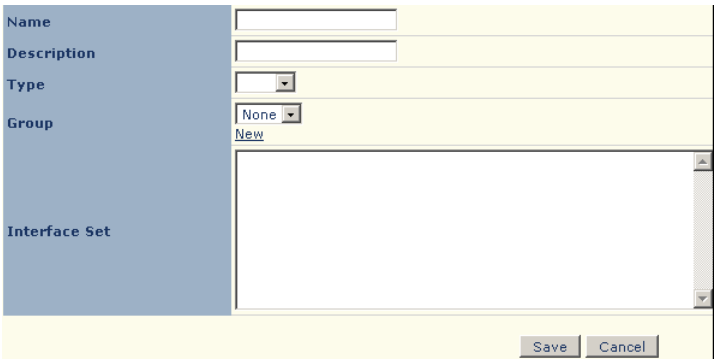
   The Detection Engines page appears. The graphic below shows the Defense Center version of the page.

Name	DE Type	Resources	Interface Set	PEP Policy	Set Type	Policy	Sensor	
Default IPS Detection Engine Default IPS Detection Engine	IPS	1	failopen interfaces	N/A	Inline With Failopen	sample intrusion policy	fir	Edit Delete Variables Reapply All
RNA DE fir	RNA	1	failopen interfaces	N/A	Inline With Failopen	*Sample RNA Detection Policy*	fir	Edit Delete
RNA DE holly	RNA	1	failopen interfaces	N/A	Inline With Failopen	*Sample RNA Detection Policy*	holly	Edit Delete
RUA DE fir	RUA	1	failopen interfaces	N/A	Inline With Failopen	N/A	fir	Edit Delete

Keysight, EX-1004
Page 206

2. Click **Create Detection Engine**.

   The Create Detection Engine page appears.



3. In the **Name** and **Description** fields, enter a name and description for the new detection engine.

   You can use alphanumeric characters, punctuation, and spaces.

4. Select the type of detection engine that you want to create from the **Type** drop-down list, **IPS**, **RNA**, or **RUA**.

5. Optionally, add the detection engine to an existing detection engine group.

   See Using Detection Engine Groups on page 211 for information on creating and modifying detection engine groups.

6. Select the interface set that you want to assign to this detection engine. Interface sets are displayed according to the detection engine type selected.

   See Using Interface Sets on page 222 for information about creating and modifying interface sets.

7. Select the number of detection resources for this detection engine.

   Note that with an 8000 Series 3D Sensor, you do not assign detection resources. An intelligent resource algorithm automatically allocates detection resources for optimal use.

IMPORTANT! On the 3D500, you can only use one of the two detection resources for IPS. The second detection resource is available only if you want to create a second detection engine for RNA or RUA. See the Detection Resources by Model table on page 203 for more information.

Keysight, EX-1004
Page 207

8. Optionally, if you are creating an IPS detection engine and if you are using a 3D Sensor other than a 3D500, 3D1000, or 3D3800, select **Inspect Traffic During Policy Apply**.

   TIP! This option may degrade performance when you apply a policy and may result in longer policy-apply periods. However, if this option is employed, the detection engine does not restart and interrupt traffic inspection when the policy is applied.

9. Click **Save**.

   The detection engine is created.

Editing a Detection Engine

Platform: DC/3D Sensor

Component: None

In some circumstances, editing an interface set or detection engine can cause the detection engines on the sensor to restart, which can cause a short pause in processing.

IMPORTANT! For most 3D Sensors with inline interface sets, a software bridge is automatically set up to transport packets when the sensor restarts. Although some packets are transmitted without inspection during this time, no packets are lost.

The following sections describe some of the cases where a detection engines is affected by changes to the detection engines and interface sets:

3Dx800 Sensors

- If you change the number of network interfaces, the interface set type, or the setting for tap mode or transparent mode for an interface set, all the detection engines using that interface set are restarted.
- If you change the number of detection resources, which interface set is used, or the detection engine type, only that detection engine is restarted (although other CPUs may be restarted to rebalance the processing load).

  IMPORTANT! If you have an 3Dx800 health policy applied to a 3D9800 sensor when you change the number of detection resources, it will generate hardware alarms. Contact Sourcefire Support for information about how to clear those hardware alarms.

- If you create a detection engine, only that detection engine is started (although other CPUs may be restarted to rebalance the processing load).

Keysight, EX-1004
Page 208

- If you delete a detection engine or interface set, all detection engines on the sensor are restarted.
- If you create an interface set, nothing is restarted.
- If you change the name or description of an interface set or detection engine, nothing is restarted.

3D9900 Sensors

- If you change any non-description field in the interface set configuration (such as link propagation state) or detection engines configuration (such as number of resources), all detections engines will restart.

Other Sensors

- If you change which network interfaces are used by an interface set, all the detection engines on the sensor are restarted.
- If you change an interface set's transparent mode setting, or interface set type, all detection engines assigned to that interface set are restarted.
- If you change a detection engine's interface set, all detection engines on the sensor are restarted.
- If you change the number of detection resources allocated to a detection engine, all the detection engines on the sensor are restarted.
- If you change the detection engine type for a detection engine, that detection engine is restarted.
- When you create a detection engine, all the detection engines on the sensor are restarted because the total number of allocated resources has changed.
- If you delete a detection engine or interface set, all detection engines on the sensor are restarted.
- If you create an interface set, nothing is restarted. A restart occurs only when you assign a detection engine to the interface set.
- If you change the name or description of an interface set or detection engine, nothing is restarted.
- If you change system settings and your sensor and Defense Center are separated by a NAT device, the automatic reconnect link will not function. Manually reconnect the appliance.
- If you have installed a custom HTTPS certificate, you may have to accept an exception when you log back in, unless you reconnect manually.

Keysight, EX-1004
Page 209

Make sure you plan these actions for times when they will have the least impact on your deployment.

TIP! On your 3D Sensor Software for Crossbeam Systems X-Series, you may want to remove any affected VAPs from the load-balanced list until the associated detection engines restart, then reinstate the VAPs. For more information, see the *Sourcefire 3D Sensor Software for X-Series Installation Guide*.

To edit an existing detection engine:

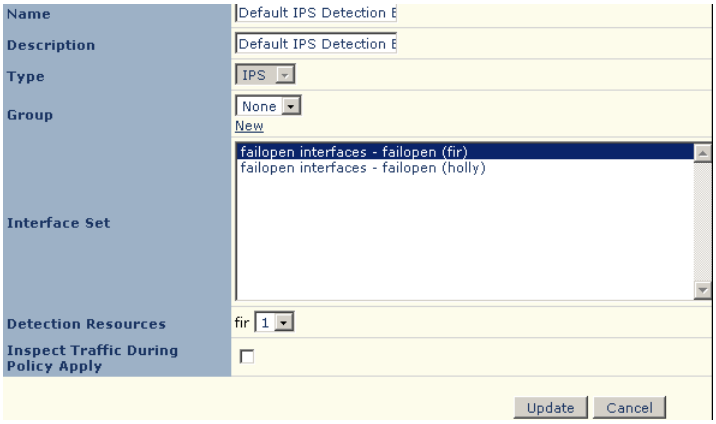
Access: Admin

1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

2. Click **Edit** next to the detection engine you want to modify.

   The Edit Detection Engine page appears.



   You can modify the name, description, group, and number of detection resources for the detection engine. You cannot modify the detection engine type. If you need to change the detection engine type, you must delete the detection engine and create a new one. In the case of an IPS detection engine you can also select if traffic is inspected while a policy is being applied.

   TIP! The **Inspect Traffic During Policy Apply** option is not available on 3D500, 3D1000, or 3D3800 sensors.

3. Click **Save**.

   Your changes are saved.

Keysight, EX-1004
Page 210

Deleting a Detection Engine

Platform: DC/3D Sensor

Component: None

Use the following procedure to delete a detection engine.

WARNING! Do **not** delete a detection engine that is in use. Also, you should not delete a detection engine that is used as a constraint in one or more compliance rules; you should first delete (or modify) the constraint in all rules in which it is used. For information on modifying compliance rules, see Modifying a Rule on page 1260.

To delete a detection engine:

Access: Admin

1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

2. Click **Delete** next to the detection engine you want to delete.

3. At the prompt, confirm that you want to delete the detection engine.

   The detection engine is deleted; however, a record of the detection engine is retained so that events generated by that detection engine are viewable.

Using Detection Engine Groups

Platform: DC/MDC/Standalone Sensor

Component: None

You can use detection engine groups to combine similar detection engines. These groups make it easier to apply policies to detection engines that have similar purposes.

See the following sections for more information:

- Creating Detection Engine Groups on page 211
- Editing Detection Engine Groups on page 212
- Deleting Detection Engine Groups on page 213

Creating Detection Engine Groups

Platform: DC/MDC/Standalone Sensor

Component: None

The following procedure explains how to create a detection engine group.

Keysight, EX-1004
Page 211

To create a detection engine group:

Access: Admin

1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

2. Click **Create Detection Engine Group**.

   The Create Detection Engine Group page appears.



3. Type a name for the detection engine group in the **Group Name** field.
4. Click **Save**.

   The Detection Engine page appears again. You can add detection engines to this group by clicking **Edit** next to a detection engine name and, on the Edit Detection Engine page, adding the detection engine to the group and clicking **Update**.

Editing Detection Engine Groups

Platform: DC/MDC/Standalone Sensor

Component: None

The following procedure explains how to edit a detection engine group. You must create a detection engine group before you can edit it. See Creating Detection Engine Groups on page 211.

To edit a detection engine group:

Access: Admin

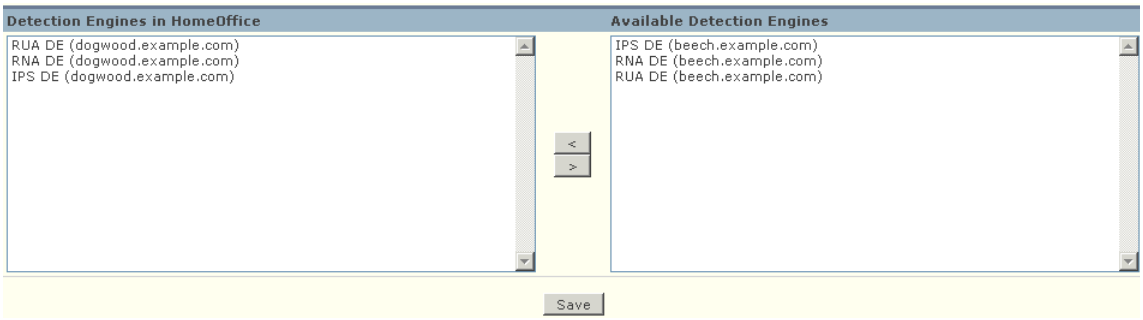
1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

Keysight, EX-1004
Page 212

2. Click **Edit** for the detection engine group.

   The Detection Engine Group Edit page appears.



3. Select available detections engines and to move them to the detection engine group with the arrow buttons.

   You can also move detection engines out of the detection engine group.

4. Click **Save** to add the selected detection engines to the detection engine group.

   The Available Detection Engines page appears.

Deleting Detection Engine Groups

Platform: DC/MDC/Standalone Sensor

Component: None

When you delete a detection engine group, any detection engines in the group are automatically ungrouped; they are not deleted.

To delete a detection engine group:

Access: Admin

1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

2. Click **Delete** next to the name of the detection engine group.

   The detection engine group is deleted.

Using Variables within Detection Engines

Platform: DC/MDC/Standalone Sensor

Component: IPS

A *system default variable* sets a variable value on your Sourcefire 3D Sensor or Defense Center that IPS uses by default unless it is overridden by a policy-specific or detection-engine-specific value for the same variable. You can associate a

Keysight, EX-1004
Page 213

system default variable with a specific detection engine and give the resulting *detection-engine-specific variable* an explicit value for that detection engine. When you apply an intrusion policy to that detection engine, IPS can use the value of the detection-engine-specific variable in rules you enable in your policy to monitor network traffic and generate events. For information on policy-specific variables, which are specific to the policy in which they are created, see Creating New Policy-Specific Variables on page 383.

For example, the intrusion rules in an intrusion policy take advantage of certain system default variables such as `HOME_NET` and `EXTERNAL_NET` to look for exploits that originate outside your network and are targeted against hosts within your network. You can define `HOME_NET` in your system default variable to encompass your internal address range (for example, 10.10.0.0/16).

However, if you have created your detection engines so that one detection engine monitors one class of hosts (in this example, hosts in your network's DMZ in the range 10.10.30.0/24) and another monitors a different class (for example, hosts in your accounting department in the address range 10.10.90.0/24), you can use detection-engine-specific variable values to tailor your detection capabilities to more closely match your infrastructure.

In the system default variable used in the intrusion policy:

```
HOME_NET = 10.10.0.0/16
```

In the detection engine named DE_DMZ:

```
HOME_NET = 10.10.30.0/24
```

In the detection engine named DE_ACCT:

```
HOME_NET = 10.10.90.0/24
```

If you later create another detection engine that monitors the rest of your network, which includes a mixed address space, you can use the system default variable value rather than creating another detection-engine-specific value for `HOME_NET`.

You can also create new variables for use only within the context of the detection engine. You can create detection-engine-specific variables and set detection-engine-specific values for system default variables within an intrusion policy or from the detection engine Variable List page. Configuration details in this section relate to the detection engine Variable List page. For configuration details related to setting detection-engine-specific variables within an intrusion policy, see Creating New Variables on page 377.

Creating a detection-engine-specific variable from the detection engine Variable List page also creates a corresponding system default variable with the value set to **any**. You can view the explicit detection-engine-specific value you configured in the list of variables for the detection engine within each policy, or on the detection engine Variable List page for the detection engine. You can view the corresponding new system default variable in the list of system default variables within each policy, and on the Variable list page for all other detection engines where it is listed with the value set to **Policy Defined**, which means that the value

Keysight, EX-1004
Page 214

specified in the policy will be used when you apply the policy. Optionally, you can modify the variable in the intrusion policies and detection engines where it is added automatically to give it a specific definition. When they exist, a detection-engine-specific variable value takes precedence over a policy-specific or system default value for the same variable. If you disable a variable defined on the Variable List page by resetting the variable, the definition reverts to the definition in the intrusion policy the next time you apply the policy.

Variables use the same syntax and must follow the same guidelines regardless of whether you create or define them from within intrusion policies or from the detection engine Variable List page. See Creating New Variables on page 377 and Modifying Variables on page 374 for more information.

IMPORTANT! You cannot use variables with RNA detection engines.

For more information, see the following sections:

- Assigning Values to System Default Variables in Detection Engines on page 215
- Creating New Variables for Detection Engines on page 217
- Deleting and Resetting Variables on page 218
- Configuring Custom Variables in Detection Engines on page 219
- Using Portscan-Only Detection Engines on page 220

Assigning Values to System Default Variables in Detection Engines

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can assign detection-engine-specific values to system default variables. For an explanation see Using Variables within Detection Engines on page 213.

To assign a detection-engine-specific value to a system default variable:

Access: Admin

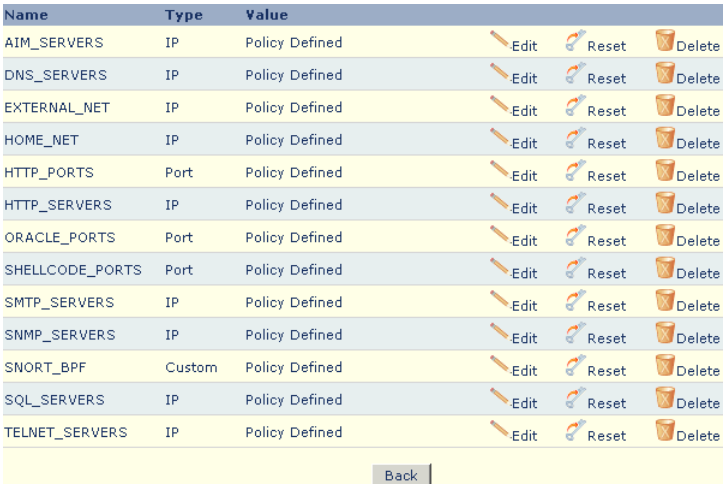
1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

Keysight, EX-1004
Page 215

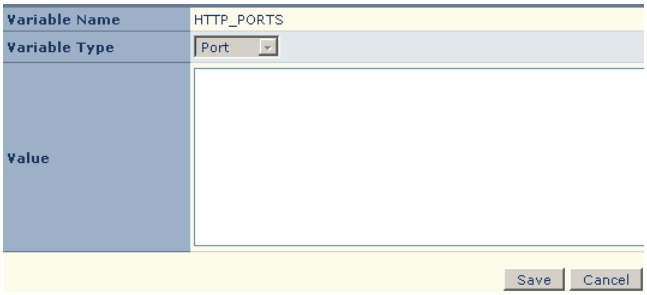
2. Click **Variables** next to the detection engine where you want to define a variable value.

   The Variable List page appears. The value for each of the variables defaults to the value within the intrusion policy that is applied to the detection engine.



3. Click **Edit** next to the variable you want to define.

   The Variable Binding page appears.



4. Enter a value for the variable and click **Save**.

   Note that you cannot edit the Variable Type. See Creating New Variables on page 377 for information about variable syntax.

   The Variable List page appears again and shows the new value for the variable. The variable takes effect the next time you apply an intrusion policy to the detection engine, as described in Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 216

Creating New Variables for Detection Engines

Platform: DC/MDC/Standalone Sensor

Component: IPS

When you create an intrusion policy, you can associate detection-engine-specific variable definitions with the policy. For an explanation see Using Variables within Detection Engines on page 213.

To create a new variable for a detection engine:

Access: Admin

1. Select **Operations > Configuration > Detection Engines**.

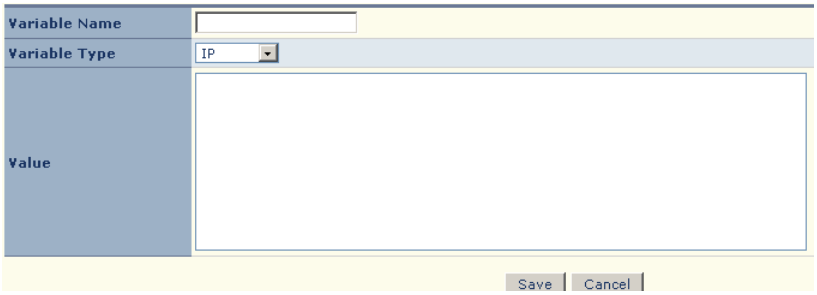
   The Detection Engines page appears.

2. Click **Variables** next to the detection engine where you want to define a variable value.

   The Variable List page appears.

3. Click **Add Variable**.

   The Variable page appears.



4. In the **Variable Name** field, enter a name for the variable.
5. From the **Variable Type** drop-down list, select `IP`, `Port`, or `Custom`.
   - See Defining IP Addresses in Variables and Rules on page 399 for more information if you are defining a IP address-based variable.
   - See Defining Ports in Variables and Rules on page 401 for more information if you are defining a port-based variable.
   - See Understanding Custom Variables on page 391 if you are defining a special-purpose custom variable with one of the reserved variable names described in the Custom Variables table on page 391.

Keysight, EX-1004
Page 217

6. In the **Value** field, enter a value for the variable and click **Save**. See Creating New Variables on page 377 for information about the syntax for variables.

   The Variable List page appears again and shows the new variable and its value.

   The variable is created and is accessible to all policies as a system default variable. It is listed in the variable list for the detection engine in all intrusion policies with the explicitly set value, and listed for all other detection engines on the Variable List page with a value of **Policy Defined**. The variable takes effect the next time you apply an intrusion policy to the detection engine, as described in Applying an Intrusion Policy on page 347.

IMPORTANT! Each new detection engine variable adds a system variable with a value of **any** that is accessible in all your intrusion policies. Creating the new detection engine variable also lists the description **Policy Defined** for all other IPS detection engines on the Variable List page, meaning that the value specified in the policy will be used when you apply the policy. In any intrusion policy that you apply to a different detection engine and do not explicitly set a policy-defined or detection-engine-specific variable to override the value of the system variable, the value **any** will be used.

Deleting and Resetting Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can reset the value of a variable on the Variable List page and the variable reverts to the value defined in the intrusion policy the next time you apply the intrusion policy to the detection engine. You can also delete variables that you created within the context of the detection engine. You cannot delete predefined system variables within an intrusion policy. You can delete predefined system variables on the detection engine Variable List page, but only if they are not used in any active or inactive rule within the system.

To delete or reset variables on a detection engine:

Access: Admin

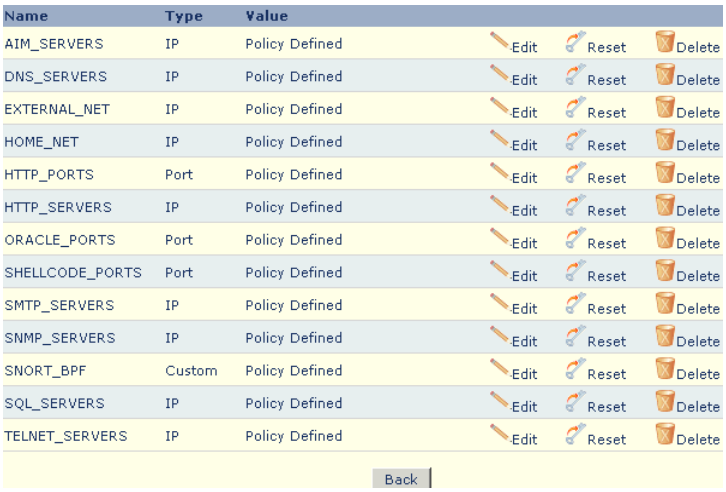
1. Select **Operations > Configuration > Detection Engines**.

   The Detection Engines page appears.

Keysight, EX-1004
Page 218

2. Click **Variables** next to the detection engine where you want to delete or reset a variable value.

   The Variable List page appears.



Name	Type	Value			
AIM_SERVERS	IP	Policy Defined	Edit	Reset	Delete
DNS_SERVERS	IP	Policy Defined	Edit	Reset	Delete
EXTERNAL_NET	IP	Policy Defined	Edit	Reset	Delete
HOME_NET	IP	Policy Defined	Edit	Reset	Delete
HTTP_PORTS	Port	Policy Defined	Edit	Reset	Delete
HTTP_SERVERS	IP	Policy Defined	Edit	Reset	Delete
ORACLE_PORTS	Port	Policy Defined	Edit	Reset	Delete
SHELLCODE_PORTS	Port	Policy Defined	Edit	Reset	Delete
SMTP_SERVERS	IP	Policy Defined	Edit	Reset	Delete
SNMP_SERVERS	IP	Policy Defined	Edit	Reset	Delete
SNORT_BPF	Custom	Policy Defined	Edit	Reset	Delete
SQL_SERVERS	IP	Policy Defined	Edit	Reset	Delete
TELNET_SERVERS	IP	Policy Defined	Edit	Reset	Delete

Back

3. You have two options:
   - To disable the variable value defined in the IPS detection engine and revert to the variable value defined in the policy, click **Reset** next to the name of the variable.

     The variable is reset and Policy Defined appears in the Value column.
   - To delete a locally created variable, click **Delete** next to the name of the variable.

     The variable is deleted from the detection engine the next time you apply an intrusion policy to the detection engine.

Configuring Custom Variables in Detection Engines

Platform: DC/MDC/Standalone Sensor

Component: IPS

Custom variables allow you to configure special IPS features that you cannot otherwise configure via the web interface. You create a detection-engine-specific custom variable by setting an explicit value for a reserved predefined system variable, or by creating a variable using a specific reserved name. You then define the variable value with a set of instructions appropriate to the function the variable provides. For more information, see Understanding Custom Variables on page 391.

You can set an explicit detection engine value for the predefined SNORT_BPF custom system variable.

Keysight, EX-1004
Page 219

You can add a new USER_CONF detection engine variable using the reserved name USER_CONF.

To configure the SNORT_BPF custom variable for a detection engine:

Access: P&R Admin/Admin

- To set an explicit detection-engine-specific value for SNORT_BPF using the existing system default variable, see Assigning Values to System Default Variables in Detection Engines on page 215.

To configure the USER_CONF custom variable for a detection engine:

Access: P&R Admin/Admin

- To create USER_CONF as a new detection-engine-specific variable using the reserved name USER_CONF, see Creating New Variables for Detection Engines on page 217.

Using Portscan-Only Detection Engines

Platform: DC/MDC/Standalone Sensor

Component: IPS

If you configure a sensor to use multiple detection resources within a single IPS detection engine, a portion of the traffic that the 3D Sensor sees is directed to each detection resource for processing. Internal logic on the sensor ensures that packets belonging to the same session are directed to the same resource for analysis. In this way, the sensor can process more packets with greater efficiency.

One downside to using multiple detection resources is that no single resource sees all the traffic on a network segment, which is a requirement for the portscan preprocessor. To overcome this issue, you can create a portscan-only intrusion policy and apply it to a portscan-only detection engine on the sensor. The following process describes how you can configure your sensor to detect portscans in addition to other exploits against your network assets.

To configure a portscan-only detection engine:

1. Using the Defense Center's web interface, create an interface set that includes the network interfaces you want to use on the sensor. Multiple detection engines will use this interface set.

   The interface set can be passive, inline, or inline with fail open depending on how your sensor is deployed.

2. Create an IPS portscan-only detection engine and assign one detection resource to it. Make sure you use the interface set that you created in step 1.

Keysight, EX-1004
Page 220

3. Create another IPS detection engine that uses up to the remaining number of detection resources and the interface set that you created in step 1.

IMPORTANT! A portscan-only intrusion policy is able to process up to three times more traffic than a more complex intrusion policy because it uses fewer CPU resources. However, Sourcefire recommends that you monitor the performance of your sensor to make sure that the portscan-only detection engine is able to keep up with the multi-resource detection engine. Depending on the traffic mix on your network, you may need to adjust the number of resources in the multi-resource detection engine. Remember that the portscan-only detection engine can use only one detection resource.

---

4. Create and apply an intrusion policy for the multi-resource detection engine.

   Make sure you match the type of intrusion policy to the type of interface set that you created in step 1. Also, make sure you disable portscan detection in this policy.

5. Create and apply an intrusion policy to the portscan-only detection engine. The policy should inherit or be set to the following settings in the layer in your intrusion policy where you enable portscan detection (See Creating an Intrusion Policy on page 338, Working with Layers, and Applying an Intrusion Policy on page 347 for more information):

   - Select the No Rules Active **Base Policy** and make sure the **Drop When Inline** check box is cleared (disabled). See Selecting the Base Policy on page 364 and Setting Drop Behavior in an Inline Deployment on page 358 for more information. Note that all rules are disabled on the Rules page.
   - Ensure that the **DCE/RPC Configuration** preprocessor, the **HTTP Configuration** preprocessor, the **SMTP Configuration** preprocessor (under **Application Layer Preprocessors**), and **Back Orifice Detection** (under **Specific Threat Detection**) are disabled. See Modifying Advanced Settings on page 478 for more information.
   - Ensure that **OPSEC Configuration** (under **External Responses**) is disabled.
   - Enable **IP Defragmentation** (under **Transport/Network Layer Preprocessors**) and make sure it is configured for your environment (using the Hosts option). See Modifying Advanced Settings on page 478 for more information.
   - You should not change the default settings for **Checksum Verification** or **Packet Decoding** (under **Transport/Network Layer Preprocessors**), items listed under **Performance Statistics**, or Rule Processing Configuration.

Keysight, EX-1004
Page 221

- Enable **Portscan Detection** and configure it for your network environment. See Detecting Portscans on page 651 for more information.
- Make sure portscan rules are enabled for the types of portscans you configure.

IMPORTANT! 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 the Portscan Detection SIDs (GID:122) table on page 659 for more information.

You do not need to set up variables for this policy.

6. Review the resulting intrusion events to ensure that you are receiving the events you expect.

Using Interface Sets

Platform: DC/3D Sensor

Component: None

An interface set refers to a grouping of one or more sensing interfaces on a sensor, although a sensing interface can belong to only one interface set at a time. The Sourcefire 3D System supports three types of interface sets, but the interface options available to you depend on the type of sensor and the capabilities of its sensing interfaces. The three interface types are passive, inline, and inline with fail open.

- **Passive**: Use a passive interface set if you deployed the sensor out of band from the flow of network traffic.
- **Inline:** Use an inline interface set if you deployed the sensor inline on your network and the sensing interfaces do **not** support automatic fail-open capabilities. Note that you can use any two of the non-fail-open interfaces on the sensor's network interface cards as part of an inline interface set. The exception is on 3D9900s, where pairs are pre-determined.
- **Inline with Fail-Open:** Use an inline with fail open interface set if you deployed the sensor inline on your network and the sensing interfaces **do** support automatic fail-open capabilities. Note that you must use paired fail-open interfaces on the sensor's network interface cards for an inline with fail open interface set.

Keysight, EX-1004
Page 222

You can use RNA or RUA to monitor the traffic that passes through any of the three types of interface sets.

IMPORTANT! On a 3D3800 or 3D5800 sensor, if you plan to use RNA to monitor either an inline or inline with fail open interface set, you **must** either configure an IPS detection engine that uses that interface set, as well as apply an intrusion policy to that detection engine, or configure the interface set in tap mode. Otherwise, the RNA detection engine monitoring that interface set will not see any traffic. If you are monitoring the same inline interface set with both IPS and RNA or RUA, and the IPS detection engine fails for any reason, the RNA or RUA detection engine monitoring that interface set will not see any traffic until the IPS detection engine restarts. Neither RNA nor RUA are supported on the 3D9800 sensor.

To list the available interface sets:

Access: Admin

▶ Select **Operations > Configuration > Detection Engines > Interface Sets**.

You can sort the available interface sets by group, set type, sensor, or PEP policy.

See the following sections for more information about interface sets:

- Understanding Interface Set Configuration Options on page 223
- Creating an Interface Set on page 230
- Creating an Inline Interface Set on page 234
- Editing an 8000 Series Interface Set on page 241
- Editing an Interface Set on page 239
- Deleting an Interface Set on page 244
- Inline Fail Open Interface Set Commands on page 246
- Using Clustered 3D Sensors on page 250

Understanding Interface Set Configuration Options

Platform: DC/3D Sensor

Component: None

There are a number of configuration variables to consider when you configure interface sets:

- With the exception of the Virtual 3D Sensor, you can set up any of your 3D Sensor interfaces in passive, inline, or inline with fail-open mode. The Virtual 3D Sensor supports only passive mode operation.
- You can configure interfaces on most sensors to transparent inline mode.
- On selected sensors you can set interfaces to tap mode.

Keysight, EX-1004
Page 223

- Some installations require that the link state be propagated, and most sensor interfaces provide that option.
- Sensors with gigabit Ethernet interfaces support jumbo frames.

  Note that Crossbeam-based software sensors also support jumbo frames, but you must configure jumbo frame support using the Crossbeam CLI. You do **not** use the Defense Center to configure jumbo frame support for these sensors.
- 3D Sensors deployed in networks that are highly sensitive to latency can use the automatic application bypass option.
- Only 3D9900 sensors provide a fail-safe option that works with inline interface sets.
- Only 3D9900 sensors provide the PEP feature. For more information on the PEP feature, see Using PEP to Manage Traffic on page 1510.

See the following table for a list of 3D Sensors and each of their applicable interfaces features.

Supported Features by 3D Sensor Model

3D Sensor Model	Transparent Inline Mode	Link State Propagation Mode	Tap Mode	Jumbo Frames	Automatic Application Bypass	Enable Fail-safe	PEP
Virtual 3D Sensor	Yes				Yes		
3D500		Yes			Yes		
3D1000	Yes	Yes			Yes		
3D2000	Yes	Yes			Yes		
3D2100	Yes	Yes			Yes		
3D2500	Yes	Yes			Yes		
3D3000	Yes	Yes			Yes		
3D3500	Yes	Yes			Yes		
3D3800	Yes	Yes	Yes				
3D4500	Yes	Yes			Yes		
3D5800	Yes	Yes	Yes				
3D6500	Yes	Yes		Yes	Yes		

Keysight, EX-1004
Page 224

Supported Features by 3D Sensor Model (Continued)

3D Sensor Model	Transparent Inline Mode	Link State Propagation Mode	Tap Mode	Jumbo Frames	Automatic Application Bypass	Enable Fail-safe	PEP
3D9800	Yes		Yes	Yes			
3D9900	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crossbeam-based software sensors	Yes			Yes			
8000 Series	Yes	Yes	Yes	Yes	Yes	Yes	Yes

See the following sections for more information:

- Types of Interface Sets on page 225
- Transparent Inline Mode on page 226
- Tap Mode on page 227
- Link State Propagation Mode on page 228
- Jumbo Frames on page 229
- Automatic Application Bypass on page 229
- Enabling Fail-Safe on page 230

Types of Interface Sets

When you create an interface set, you can choose one of three types:

- Passive

  For most sensors, a passive interface set can encompass any number of the available sensing interfaces on a sensor. On Crossbeam-based software sensors, a passive interface set can include up to 16 interfaces, though Sourcefire recommends you use no more than eight.

Keysight, EX-1004
Page 225

- Inline

  For most sensors, an inline interface set can include any two interfaces. The interfaces do not have to be on the same network cards, but you should avoid using an on-board interface.

  On a 3D3800 or 3D5800 sensor, an inline interface set can include up to four interface pairs, and an inline interface set on a 3D9800 sensor can include up to the total number of interface pairs on the sensor. Note that interface pairs on the same fiber-based NIM will act as fail open interfaces even if you assign them to an inline interface set. That is, if the power fails or the Snort process halts, network traffic continues to flow through the sensor as it would for an inline with fail open interface set.

  On Crossbeam-based software sensors, an inline interface set can include up to eight interface pairs, though Sourcefire recommends you use no more than four.

- Inline with Fail Open

  For most sensors, an inline with fail open interface set must include exactly one interface pair. However, an inline with fail open interface set on a 3D3800 or 3D5800 sensor can include up to four interface pairs, and an inline with fail open interface set on a 3D9800 sensor can include up to the total number of interface pairs on the sensor.

You can set up multiple detection engines to use a single interface set, except on the 3D9800 sensor, which only supports a single IPS detection engine. For example, you could create a single passive interface set and create two detection engines, one for an IPS and the other for RNA, then apply different policies to the detection engines.

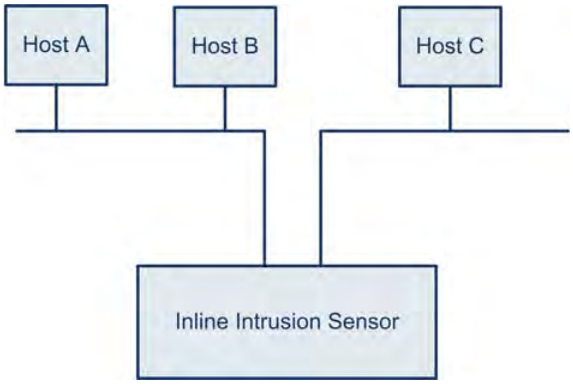
Transparent Inline Mode

Transparent inline mode is a feature for inline interface sets and is not available for Passive interface sets.

If you choose the Inline or Inline with Fail Open option, the Transparent Inline Mode option is enabled by default, except for the 3D500 and the Virtual 3D Sensor. It is not available on the 3D500 and available but not a default configuration on the Virtual 3D Sensor. This allows the sensor to act as a "bump in the wire" and means that the sensor forwards all the network traffic it sees regardless of its source and destination.

Keysight, EX-1004
Page 226

If you disable this option, a sensor acts as a bridge. Over time, the sensor learns which hosts are on which side of the inline interface, and forwards packets accordingly. For example, consider the following diagram.



If your sensor is deployed inline (or more precisely, if your sensor includes a detection engine with an inline interface set) and the Transparent Inline Mode option **is** selected, then if the sensor sees network traffic from Host A to Host B, it allows the traffic to pass through the interface even though Host A and Host B are on the same side of the sensor.

If the sensor is inline and you are **not** using transparent inline mode, when the sensor sees traffic from Host A to Host B, it does **not** allow the traffic to pass through the interface to the side of the network with Host C. Only traffic between Host A and Host C or between Host B to Host C is allowed to pass.

Keep in mind that if you create an inline interface set but do not use transparent inline mode, you must be especially careful not to create loops in your network infrastructure.

3Dx800 sensors run in transparent inline mode, and you cannot disable it.

Tap Mode

Tap mode is available for the 3D3800, 3D5800, 3D9900, 8000 Series, and on later versions of 3D9800 3D Sensor when you create an inline or inline with fail open interface set.

TIP! 3D9800 sensors with earlier versions of firmware do not support tap mode. The Sourcefire 3D System checks the 3D9800 firmware version and displays the optional tap mode check box in the Create Interface Set page when appropriate.

With tap mode, the sensor is deployed inline, but instead of the packet flow passing through the sensor, a copy of each packet is sent to the sensor 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

Keysight, EX-1004
Page 227

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 sensors that are deployed inline. For example, you can set up the cabling between the sensor and the network as if the sensor were inline and analyze the kinds of intrusion events the sensor 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 sensor inline, you can disable tap mode and begin dropping suspicious traffic without having to reconfigure the cabling between the sensor and the network.

IMPORTANT! On a 3D3800 or 3D5800 sensor, if you plan to use RNA to monitor either an inline or inline with fail open interface set, you **must** either configure an IPS detection engine that uses that interface set, as well as apply an intrusion policy to that detection engine, or configure the interface set in tap mode. Otherwise, the RNA detection engine monitoring that interface set will not see any traffic. If you are monitoring the same inline interface set with both IPS and RNA or RUA, and the IPS detection engine fails for any reason, the RNA or RUA detection engine monitoring that interface set will not see any traffic until the IPS detection engine restarts. Neither RNA nor RUA are supported on the 3D9800 sensor.

Link State Propagation Mode

Link state propagation mode is a feature for interface sets in the inline fail-open mode so both pairs of an inline pair track state. It is also available on 3D9900s in both the inline and inline fail-open mode. It is not available for passive interface sets.

IMPORTANT! Fiber interface sets configured as inline fail-open, other than those on 3D9900s, must be in hardware bypass mode for link state propagation to function correctly. For more information about fiber interface sets and hardware bypass, see Removing Bypass Mode on Inline Fail Open Fiber Interfaces on page 246.

Link state propagation mode automatically brings down the second interface in the interface pair when one of the interfaces in an inline interface set goes down. When the downed interface comes back up, the second interface automatically comes back up, too. In other words, if the link state of one interface changes, the

Keysight, EX-1004
Page 228

link state of the other interface is changed automatically to match it. Link state propagation is available for both copper and fiber fail-open NIMs.

IMPORTANT! Crossbeam-based software sensors and 3D9800 sensors do **not** support link state propagation.

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.

Jumbo Frames

Jumbo frames are Ethernet frames with a frame size greater than the standard 1518 bytes. Typical maximum-sized jumbo frames are 9018 bytes. Most gigabit Ethernet network interface cards support jumbo frames to increase efficiency.

The following table lists the 3D Sensors that support jumbo frames and describes how and where you configure jumbo frame support.

Jumbo Frame Support by Sensor Model

For this sensor model...	You can...
3D6500 and 3D9900 with gigabit Ethernet interfaces, and 8000 Series	set the maximum frame size for the appropriate interfaces when you use the Create Interface Set page to create an interface set.
3D9800	do nothing. The 3D9800 accepts the maximum size frame by default.
Crossbeam-based software sensors	enable jumbo frame support and set the maximum frame size using the Crossbeam CLI. Do **not** use the Defense Center to configure jumbo frame support for these sensors. For more information, see Sourcefire's *XOS Installation and Configuration Guide*.

Note that frames larger than the configured maximum frame size are silently dropped by the sensor.

Automatic Application Bypass

The automatic application bypass feature allows you to balance packet processing delays with your network's tolerance for packet latency. You can apply automatic application bypass on an interface set basis. The feature functions with both passive and inline interface sets; however, it is most valuable in inline deployments.

Keysight, EX-1004
Page 229

Automatic application bypass limits the time allowed to process packets through an IPS, RNA, or RUA detection engine and allows packets to bypass the detection engine if the time is exceeded. If the automatic application bypass is triggered, Snort automatically restarts within ten minutes.

The automatic application bypass option is off by default. You can change the bypass threshold if the option is selected. The default setting is 750 milliseconds (ms). The valid range is from 250 ms to 60,000 ms.

To see a list of which 3D Sensors you can use Automatic Application Bypass Monitoring on, see the Supported Features by 3D Sensor Model table on page 224.

If a detection engine is bypassed, 3D Sensors generate a health monitoring alert. For more information on the health monitoring alert, see Configuring Automatic Application Bypass Monitoring on page 1907.

Enabling Fail-Safe

The Create Interface Set page includes an additional option for 3D9900 sensors: the Enable Fail-Safe option. The Enable Fail-Safe option is only available on inline interface configurations. When you enable the Enable Fail-Safe option, traffic is allowed to bypass detection and continue through the sensor. 3D9900 sensors monitor internal traffic buffers and bypass detection engines if those buffers are full.

Load Balancing

You can choose to load balance your IP-to-IP tunneled traffic on either the inner header or outer header. For example, if all your traffic is tunneled, you should load balance on the IP header inside the tunnel because there is only one IP header outside the tunnel.

Creating an Interface Set

Platform: DC/3D Sensor

Component: None

An interface set is a collection of one or more sensing interfaces on your appliance. For information about their use, see Using Interface Sets on page 222.

Note that 8000 Series sensors are fully configured by default. To create a new interface set, you edit the default configuration. See Editing an 8000 Series Interface Set on page 241.

IMPORTANT! The procedure for creating an inline interface set for 3Dx800 sensors is slightly different. For more information, see the next section, Creating an Inline Interface Set.

Keysight, EX-1004
Page 230

To create an interface set:

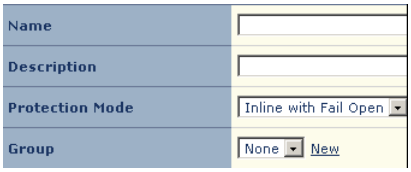
Access: Admin

1. Select **Operations > Configuration > Interface Sets**.

   The Interface Sets page appears.

2. Click **Create Interface Set**.

   The Create Interface Set page appears.



3. Type a name and description for the new interface set in the **Name** and **Description** fields.You can use alphanumeric characters and spaces.

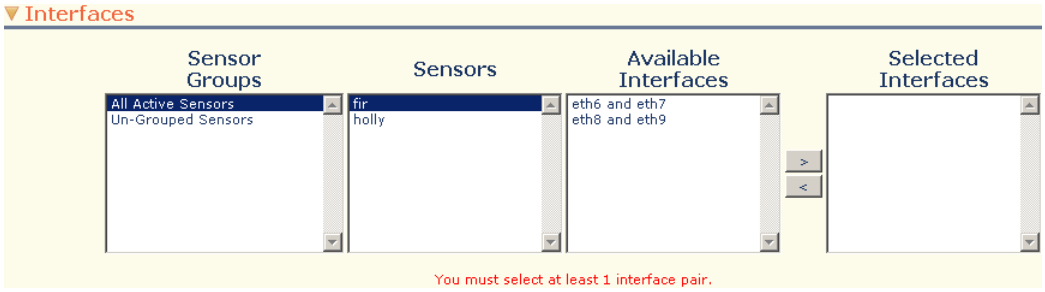
   TIP! Using a descriptive name, such as `chestnut with failopen` (sensor name and interface set type) is more helpful than `interface set 2`.

4. Select the type of interface you want to create, **Passive**, **Inline**, or **Inline with Fail Open**, from the **Interface Set Type** drop-down list.

   Note that some sensors do not support every interface set type.

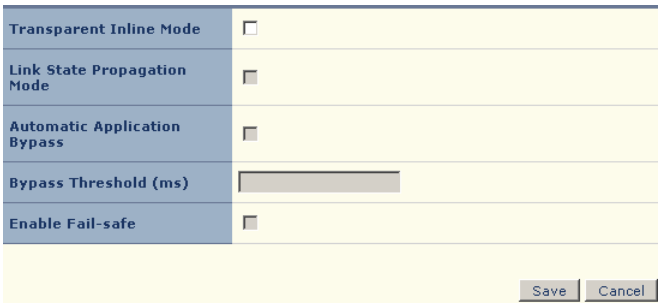
5. Optionally, select an existing interface set group or select **Create New Group** to create a new interface set group. See Using Interface Set Groups on page 244 for more information.

6. From the Interfaces menu, select a Sensor Group and Sensor. The available interfaces are displayed.



Keysight, EX-1004
Page 231

7. Select an interface and move it from Available to Selected by highlighting the interface and clicking the arrow. Interfaces vary depending on Protection Mode.
   - Passive: Select one interface.
   - Inline: Select two or more interfaces. Use the control key to select multiple interfaces.
   - Inline with Fail-Open: Select one or more preconfigured interface pairs.
8. Optionally, from the Advanced menu, you can change any field that is not grayed out. Options vary depending on your Protection Mode.



9. Optionally, if you selected the **Inline** or **Inline with Fail Open** option, clear the **Transparent Inline Mode** check box to disable transparent mode.
10. If you selected either the Inline or Inline with Fail Open option and you are not configuring a Crossbeam-based software sensor, then optionally, select **Link State Propagation Mode**. This option is especially useful if the routers on your network are able to re-route traffic around a network device that is down.

IMPORTANT! Link state propagation and automatic application bypass are not supported on Sourcefire 3D Sensor Software for X-Series platforms. You can, however, set jumbo frame options on the Crossbeam CLI.

---

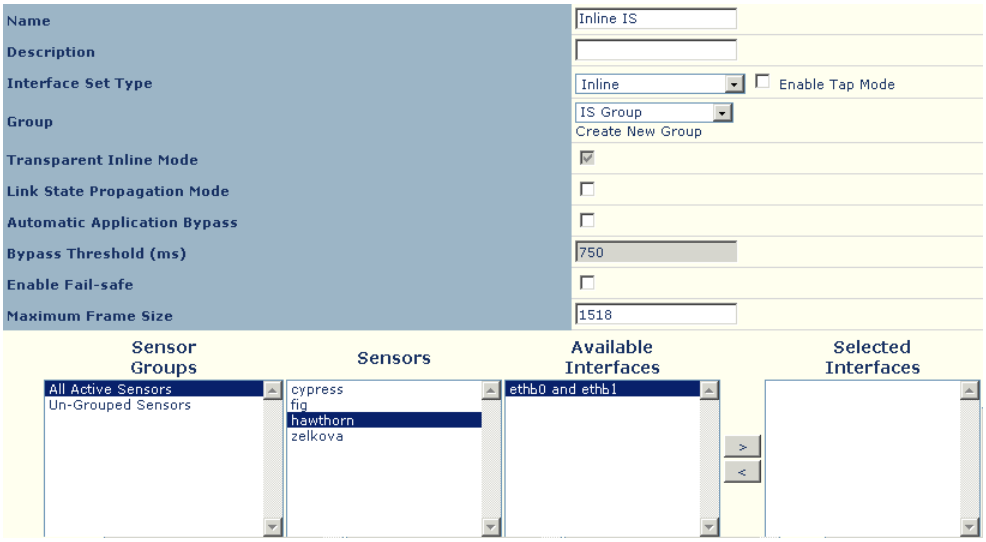
11. Optionally, select **Automatic Application Bypass** if your network is sensitive to latency. When the option selected, you can select a **Bypass Threshold** in milliseconds (ms). The default setting is 750 ms and the valid range is from 250 ms to 60,000 ms. Automatic Application Bypass is most useful in inline applications.
12. Optionally, and if you are configuring an inline interface set on a 3D9900, you can select the **Enable Fail-safe** check box to enable traffic pass-through during application bypass.

Keysight, EX-1004
Page 232

13. Optionally, and if you are configuring an interface set on a 3D6500 or 3D9900 or 8000 Series, type a maximum frame size for your IP traffic in the **Maximum Frame Size** field. You can set any jumbo frame size between 1518 and 9018 bytes, inclusive.

    On the Defense Center only, a list of sensor groups appears, including a list of ungrouped sensors.

    The following shows a 3D9900 interface set.



14. If you are using a Defense Center, select the sensor group containing the sensors where you want to create the interface set. You can also select the ungrouped sensors.

    A list of sensors appears.

15. If you are using a Defense Center, select one of the sensors from the list.

    A list of network interfaces on the sensor appears.

Keysight, EX-1004
Page 233

16. Select the interfaces that you want to add from the Available Interfaces list and click the arrow button to add the interface to the Selected Interfaces list. You can use the Shift and Ctrl keys to select multiple interfaces at once.

    Determining which interface name corresponds with a physical interface on your sensor depends on the model:

    - For most 3D Sensors, log into the console and disconnect the network cable from the interface. A message appears on the console indicating the name of the interface (`eth1`, `eth2`, and so on). Remember to reconnect the network cable when you are finished.
    - For 3Dx800 sensors, the names that appear in the Available Interfaces list correspond to the slot number and interface location. For example, `s0.e0` corresponds to the leftmost interface on the network interface module (NIM) in I/O Slot 0 on the back of your appliance.
    - For 3D Sensor Software for Crossbeam Systems X-Series, the names that appear in the Available Interfaces list correspond to the device names you assigned to the circuits you created on the X-Series.

    For more information, see the *Installation Guide* for your sensor or sensor software.

    Different types of interface sets have different requirements. For example, for most sensors you can include all of the available interfaces in a passive interface set, but most inline interface sets must contain exactly two interfaces. Inline with fail open interface sets must contain one pair of interfaces from the same fail-open network card. For more information, see Types of Interface Sets on page 225.

    IMPORTANT! If you select an on-board interface rather than an interface on a network card, your sensor may not provide optimum performance.

17. Click **Save**.

    The interface set is created.

    TIP! After you create an interface set, make sure you reapply intrusion policies to the IPS detection engines on the affected sensor.

Creating an Inline Interface Set

Platform: DC/3D Sensor

Component: None

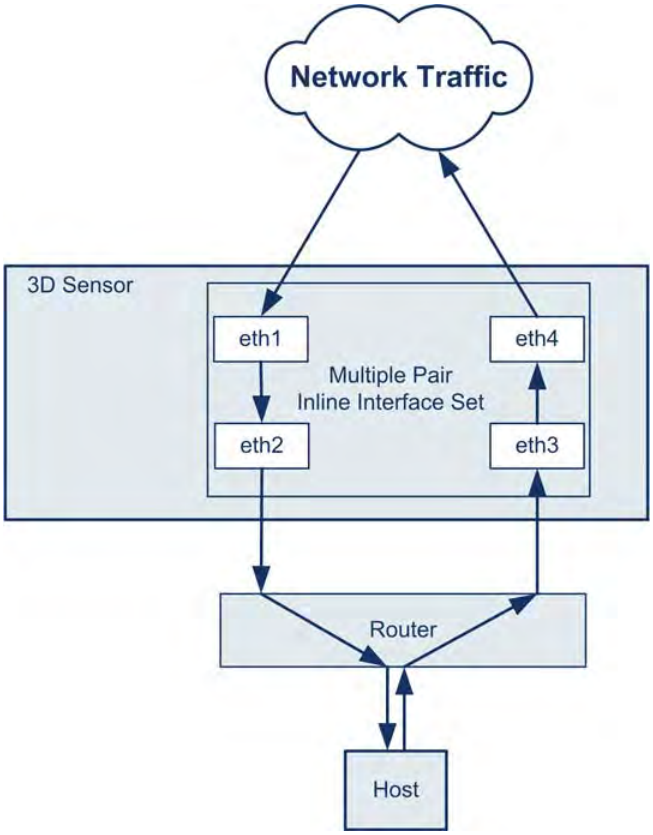
You can add multiple interface pairs to an inline interface set on 3D Sensors and Crossbeam-based software sensors. This is the default behavior during 3D Sensor installations. Using one interface set that includes all available inline interface pairs, you can apply a single policy and rapidly complete your initial

Keysight, EX-1004
Page 234

3D Sensor deployment. Later, you can refine policies for specific connected network segments and their requirements. Note that the interface sets are preconfigured on an 8000 Series sensor. To modify the configuration, edit the interface sets. See Editing an 8000 Series Interface Set on page 241.

TIP! Although the default interface set on 3D Sensors includes all the available inline interface pairs, in many cases you can improve performance by modifying the interface set to include only the inline interface pairs your network requires.

You can also use multiple interface pairs when your network employs asynchronous routing, as shown in the following graphic.



Your network may be set up to route traffic between a host on your network and external hosts through different interface pairs depending on whether the traffic is inbound or outbound. If you include only one interface pair in an interface set, the sensor might not correctly analyze your network traffic because a detection engine might see only half of the traffic.

Keysight, EX-1004
Page 235

For most 3D Sensors with inline interface sets, a software bridge is automatically set up to transport packets when the sensor restarts. Although some packets are transmitted without inspection during this time, no packets are lost.

IMPORTANT! On a 3D3800 or 3D5800 sensor, if you plan to use RNA to monitor either an inline or inline with fail open interface set, you **must** either configure an IPS detection engine that uses that interface set, as well as apply an intrusion policy to that detection engine, or configure the interface set in tap mode. Otherwise, the RNA detection engine monitoring that interface set will not see any traffic. If you are monitoring the same inline interface set with both IPS and RNA or RUA, and the IPS detection engine fails for any reason, the RNA or RUA detection engine monitoring that interface set will not see any traffic until the IPS detection engine restarts. Neither RNA nor RUA are supported on the 3D9800 sensor.

To create an inline interface set:

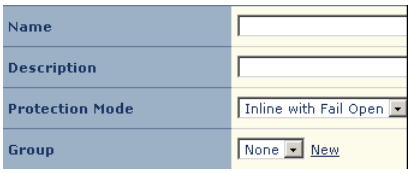
Access: Admin

1. Select **Operations > Configuration > Interface Sets**.

   The Interface Sets page appears.

2. Click **Create Interface Set**.

   The Create Interface Set page appears.



3. Type a name and description for the new interface set in the **Name** and **Description** fields.

   You can use alphanumeric characters and spaces.

4. Select the type of inline interface you want to create.

   - For an 3Dx800 sensor, choose either **Inline** or **Inline with Fail Open**, from the **Interface Set Type** drop-down list.
   - For Crossbeam-based software sensors, choose **Inline** from the **Interface Set Type** drop-down list.

   A list of sensor groups appears, including a list of ungrouped sensors.

5. Optionally, select an existing interface set group or select **Create New Group** to create a new interface set group. See Using Interface Set Groups on page 244 for more information.

Keysight, EX-1004
Page 236

6. Optionally, select **Automatic Application Bypass** if your network is sensitive to latency. When the option selected, you can select a **Bypass Threshold** in milliseconds (ms). The default setting is 750 ms and the valid range is from 250 ms to 60,000 ms.

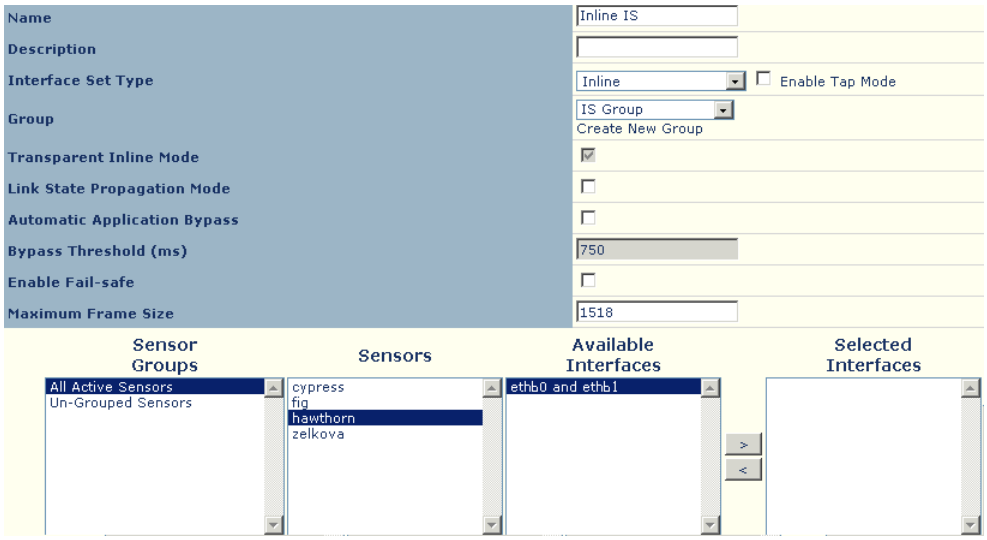
   IMPORTANT! Link state propagation and automatic application bypass are not supported on Sourcefire 3D Sensor Software for X-Series platforms. You can, however, set jumbo frame options on the Crossbeam CLI.

7. Optionally, and if you are configuring an interface set on a 3D9900, you can select the **Enable Fail-safe** check box to enable traffic pass-through during application bypass.

8. Optionally, and if you are configuring an interface set on a 3D6500 or 3D9900 type a maximum frame size for your IP traffic in the **Maximum Frame Size** field. You can set any jumbo frame size between 1518 and 9018 bytes, inclusive.

   On the Defense Center only, a list of sensor groups appears, including a list of ungrouped sensors.

   The following shows a 3D9900 interface set.



9. Select one of the sensors from the list.

   If you are creating an **inline** interface set, a list of network interfaces on the sensor appears.

   If you are creating an **inline with fail open** interface set, a list of paired network interfaces on the sensor's fail-open cards appears.

Keysight, EX-1004
Page 237

10. Add the interfaces to your interface set.
    - If you are creating an **inline** interface set, select two interfaces that you want to designate as an inline pair from the Available Interfaces list and click the arrow button to add the interface to the Selected Interfaces list. Repeat to add additional interface pairs.
    - If you are creating an **inline with fail open** interface set, select at least one interface pair from the Available Interfaces list and click the arrow button to add the interface to the Selected Interfaces list.

    Use the Shift and Ctrl keys to select multiple interfaces or interface pairs at once.

    Determining which interface name corresponds with a physical interface on your sensor depends on the model:
    - For 3Dx800 sensors, the names that appear in the Available Interfaces list correspond to the slot number and interface location. For example, `s0.e0` corresponds to the leftmost interface on the network interface module (NIM) in I/O Slot 0 on the back of your appliance.
    - For 3D Sensor Software for Crossbeam Systems X-Series, the paired interface names that appear in the **Available Interfaces** list correspond to the device names you assigned to the transparent bridge-mode bridge circuits you created on the X-Series. Note that 3D Sensor Software for Crossbeam Systems X-Series does not support inline with fail open interface sets.

    For more information, see the installation guide for your sensor or sensor software.

    You can configure inline interface sets on 3D3800 and 3D5800 sensors to contain up to four pairs of interfaces. Inline with fail open interface sets on 3D3800 and 3D5800 sensors can also contain up to four pairs of interfaces, but each pair must reside on a single fail-open network card. On the 3D9800 sensor, inline and inline with fail open interface sets can include up to the total number of interface pairs on the sensor. On Crossbeam-based software sensors, an inline interface set can include up to eight interface pairs, though Sourcefire recommends you use no more than four.

11. Optionally, for a 3Dx800 or 3Dx900 sensor, select the **Enable Tap Mode** check box to use tap mode.

TIP! 3D9800 sensors with earlier versions of firmware do not support tap mode. The Sourcefire 3D System checks the 3D9800 firmware version and displays the optional tap mode check box in the Create Interface Set page when appropriate.

Keysight, EX-1004
Page 238

12. Optionally, for a 3D3800 or 3D5800 sensor, select **Link State Propagation Mode**.

    This option is especially useful if the routers on your network are able to re-route traffic around a network device that is down.

    TIP! The link lights on fiber fail-open NIMs remain lighted even when the link state is **down** on 3D3800 or 3D5800 sensors with link state propagation enabled.

    Note that link state propagation is **not** available for Crossbeam-based software sensors or 3D9800 sensors.

13. Click **Save**.

    The interface set is created.

    TIP! After you create an interface set, make sure you reapply intrusion policies to the IPS detection engines on the affected sensor.

Editing an Interface Set

Platform: DC/3D Sensor other than 8000 Series sensors

Component: None

In some circumstances, editing an interface set or detection engine can cause the detection engines on the sensor to restart, which can cause a short pause in processing.

IMPORTANT! For most 3D Sensors with inline interface sets, a software bridge is automatically set up to transport packets when the sensor restarts. Although some packets are transmitted without inspection during this time, no packets are lost.

The following sections describe some of the cases where a detection engine is affected by changes to the detection engines and interface sets:

Keysight, EX-1004
Page 239

3Dx800 Sensors

- If you change the number of network interfaces, the interface set type, or transparent mode for an interface set, all the detection engines using that interface set are restarted.
- If you change an interface set's tap mode setting, all detection engines assigned to that interface set are restarted.

> **TIP!** 3D9800 sensors with earlier versions of firmware do not support tap mode. The Sourcefire 3D System checks the 3D9800 firmware version and displays the optional tap mode check box in the Create Interface Set page when appropriate.

- If you change the number of detection resources, which interface set is used, or the detection engine type, only that detection engine is restarted (although other CPUs may be restarted to rebalance the processing load).
- If you create a detection engine, only that detection engine is started (although other CPUs may be restarted to rebalance the processing load).
- If you delete a detection engine or interface set, all detection engines on the sensor are restarted.
- If you create an interface set, nothing is restarted.
- If you change the name or description of an interface set or detection engine, nothing is restarted.

Other Sensors

- If you change which network interfaces are used by the interface set, all the detection engines on the sensor are restarted.
- If you change an interface set's transparent mode setting or interface set type, all detection engines assigned to that interface set are restarted.
- If you change a detection engine's interface set, all detection engines on the sensor are restarted.
- If you change the number of detection resources allocated to a detection engine, all the detection engines on the sensor are restarted.
- If you change the detection engine type for a detection engine, that detection engine is restarted.
- When you create a detection engine, all the detection engines on the sensor are restarted because the total number of allocated resources has changed.
- If you delete a detection engine or interface set, all detection engines on the sensor are restarted.

Keysight, EX-1004
Page 240

- If you create an interface set, nothing is restarted. A restart occurs only when you assign a detection engine to the interface set.
- If you change the name or description of an interface set or detection engine, nothing is restarted.

Make sure you plan these actions for times when they will have the least impact on your deployment.

To edit an interface set:

Access: Admin

1. Select **Operations > Configuration > Detection Engines > Interface Sets**.

   The Interface Sets page appears.

2. Click **Edit** next to the interface set that you want to modify.

   The Create Interface Set page appears.

3. Make any changes to the interface set and click **Update**.

   Your changes are saved.

TIP! After you edit an interface set used by an IPS detection engine, make sure you reapply your intrusion policy on the affected sensor.

Editing an 8000 Series Interface Set

Platform: DC/8000 Series 3D Sensor

Component: None

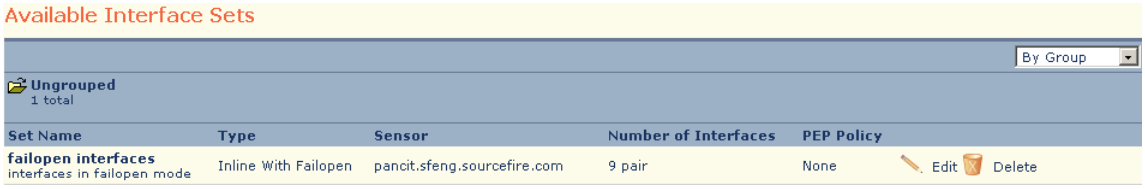
An interface set is a collection of one or more sensing interfaces on your appliance. For information about their use, see Using Interface Sets on page 222.

To edit an 8000 Series interface set:

Access: Admin

1. On the Defense Center that manages the sensor, select **Operations > Configuration > Interface Sets**.

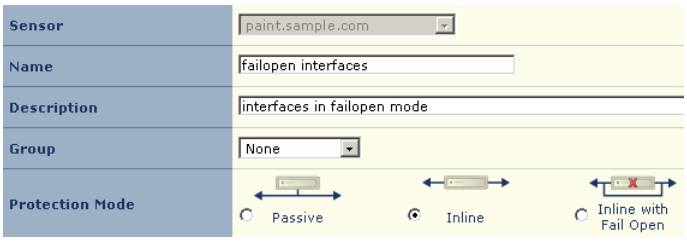
   The Interface Sets page appears.



Keysight, EX-1004
Page 241

2. Click **Edit** next to the interface set you want to modify.

   The Create Interface Set page appears.

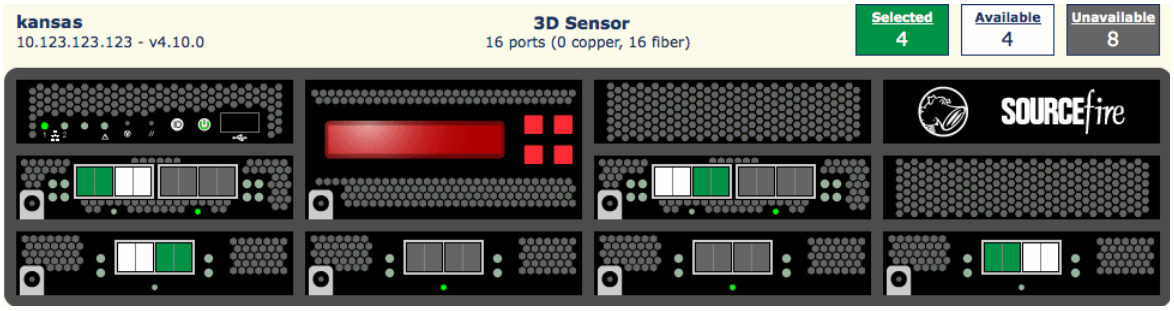


3. Type a name and description for the new interface set using alphanumeric characters and spaces.

   Optionally, select an existing interface set group or **New Group** from the drop-down menu. See Using Interface Set Groups on page 244 for more information.

4. For **Protection Mode**, select passive, inline, or inline with fail-open. See Types of Interface Sets on page 225 for more information.

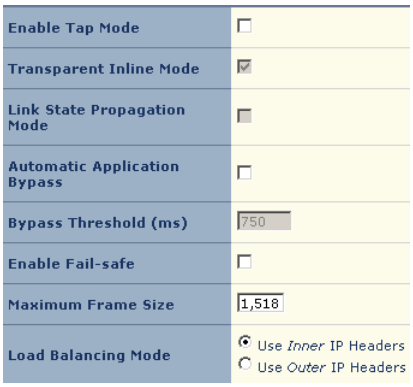
5. Click an interface set to select or deselect. The interface set color indicates whether it is selected, available, or unavailable.



   Place your cursor over an interface set to display the interface set name, type (for example, LR or SR fiber), link, speed, and bypass.

Keysight, EX-1004
Page 242

6. Optionally, expand the **Advanced** option to display Advanced options.



7. Optionally, select the **Enable Tap Mode** check box to use tap mode.
8. Optionally, if you selected the **Inline** or **Inline with Fail Open** option, clear the **Transparent Inline Mode** check box to disable transparent mode.
9. If you selected either the Inline or Inline with Fail Open option, then, optionally, select **Link State Propagation Mode**. This option is especially useful if the routers on your network are able to re-route traffic around a network device that is down.
10. Optionally, select **Automatic Application Bypass** if your network is sensitive to latency.
11. If you selected **Automatic Application Bypass**, you can set a **Bypass Threshold** in milliseconds (ms). The default setting is 750 ms and the valid range is from 250 ms to 60,000 ms. Automatic Application Bypass is most useful in inline applications.
12. Optionally, you can select the **Enable Fail-safe** check box to enable traffic pass-through during application bypass.
13. Optionally, you can configure a maximum frame size for your IP traffic in the **Maximum Frame Size** field. You can set any jumbo frame size between 1518 and 9018 bytes, inclusive.
14. Optionally, you can set the **Load Balancing Mode** to use inner IP headers (default) or outer headers. See Load Balancing on page 230.
15. Click **Save**.

    The interface set is saved.

TIP! After you create an interface set, make sure you reapply intrusion policies to the IPS detection engines on the affected sensor.

Keysight, EX-1004
Page 243

Deleting an Interface Set

Platform: DC/3D Sensor

Component: None

You **cannot** delete an interface set that is being used by a detection engine. You must delete the detection engine before you can delete the interface set.

Note also that you cannot delete an interface set on a clustered 3D9900 sensor. See Using Clustered 3D Sensors on page 250 for more information.

To delete an interface set:

Access: Admin

1. Select **Operations > Configuration > Interface Sets**.

   The Interface Sets page appears.

2. Click **Delete** next to the interface set that you want to delete, and, at the prompt, confirm that you want to delete the interface set.

   The interface set is deleted.

Using Interface Set Groups

Platform: DC

Component: None

You can use interface set groups to combine similar interface sets. These groups make it easier to apply PEP policies to interface sets that have similar purposes. For more information on PEP policies, see Understanding PEP Traffic Management on page 1511.

See the following sections for more information:

- Creating Interface Set Groups on page 244
- Editing Interface Set Groups on page 245
- Deleting Interface Set Groups on page 246

Creating Interface Set Groups

Platform: DC

Component: None

The following procedure explains how to create an interface set group.

Keysight, EX-1004
Page 244

To create a interface set group:

Access: Admin

1. Select **Operations > Configuration > Interface Sets**.
2. Click **Create Interface Set Group** or click **Create Interface Set** then click **Create New Group** in the **Group** field.

   The Create Interface Set Group page appears.



   Type a name for the interface set group in the **Group Name** field.

3. Click **Save**.

   The Interface Set page appears again.

   You can add interface sets to an interface set group by clicking **Edit** next to a interface set group name and, on the Interface Group Edit page, adding available interfaces to the group and clicking **Save**.

Editing Interface Set Groups

Platform: DC

Component: None

The following procedure explains how to edit an interface set group. You must create an interface set group before you can edit it. See Creating Interface Set Groups on page 244.

To edit an interface set group:

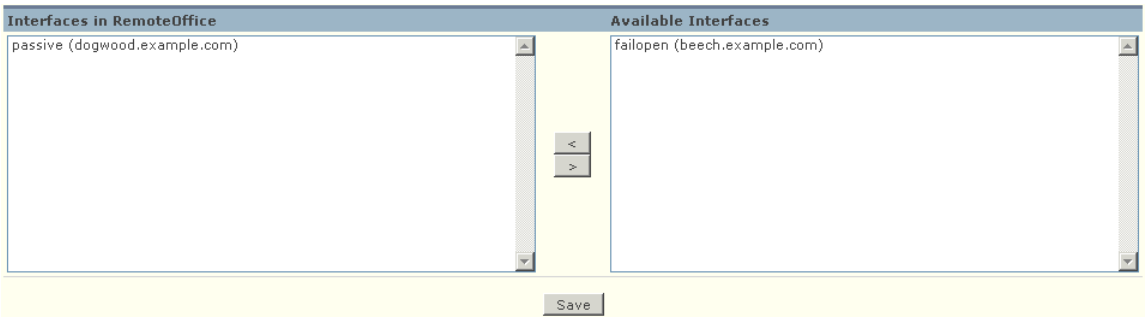
Access: Admin

1. Select **Operations > Configuration > Interface Sets**.

   The Available Interface Sets page appears.

2. Click **Edit** for the interface set group.

   The Interface Group Edit page appears.



Keysight, EX-1004
Page 245

3. Select available interface sets and to move them to the interface set group with the arrow buttons.

   You can also move interface sets out of the interface set group.

4. Click **Save** to add the selected interfaces to the interface set group.

   The Available Interface Sets page appears.

Deleting Interface Set Groups

Platform: DC

Component: None

When you delete an interface set group, any interface sets in the group are automatically ungrouped; they are not deleted.

To delete a interface set group:

Access: Admin

1. Select **Operations > Configuration > Interface Sets**.

   The Interface Sets page appears.

2. Click **Delete** next to the name of the interface set group.

   The interface set group is deleted.

Inline Fail Open Interface Set Commands

Platform: 3D Sensor

Component: None

When you use fiber inline fail open interfaces sets and the interface set goes into bypass, you can force the interface out of bypass mode. See Removing Bypass Mode on Inline Fail Open Fiber Interfaces. You can force a copper or fiber inline fail open interface in or out of bypass. See Forcing an Inline Fail Open Interface Set into Bypass Mode on page 247.

Removing Bypass Mode on Inline Fail Open Fiber Interfaces

Platform: 3D Sensor

Component: None

When link state propagation is enabled on a sensor with an inline fail open interface set and the sensor goes into bypass mode, all network traffic passes through the interface pair without being analyzed. When the links restore, most **fiber** inline fail open interface sets do not return from bypass automatically. You can use a command line tool to force the interface set out of bypass mode.

Keysight, EX-1004
Page 246

This tool works on most 3D Sensors with inline with fail open **fiber** interface pairs. It is not necessary to use this tool on inline with fail open **copper** interface pairs or to use this tool with 3D9900 sensors.

IMPORTANT! Make sure you contact Technical Support if you are having issues with the fail open interfaces on your sensor.

To force a fiber inline fail open interface set out of bypass mode:

Access: Admin

1. Open a terminal window on your 3D Sensor and enter the command `su` and the `admin` password to switch to the admin user.
2. Enter the following at the command line:
   ```
   su /var/sf/bin/unbypass_cards.sh
   ```
   You will be prompted for your password.
3. When the interfaces switch out of bypass mode, a message in syslog indicates the 3D Sensor is analyzing traffic. For example:
   ```
   Fiber pair has been reset by un_bypass
   ```

To force a fiber inline fail open interface set out of bypass mode on an 8000 Series sensor:

Access: Admin

1. Open a terminal window on your 8000 Series 3D Sensor and enter the command `su` and the `admin` password to switch to the admin user.
2. Enter `expert` at the prompt.
3. Enter the following at the command line:
   ```
   su /var/sf/bin/unbypass_cards.sh
   ```
   You will be prompted for your password.
4. When the interfaces switch out of bypass mode, a message in syslog indicates the 3D Sensor is analyzing traffic. For example:
   ```
   Fiber pair has been reset by un_bypass
   ```

Forcing an Inline Fail Open Interface Set into Bypass Mode

Platform: 3D Sensor

Component: None

When the sensor with an inline fail open interface set fails, it goes into bypass mode, a state where all network traffic passes through the interface pair without being analyzed. If you are troubleshooting an interface set, or if the interface card does not fail open on its own, you can use a command line tool to force the interface set into bypass mode. Note that the procedure is slightly different for an 8000 Series sensor.

Keysight, EX-1004
Page 247

Note that this tool works only with inline with fail open interface pairs. You **cannot** use it with non-fail open inline interface sets.

To force an inline fail open interface set into bypass mode, you must know which two interfaces are included in the interface set. You can determine this information on the Interface Sets page.

IMPORTANT! Make sure you contact Technical Support if you are having issues with the fail open interfaces on your sensor.

To force an inline fail open interface set into bypass mode:

Access: Admin

1. On the appliance's web interface, select **Operations > Configuration > Interface Sets**.

   The Interface Sets page appears.

2. Under Available Interface Sets, click **Edit** next to the inline with fail open interface set you are investigating.

   The Create Interface Set page appears. Place your cursor over the interfaces to display the names of the interfaces in the interface set.

3. Using the command line interface, log into the admin account on the sensor.

4. Enter the following at the command line:

   ```
   su failopen_pair.pl open eth#:eth#
   ```

   For example, if the interfaces in the interface set are `eth2` and `eth3`, enter the following:

   ```
   failopen_pair.pl open eth2:eth3
   ```

   The following message appears:

   ```
   NOTE:  You must already have a failopen interface set
          and detection engine configured on the pair
          you are forcing open or closed for this utility
          to work.
   ```

   Then, if you specified the correct interfaces, the following message appears:

   ```
   Mode changed for interfaces eth2:eth3
   ```

   The interfaces switch to bypass mode and the traffic is no longer analyzed.

   If you did not specify the correct interfaces, the following message appears:

   ```
   No failopen interface set configured for interfaces
   eth2:eth3...
   ```


Keysight, EX-1004
Page 248

To force an inline fail open interface set into bypass mode on an 8000 Series sensor:

Access: Admin

1. On the appliance's web interface, select **Operations > Configuration > Interface Sets**.

   The Interface Sets page appears.

2. Under Available Interface Sets, click **Edit** next to the inline with fail open interface set you are investigating.

   The Create Interface Set page appears. Place your cursor over the interface set to display the names of the interfaces in the interface set.

3. Use the command line interface on the sensor, and log into the `admin` account.

4. Enter `expert` at the prompt.

5. Enter the following at the command line:

   ```
   su failopen_pair.pl open eth#:eth#
   ```

   For example, if the interfaces in the interface set are `eth2` and `eth3`, enter the following:

   ```
   failopen_pair.pl open eth2:eth3
   ```

   The following message appears:

   ```
   NOTE:  You must already have a failopen interface set
          and detection engine configured on the pair
          you are forcing open or closed for this utility
          to work.
   ```

   Then, if you specified the correct interfaces, the following message appears:

   ```
   Mode changed for interfaces eth2:eth3
   ```

   The interfaces switch to bypass mode and the traffic is no longer analyzed.

   If you did not specify the correct interfaces, the following message appears:

   ```
   No failopen interface set configured for interfaces
   eth2:eth3...
   ```


Keysight, EX-1004
Page 249

To return an inline fail open interface set to normal mode:

Access: Admin

1. Log in as admin onto the sensor and, at the prompt, enter the correct password.
2. Enter the following at the command line:

   ```
   failopen_pair.pl close eth#:eth#
   ```

   For example, if the interfaces in the interface set are `eth2` and `eth3`, enter the following:

   ```
   failopen_pair.pl close eth2:eth3
   ```

   The following message appears:

   ```
   Mode changed for interfaces eth2:eth3
   ```

   The interfaces return to normal mode and the traffic flowing through the detection engines on the interface set is analyzed as you would expect.

To return an inline fail open interface set to normal mode on an 8000 Series sensor:

Access: Admin

1. Log in as admin onto the sensor and, at the prompt, enter the correct password.
2. Enter `expert` at the prompt.
3. Enter the following at the command line:

   ```
   failopen_pair.pl close eth#:eth#
   ```

   For example, if the interfaces in the interface set are `eth2` and `eth3`, enter the following:

   ```
   failopen_pair.pl close eth2:eth3
   ```

   The following message appears:

   ```
   Mode changed for interfaces eth2:eth3
   ```

   The interfaces return to normal mode and the traffic flowing through the detection engines on the interface set is analyzed as you would expect.

Using Clustered 3D Sensors

Platform: DC + 3D9900s or 3D8250s

Component: None

When you connect two fiber- based 3D9900s or two 3D8250s in a clustered pair configuration, you combine the detection resources of both sensors into a single, shared configuration and increase the amount of traffic inspected on a network segment.

You must use a Defense Center to manage the clustered pair, rather than the individual sensor's web interface. Use the Defense Center to create and edit the

Keysight, EX-1004
Page 250

detection engines of clustered 3D Sensors. You cannot manage detection engines on the local web interface of a clustered 3D Sensor; the Edit page is replaced with an informational page. You can use the detection engines configured on a clustered pair in:

- IPS policies
- PEP policies
- RNA detection policies
- compliance rules

Although a clustered pair is managed as a single unit, Analysis & Reporting tools display information from each half of the detection engine independently. Add data from both sensors of the detection engine to measure the total.

A sensor in a clustered pair collects and reports information on:

- intrusion event statistics
- intrusion events
- event graphs
- dashboards
- RNA statistics
- network map
- searches

IMPORTANT! If you use eStreamer to stream event data from a clustered pair of sensors to an external client application, collect the data from both sensors and ensure that you configure each sensor identically. The eStreamer settings are not automatically synchronized over the pair.

For information on clustered pairs, including establishing and creating clustered pairs, see Managing a Clustered Pair on page 146.

For information about creating, editing, and deleting detection engines, see Managing Detection Engines on page 206.

Identifying Detection Engines and Interface Sets on a Clustered Pair

The Defense Center displays and manages the clustered pair as a single interface set of the master sensor.

You can identify a clustered pair on the Sensor list page. Select **Operations > Sensors** and note that clustered sensors display a peer icon (). Place your cursor over the icon to display whether the sensor is the master or the slave.

When you create a detection engine for a clustered pair, both sensors are listed in the interface set. The format is *DetectionEngineName*(*MasterSensorName,*

Keysight, EX-1004
Page 251

SlaveSensorName). For example, a clustered 3D Sensors detection engine could be: `Sample inline DE (birch.example.com, fir.example.com)`.

- `Sample inline DE` is the name of the detection engine.
- `birch.example.com` is the name of the master in the pair.
- `fir.example.com` is the name of the slave in the pair of sensors.

The following graphic displays the detection engine of the clustered pair as it appears in the list of detection engines on the Defense Center.



When you create or edit a detection engine formed by the clustered 3D Sensors, the detection resources are listed as from both sensors. The following graphic displays the detection engine resources from the clustered pair.



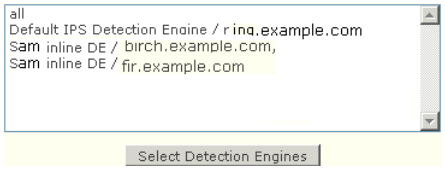
The detection engine is displayed each sensor in a clustered pair is identified as:

```
DetectionEngineName/MasterSensorName
DetectionEngineName/SlaveSensorName
```

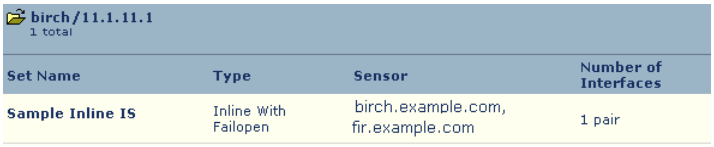
Following the `Sample inline DE (birch.example.com, fir.example.com)` example, information collected from the sensors is attributed to each sensor:

```
Sample inline DE / birch.example.com
Sample inline DE / fir.example.com
```

An example of a list (this one from the Intrusion Event Statistics page) is shown below.



The following example displays a clustered pair interface set. Note that both the master and the slave in the **Sensor** column.



Do **not** attempt to change the interface settings on a sensor in a clustered pair; results are unpredictable.

Keysight, EX-1004
Page 252

Chapter 6
Introduction to Sourcefire IPS

You can deploy your Sourcefire 3D System either inline or passively on your network to help you monitor for 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 IPS component of the Sourcefire 3D System can perform as both an intrusion detection system and an intrusion prevention system depending on:

- how you attach 3D Sensors with IPS to your network: inline or out of band
- how you configure the sensors' interface sets: passive or inline
- the drop behavior of rules set to Drop and Generate Events in an inline deployment: enabled or disabled

After you deploy your 3D Sensors with IPS and configure them according to your needs, IPS has several mechanisms it uses to look for the broad range of exploits that attackers have developed. You can then use the broad range of tools in the Sourcefire 3D System to analyze and respond to the intrusion events that are generated by IPS.

IMPORTANT! All of the intrusion detection and prevention features in the Sourcefire 3D System require that you deploy 3D Sensors with the IPS component. IPS adds a full web interface to the 3D Sensor, but you can gain additional insight into your network and the attacks on it by managing your 3D Sensors with a Defense Center.

Keysight, EX-1004
Page 253

Sensing Intrusions

3D Sensors use *detection engines,* which are a collection of one or more sensing interfaces on a sensor plus a portion of the sensor's computing resources, to analyze network traffic for evidence of attacks on network resources. Each detection engine on the 3D Sensor has several methods that it uses to look for a broad range of exploits. First, 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, the detection engine uses intrusion rules to 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.

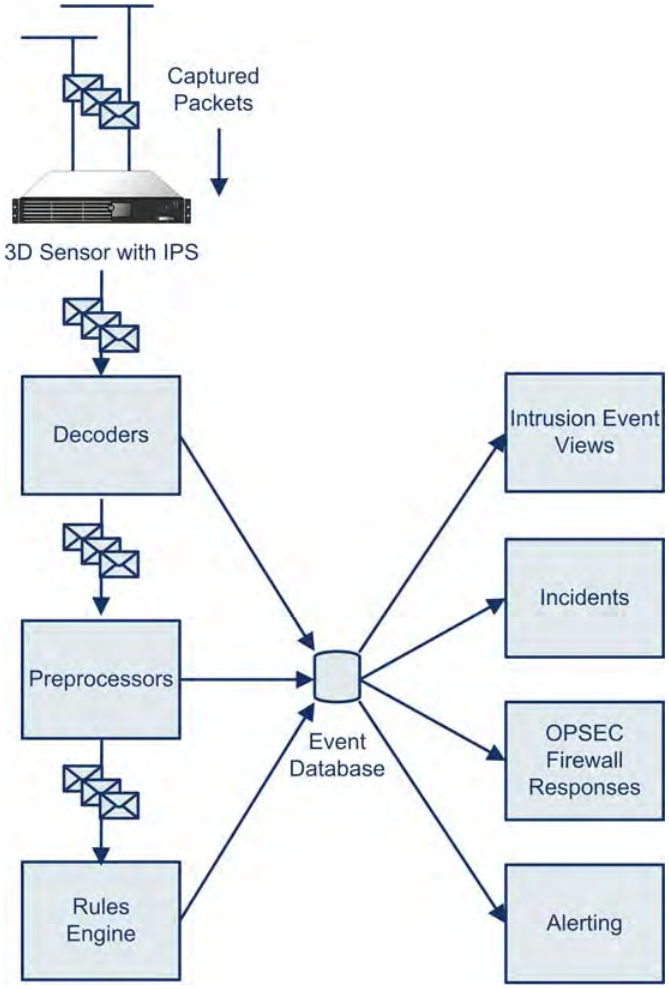
The Sourcefire Vulnerability Research Team (VRT) regularly sends out updates called Security Enhancement Updates, or SEUs, that can contain new intrusion rules, decoders, and preprocessors and other intrusion detection and prevention features, so you can be sure that you are detecting the most recently released attacks.

Responding to Intrusions

When a packet travels over a segment monitored by a detection engine, the IPS detection engine captures and analyzes it using a series of decoders and preprocessors and then a rules engine. When a detection engine 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. If the IPS detection engine is deployed inline and traffic flows through a pair of interfaces on the sensor, the detection engine 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. You can also generate alerts in response to events, or launch a remediation using a firewall. In a Sourcefire 3D System deployment that includes 3D Sensors with IPS and a Defense Center, the sensors transmit their events to the Defense Center where

Keysight, EX-1004
Page 254

you can view the aggregated data and gain a greater understanding of the attacks against your network assets.



To learn more about IPS and how a Sourcefire 3D System deployment that includes IPS can help protect your network, see the following sections:

- Understanding How Traffic Is Analyzed on page 256
- Analyzing Intrusion Event Data on page 261
- Using Intrusion Event Responses on page 261
- Understanding IPS Deployments on page 262
- Using IPS Software Sensors on page 265
- The Benefits of Custom Intrusion Policies on page 265

Keysight, EX-1004
Page 255

Understanding How Traffic Is Analyzed

Platform: DC/MDC/Standalone Sensor

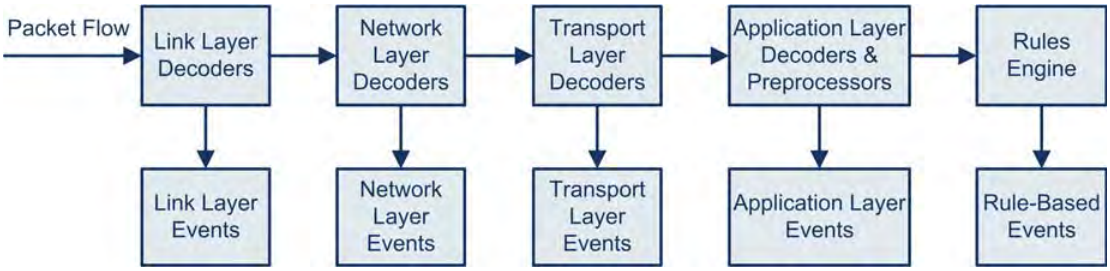
Component: IPS

3D Sensors with IPS use 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 a detection engine on the sensor 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

The following diagram illustrates the process:



Events include:

- the date and time the event was generated
- the event priority
- Requires: DC + IPS + RNA the impact flag associated with the event
- whether the packet that caused the event was dropped or would have been dropped in an inline deployment
- the name of the detection engine 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 name of the user logged into the destination host
- the ICMP type and code (for ICMP traffic)

Keysight, EX-1004
Page 256

- 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 details on the basic information included in an intrusion event, see Understanding Intrusion Events on page 276.

IMPORTANT! For events generated by shared object rules, the rule itself is not available.

The following sections describe more about how IPS acquires and processes information:

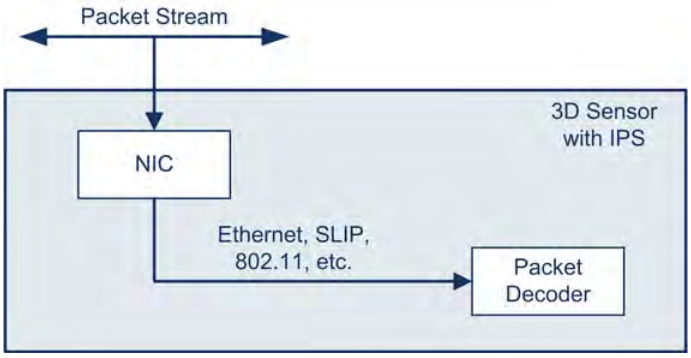
- Capturing and Decoding Packets on page 257
- Processing Packets on page 258
- Generating Events on page 260

Capturing and Decoding Packets

Platform: DC/MDC/Standalone Sensor

Component: IPS

Before packets can be inspected by the IPS detection engine, the packets must be captured from the network. The following illustration shows how the detection engine sniffs packets, then decodes them prior to any further analysis.



As the detection engine 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

Keysight, EX-1004
Page 257

through the network and transport layers, as described in the Decoded Packets table.

Decoded Packets

In this TCP/IP layer...	These packets are decoded...
Data Link	• raw packets • loopback interface • Ethernet • Token Ring • Fiber Distributed Data Interface (FDDI) • Linux SLL (cooked mode) • IEEE 802.11 • Point-to-Point Protocol (PPP) • Point-to-Point Protocol over Ethernet (PPPOE) • Serial Line Internet Protocol (SLIP) • ISDN for Linux (I4L) • Address Resolution Protocol (ARP) • Extensible Authentication Protocol (EAP) • EAP over LAN (EAPOL)
Network	• Internet Protocol (IP) • Internet Control Message Protocol (ICMP)
Transport	• Transmission Control Protocol (TCP) • User Datagram Protocol (UDP)

Processing Packets

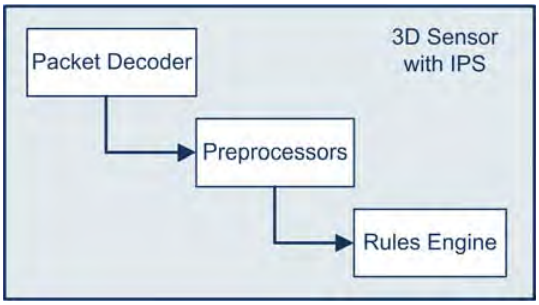
Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 258

and payloads to determine whether they trigger any of the 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 477.

When you apply an intrusion policy to the detection engines, and before any packets are processed, the rule optimizer reads in all the activated rules and then classifies them 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 event intrusion statistics and intrusion event reports.

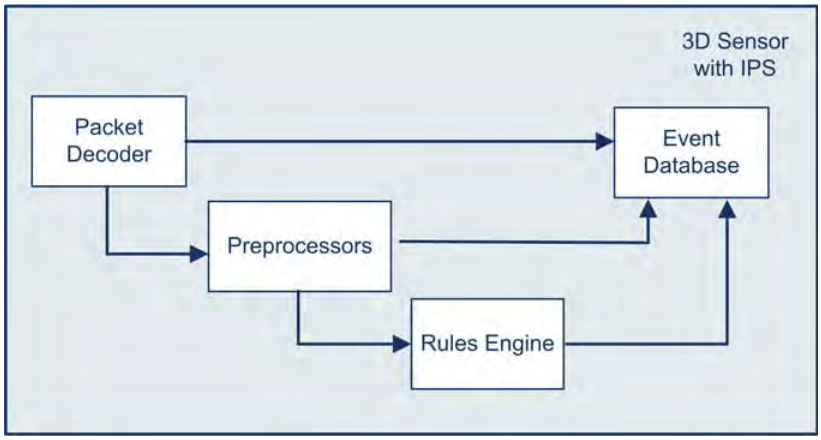
Keysight, EX-1004
Page 259

Generating Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

Packets are evaluated by the packet decoder, the preprocessors, and the rules engine. At each step of the process, a packet could cause Sourcefire 3D 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 types 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 and 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.

Keysight, EX-1004
Page 260

Analyzing Intrusion Event Data

Platform: DC/MDC/Standalone Sensor

Component: IPS

For more information, see Working with Event Reports on page 1531. As IPS accumulates intrusion events, you can begin your analysis of potential attacks. The Sourcefire 3D System provides you with the tools you need to review the 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 3D Sensor

  For more information, see Viewing Intrusion Event Statistics on page 269.
- 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 Event Reports on page 1531.
- 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 319.
- 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 1573 and Working with Intrusion Events on page 267.

Using Intrusion Event Responses

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 734.

Keysight, EX-1004
Page 261

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 734.
- automated firewall response that you can configure

  For more information, see Using Check Point OPSEC Responses on page 735.
- on a Defense Center, automated compliance policies that you can use to respond to and remediate specific intrusion events

  For more information, see Configuring Responses for Compliance Policies on page 1281.

Understanding IPS Deployments

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can assign a passive interface set to an IPS detection engine so that the detection engine senses traffic out of band from the packet stream. Similarly, you can assign an inline interface set to an IPS detection engine so that the traffic flows between a pair of sensing interfaces on the sensor. An inline interface set on an IPS detection engine with an intrusion policy set to drop packets allows you to drop or replace packets that you know to be harmful.

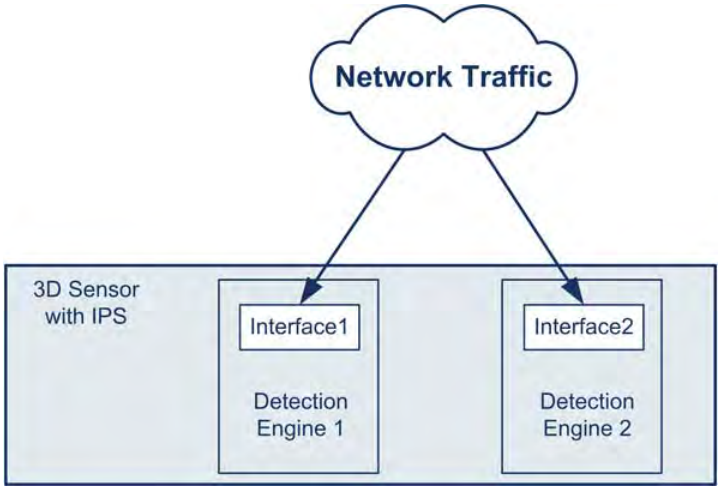
You can tailor intrusion policies for each IPS detection engine so that it generates events only for the attacks that are likely to affect the security of the hosts on specific portions of your network. You can specify which rules do not alert, which rules generate events and, for an inline deployment, 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 3D Sensor 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 detection engine 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 deployment lies in the interface sets used by each system. An out-of-band deployment uses a passive interface set; an inline deployment uses an inline or inline with fail open interface 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 or inline with fail open interface set.

Keysight, EX-1004
Page 262

The following illustration shows an example of a 3D Sensor with IPS deployed passively and configured with two detection engines using passive interface sets. Each IPS detection engine is using a different network interface on the sensor and is monitoring a different network segment.

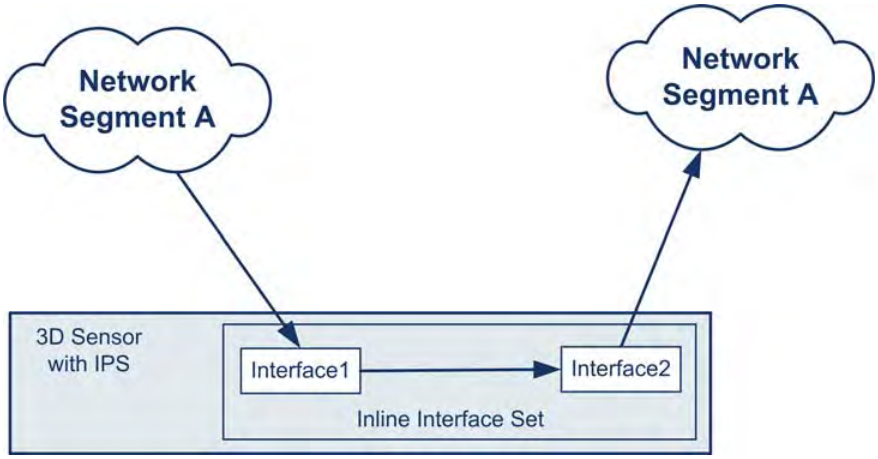


Configuring an out-of-band IPS detection engine allows you to devote almost all of the detection engine's bandwidth and computational power to monitoring traffic, reconstructing datagrams and streams, normalizing packets, detecting anomalies, and alerting you to possible intrusions. Moreover, because the detection engine 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 an IPS detection engine that uses an inline interface set. To do this, you connect the 3D Sensor with IPS to your network so that traffic flows between the network interfaces that are assigned to your detection engine. When the interface set is configured as Inline or Inline with Fail Open, 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 detection engine 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.

Keysight, EX-1004
Page 263

For example, the following illustration shows an IPS deployed inline. The detection engine uses an interface set containing two of the network interfaces on the back of the sensor and is monitoring a single network segment.



Similar to a detection engine using a passive interface set, the detection engine 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 detection engine detects a packet whose payload contains a known exploit to which your network may be vulnerable, you can configure the engine to drop the packet. In this case, the malicious packet never reaches its intended target.

When using IPS 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 detection engine detects a packet that contains `bin/sh`, which can often indicate 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 an IPS, you should weigh some of the trade-offs. First, if you deploy an IPS, you must choose a 3D Sensor 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 sensor with the optional *fail-open* network card. The fail-open 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 detection engines and

Keysight, EX-1004
Page 264

interface sets, see Using Detection Engines and Interface Sets on page 198. You can learn more about deployment options in your sensor's *Installation Guide*.

Using IPS Software Sensors

Platform: DC/MDC/Standalone Sensor

Component: IPS

Sourcefire offers Intrusion Agents that you can deploy throughout your network environment, as well as 3D Sensor Software with IPS for Crossbeam Systems X-Series.

Because they do not have a web interface, your Sourcefire 3D System deployment must include at least one Defense Center if you want to deploy Intrusion Agents or Crossbeam-based software sensors.

The Benefits of Custom Intrusion Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS 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* the 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:

- Deactivate 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.
- Determine whether you want the preprocessor to generate an event when it acts on a packet.
- Configure preprocessors to generate events when they encounter certain features in packets, for example, state problems or certain combinations of TCP flags.
- If your Sourcefire 3D System includes RNA resources, onfigure adaptive profiles to use information about host operating systems from the RNA network 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 477.

Keysight, EX-1004
Page 265

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 detection engines, specify which rules should direct the detection engine to drop malicious packets from the packet stream.

> **TIP!** If your Sourcefire 3D System deployment includes IPS and RNA, you can use RNA 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 that 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 chose 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 297, Setting Suppression Options within the Packet View on page 298, and Filtering Intrusion Event Notification Per Policy on page 439.

Specifying the protocol analysis, data normalization, and traffic inspection performed by IPS and saving this configuration as a whole allows you to control the kind of information IPS 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 761.

Keysight, EX-1004
Page 266

Chapter 7
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 3D Sensors that are licensed for the IPS component on key network segments, you can examine the packets that traverse your network for malicious activity. IPS has several mechanisms it uses to look for the broad range of exploits that attackers have developed.

When IPS 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. In a Sourcefire 3D System deployment that includes a Defense Center managing 3D Sensors with IPS, the sensors 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 IPS as an inline intrusion system, which allows you to configure the sensor 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 the events that IPS generates 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 3D Sensors with IPS
- 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

Keysight, EX-1004
Page 267

- 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
- on a Defense Center that manages 3D Sensors with IPS and RNA, automated compliance 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 269 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 Graphs on page 273 explains how to generate charts that show event trends over time.
- Viewing Intrusion Events on page 274 describes how to use the web interface to view and investigate your intrusion events.
- Understanding Workflow Pages for Intrusion Events on page 280 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 284 describes the features of two of the types of pages in an intrusion event workflow.
- Using the Packet View on page 290 explains how to use the packet view of intrusion events.
- Using Impact Flags to Evaluate Events on page 307 describes how you can use the impact flag to evaluate intrusion events if your Defense Center manages 3D Sensors with RNA and IPS deployed on the same network segment.
- Searching for Intrusion Events on page 309 explains how you can use the search feature to constrain a list of intrusion events to specific criteria.
- Using the Clipboard on page 316 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 319 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 734 for more information about automated alerting.
- Working with Event Reports on page 1531 for more information about intrusion event reports.

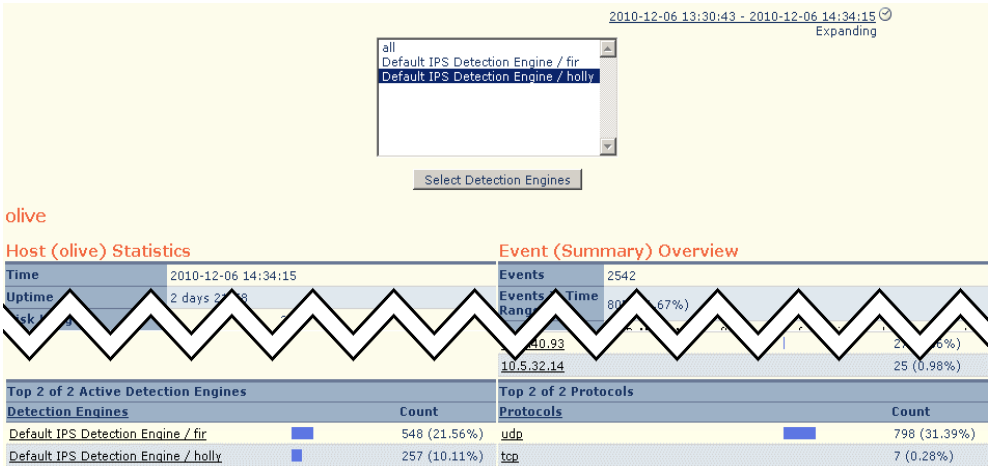
Keysight, EX-1004
Page 268

Viewing Intrusion Event Statistics

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Intrusion Event Statistics page provides you with a quick summary of the current state of your appliance and any intrusion events generated by IPS for your network.



The Intrusion Event Statistics page has three main areas:

- Host Statistics on page 271 describes the Host Statistics section, which provides information about the appliance and, for Defense Centers, their managed 3D Sensors with IPS.
- Event Overview on page 271 describes the Event Overview, which provides an overview of the information in the event database.
- Event Statistics on page 272 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, and event messages 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 detection engines 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

Keysight, EX-1004
Page 269

click the First Event link, the resulting event pages will not show the event until you change the time range.

TIP! If your 3D Sensor is configured to store events only on the Defense Center, then you must use the Defense Center to view the intrusion events.

To view intrusion event statistics:

Access: Any IPS/Admin

1. Select **Analysis & Reporting > Event Summary > Intrusion Event Statistics**.

   The Intrusion Event Statistics page appears.

2. From the list of detection engines, select the detection engine whose statistics you want to view, or select **all** to view statistics for all the detection engines that are collecting intrusion events.

   TIP! On the Master Defense Center, the list includes all the detection engines from the sensors that are managed by Defense Centers that are, in turn, managed by the Master Defense Center.

3. Click **Select Detection Engines**.

   The Intrusion Event Statistics page refreshes with data from the detection engines you selected. On the Defense Center, note that if you selected a single detection engine, the Event Overview section for that detection engine is repeated.

   TIP! To view data from a custom time range, click the link next to **Events in Time Range** and follow the directions in Setting Event Time Constraints on page 1597.

4. See the following sections for more information about the statistics that appear on the Intrusion Event Statistics page:
   - Host Statistics on page 271
   - Event Overview on page 271
   - Event Statistics on page 272

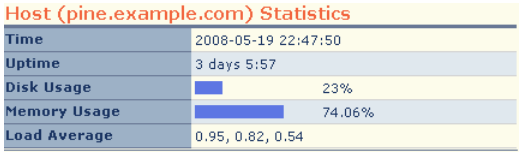
Keysight, EX-1004
Page 270

Host Statistics

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 3D Sensors with IPS.



This information includes the following:

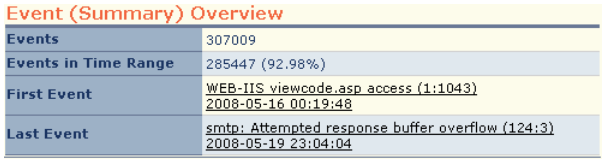
- **Time** shows the current time on the appliance.
- **Uptime** shows the number of days, hours, and minutes since the appliance itself was restarted. On the Defense Center, the uptime also shows the last time each managed sensor was rebooted, the number of users logged in, and the load average.
- **Disk 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Event Overview section of the Intrusion Event Statistics page provides an overview of the information in the intrusion event database.



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.

Keysight, EX-1004
Page 271

- **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.

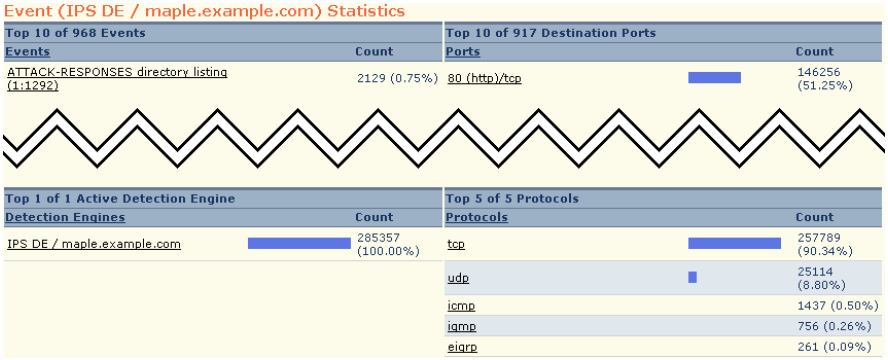
IMPORTANT! On the Defense Center, note that if you selected a single detection engine, the Event Overview section for that detection engine appears instead.

Event Statistics

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Event Statistics section of the Intrusion Event Statistics page provides more specific information about of 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 detection engines and protocols with the greatest number of events

Keysight, EX-1004
Page 272

Viewing Intrusion Event Graphs

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Sourcefire 3D System provides graphs that show you intrusion event trends over time. You can generate intrusion event graphs over time periods ranging from the last hour to the last month, for the following:

- one or all detection engines
- top 10 destination ports
- top 10 source IP addresses
- top 10 event messages

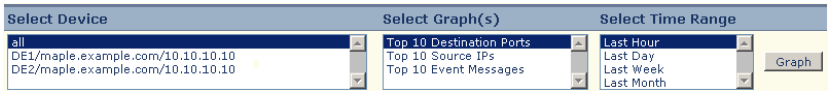
TIP! You can also view performance statistics such as the number of alerts per second, number of megabits per second, average number of bytes per packet, and the percent of packets dropped for your detection engines. See Viewing IPS Performance Statistics on page 1883 for more information.

To generate an event graph:

Access: Any IPS/Admin

1. Select **Analysis & Reporting > Event Summary > Event Graphs**.

   The Event Graphs page appears. Three selection boxes at the top of the page control which graph is generated. Note that the following graphic shows the Defense Center version of the page.



2. Under **Select Device**, select **all** to include all detection engines or select the specific detection engine you want to include in the graph.

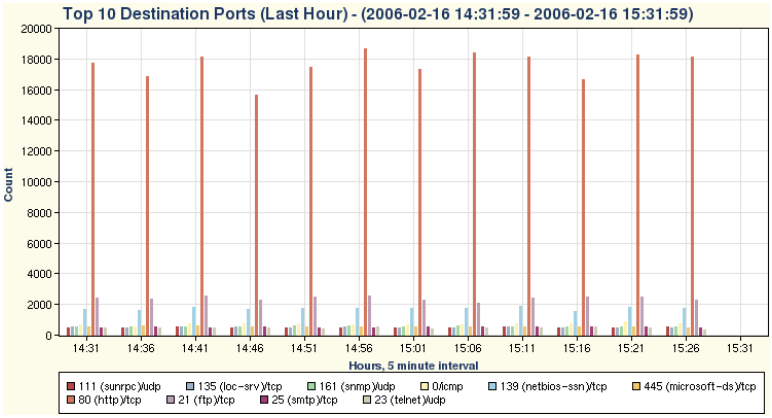
   IMPORTANT! On the Master Defense Center, the device list includes all the detection engines from the sensors that are managed by Defense Centers that are, in turn, managed by the Master Defense Center.

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.

Keysight, EX-1004
Page 273

5. Click **Graph**.

   The graph is generated.



Viewing Intrusion Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

When a detection engine on your IPS recognizes a packet that is potentially malicious, it generates an intrusion event and adds the event to the database.

TIP! If your 3D Sensor is configured to store events only on the Defense Center, then you must use the Defense Center to view intrusion events.

The packet view includes 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.

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.

The initial intrusion events view differs depending on the workflow you use to access the page. On your sensor, 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. On the Defense Center, you can view workflows based on custom tables, which may include intrusion events. Note that an event view can 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 Event Preferences on page 46 for more information.

Keysight, EX-1004
Page 274

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 278 for more information.

If you perform a backup and 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.

For more information, see the following sections:

- Understanding Intrusion Events on page 276
- Creating Custom Workflows on page 1618
- Using Drill-Down and Table View Pages on page 284
- Using the Packet View on page 290
- Reviewing Intrusion Events on page 278
- Viewing a Workflow Based on a Custom Table on page 1569

To view intrusion events:

Access: Any IPS/Admin

▶ Select **Analysis & Reporting > IPS > Intrusion Events**.

The first page of the default intrusion events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

TIP! If you are using a custom workflow that does not include the table view of intrusion events, from the **Workflows** menu on the toolbar, select any of the workflows that ship with the appliance.

See Understanding Intrusion Events on page 276 to learn more about the events that appear in intrusion event views. See Understanding Workflow Pages for Intrusion Events on page 280 to learn more about how to narrow your view to the intrusion events that are important to your analysis.

Keysight, EX-1004
Page 275

Understanding Intrusion Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS 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 IPS 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 describes the information that an intrusion event contains.

Time

The date and time of the event.

Priority

The event priority as determined by the Sourcefire VRT.

Impact Flag

On a Defense Center managing a 3D Sensor with IPS, the impact flag in this field indicates the correlation between IPS data, RNA data, and vulnerability information. For more information, see Using Impact Flags to Evaluate Events on page 307.

Note that intrusion events are qualified with impact flags only if you deploy sensors with IPS and RNA on the same network segment, and then view those events on a Defense Center that manages those sensors.

Note also 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 red (Vulnerable) 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 IPS dropped the packet that triggered the rule
- a gray down arrow, indicating that IPS 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

Keysight, EX-1004
Page 276

Note that IPS 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 435, Setting Drop Behavior in an Inline Deployment on page 358, and Tap Mode on page 227.

Detection Engine

The name of the detection engine that generated the event.

Note that the master and slave in a clustered pair report intrusion events as if they were separate detection engines. See Identifying Detection Engines and Interface Sets on a Clustered Pair on page 251 for more information.

Protocol

The protocol used by the event.

Source IP

The IP address of the sending host.

Destination IP

The IP address of the receiving host.

Original Client IP

The original client IP address that was extracted from the X-Forwarded-For (XFF) or True-Client-IP HTTP headers. To display a value for this field, you must enable the HTTP Inspect preprocessor **Extract Original Client IP Address** option.

Note that the Original Client IP field is disabled by default. To enable it for the duration of your session, click the Expand arrow ( ) to expand the search constraints, then click the column name under Disabled Columns.

See Selecting Server-Level HTTP Normalization Options on page 564 for more information.

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.

Source Port/ICMP Type

The port number on the sending host. For ICMP traffic, the ICMP type appears instead.

Keysight, EX-1004
Page 277

Destination Port/ICMP Code

The port number for the host receiving the traffic. For ICMP traffic, the ICMP code appears instead.

Generator

The component that generated the event. See the Generator IDs table on page 489 for a list of intrusion event generator IDs.

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 314 for list of rule classification names and numbers.

VLAN ID

The innermost VLAN ID associated with the packet that triggered the intrusion event.

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.

Reviewed By

The name of the user who reviewed the event. Note that if an event was marked reviewed prior to an upgrade to version 4.10, the Reviewed By name is `unknown`.

Reviewing Intrusion Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 host 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. Note that if an intrusion event was marked reviewed prior to version 4.10,

Keysight, EX-1004
Page 278

the reviewer is listed as unknown, and unknown is a valid entry in the Reviewed By search field.

To mark an intrusion event reviewed:

Access: IPS/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 276 to learn more about the events that appear in intrusion event views. See Understanding Workflow Pages for Intrusion Events on page 280 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: Any IPS/Admin

▶ Select **Analysis & Reporting > IPS > Reviewed Events**.

The first page of the default reviewed intrusion events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

TIP! If you are using a custom workflow that does not include the table view of intrusion events, from the **Workflows** menu on the toolbar, select any of the workflows that ship with the appliance.

See Understanding Intrusion Events on page 276 to learn more about the events that appear in reviewed intrusion event views. See Understanding Workflow Pages for Intrusion Events on page 280 to learn more about how to narrow your view to the intrusion events that are important to your analysis.

Keysight, EX-1004
Page 279

To mark reviewed events unreviewed:

Access: IPS/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

Platform: DC/MDC/Standalone Sensor

Component: IPS

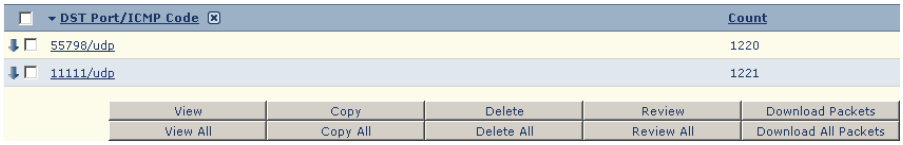
The preprocessor, protocol decoder, and intrusion rules that are enabled in the current intrusion policy generate intrusion events whenever the traffic that you monitor violates the policy.

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

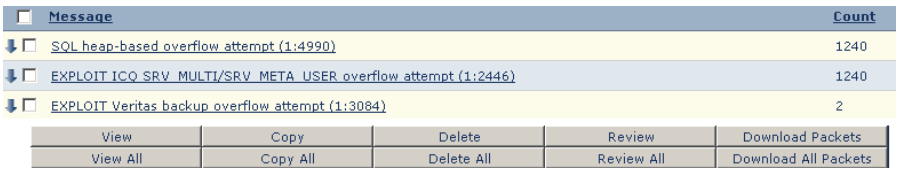
Drill-down pages generally include two or more columns in a table (and, for some drill-down views, more than one table) that allow you to view one specific type of information. For example, the following graphic shows a drill-down page with the number of events generated for each destination port.



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 21 (ftp)/tcp

Keysight, EX-1004
Page 280

and click **View**, the next page in the workflow appears (in this case, a page showing the number of events with each event message).

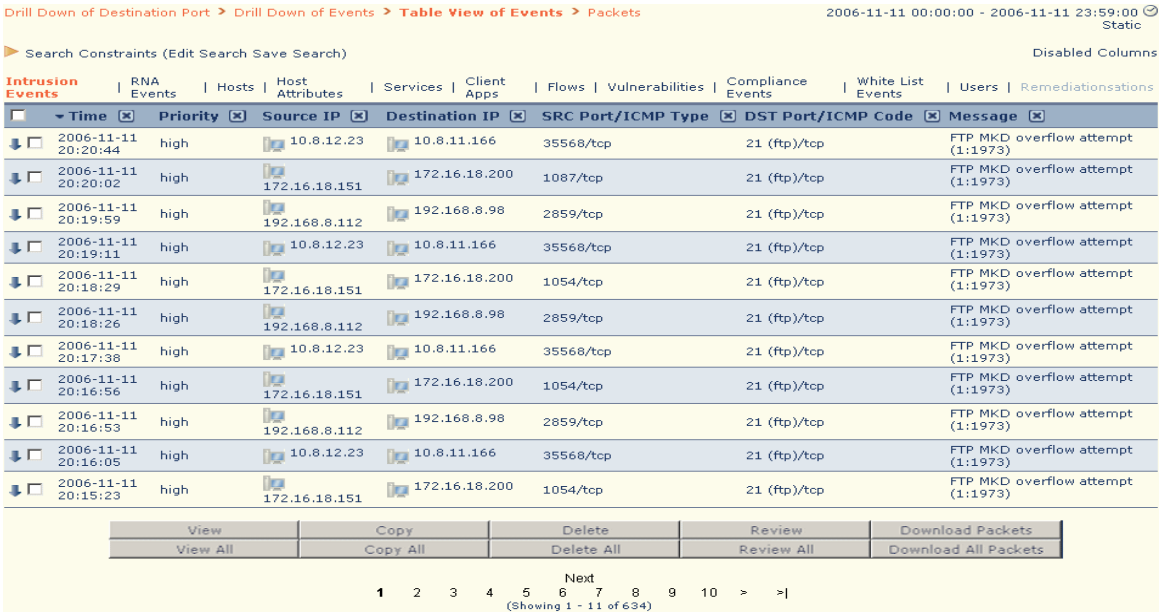


Message	Count
SQL heap-based overflow attempt (1:4990)	1240
EXPLOIT ICQ SRV_MULTI/SRV_META_USER overflow attempt (1:2446)	1240
EXPLOIT Veritas backup overflow attempt (1:3084)	2

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 4985 events from your analysis. On the second page, if you select the check box for the first row of events (FTP MKD overflow attempt (1:1973)) and click **View**, then the next page in the workflow appears and you have narrowed your list to 362 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.



Drill Down of Destination Port > Drill Down of Events > **Table View of Events** > Packets — 2006-11-11 00:00:00 - 2006-11-11 23:59:00 Static

Search Constraints (Edit Search Save Search) — Disabled Columns

Intrusion Events | RNA Events | Hosts | Host Attributes | Services | Client Apps | Flows | Vulnerabilities | Compliance Events | White List Events | Users | Remediationsations

Time	Priority	Source IP	Destination IP	SRC Port/ICMP Type	DST Port/ICMP Code	Message
2006-11-11 20:20:44	high	10.8.12.23	10.8.11.166	35568/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:20:02	high	172.16.18.151	172.16.18.200	1087/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:19:59	high	192.168.8.112	192.168.8.98	2859/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:19:11	high	10.8.12.23	10.8.11.166	35568/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:18:29	high	172.16.18.151	172.16.18.200	1054/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:18:26	high	192.168.8.112	192.168.8.98	2859/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:17:38	high	10.8.12.23	10.8.11.166	35568/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:16:56	high	172.16.18.151	172.16.18.200	1054/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:16:53	high	192.168.8.112	192.168.8.98	2859/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:16:05	high	10.8.12.23	10.8.11.166	35568/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)
2006-11-11 20:15:23	high	172.16.18.151	172.16.18.200	1054/tcp	21 (ftp)/tcp	FTP MKD overflow attempt (1:1973)

View | Copy | Delete | Review | Download Packets
View All | Copy All | Delete All | Review All | Download All Packets

Next
1 2 3 4 5 6 7 8 9 10 > >|
(Showing 1 - 11 of 634)

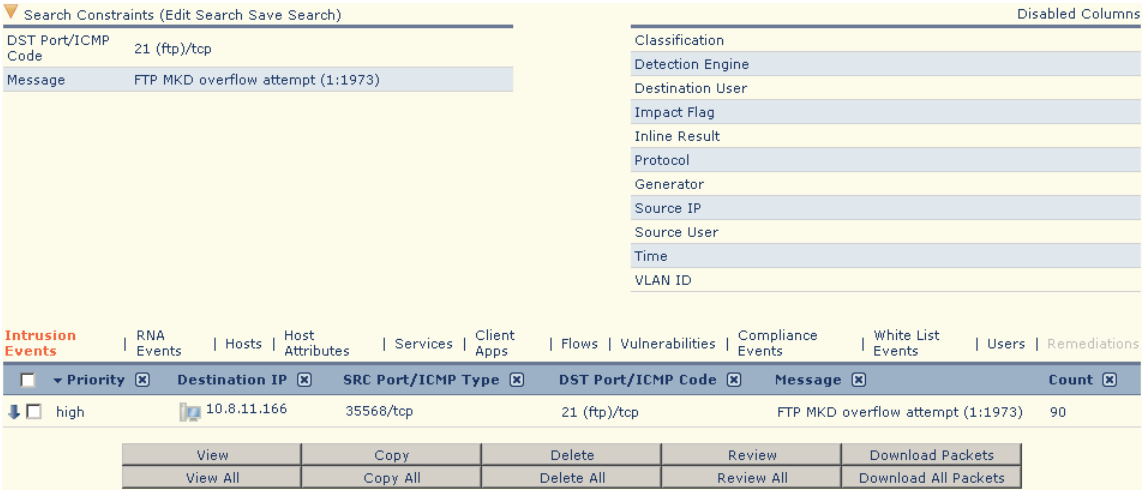
When you select events on a table view, instead of selecting events and displaying the next page in the workflow, you add to what are called *constraints*. Constraints are limits that you impose on the types of events that you want to analyze. In the previous two drill-down pages, the events were constrained in two ways: by destination port (21/tcp) and by event message (FTP MKD overflow

Keysight, EX-1004
Page 281

attempt (1:1973)). These constraints are carried forward to the table view. You can see the constraints by clicking the orange arrow at the top of the table.



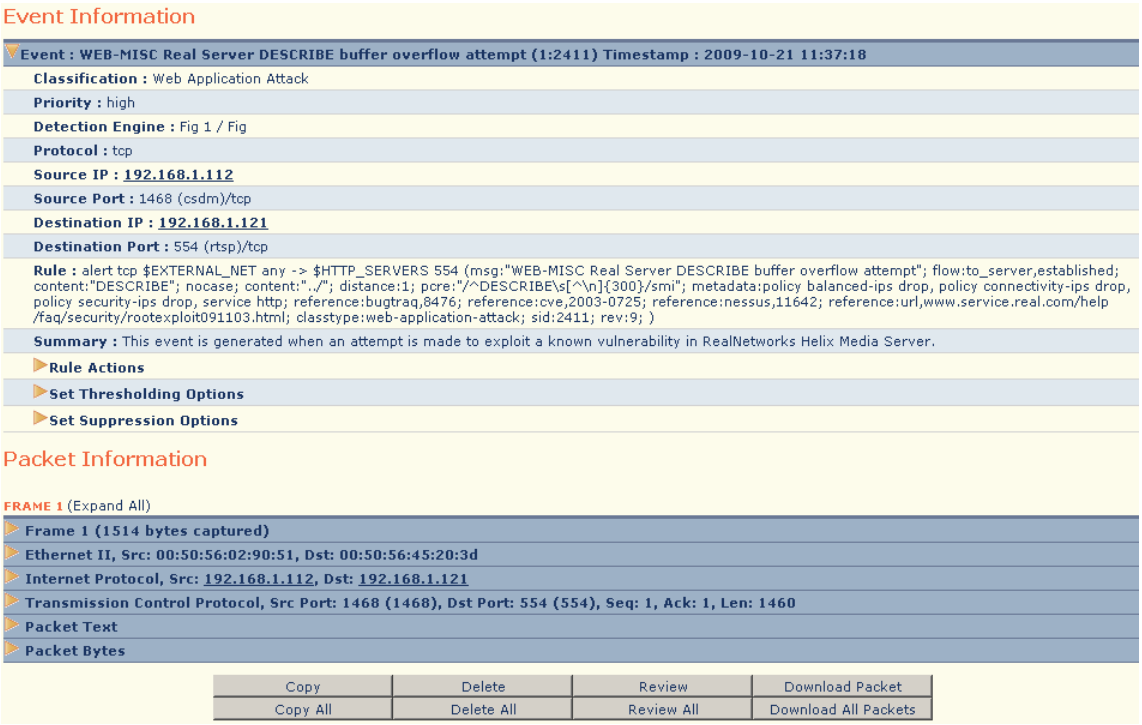
If you click the close column icon (✖) in the Time column from the table view, you can remove Time as one of the columns. To narrow the list of events in your analysis, you can click the link for a value in one of the rows in the table view. For example, to limit your analysis to the events generated 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, a single row that provides a count of the number of times the attacker (10.8.12.23) attempted to exploit the specific vulnerability (FTP MKD overflow attempt) against, in this case, a single destination IP address (10.8.11.166).



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

Keysight, EX-1004
Page 282

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 651 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; IPS automatically includes a packet view as the last page. You can easily switch between the predefined workflows and your own custom workflows depending upon how you want to investigate events.

TIP! Understanding and Using Workflows on page 1573 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.

Keysight, EX-1004
Page 283

For more information, see:

- Using Drill-Down and Table View Pages on page 284, 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 290, which explains how to use the features in the packet view.
- Searching for Intrusion Events on page 309 explains how to search the event database for specific intrusion events.

Using Drill-Down and Table View Pages

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 280.

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

Keysight, EX-1004
Page 284

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 276.
view a host's profile	click the host profile icon that appears next to the IP address of the host. This is available on the Master Defense Center and Defense Center only.
modify the time and date range for displayed events	find more information in see Setting Event Time Constraints on page 1597. 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 events on the current workflow page	find more information in: • Sorting Drill-down Workflow Pages on page 1612 • the Constraining Events on Drill-Down Pages table on page 288 • the Constraining Events on the Table View of Events table on page 289
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1613. **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 1597.
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 1593.

Keysight, EX-1004
Page 285

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	use one of the following methods: • 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**. The clipboard stores up to 25,000 events per user. For more information, see Using the Clipboard on page 316.
delete events from the event database	use one of the following methods: • 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. **IMPORTANT!** If you delete an intrusion events on the Defense Center, the event is only removed from the Defense Center database. If you have enabled event storage on a managed 3D Sensor and want to delete events from the sensor database, you must log into the sensor and delete them there.
mark events reviewed to remove them from intrusion event pages, but not the event database	use one of the following methods: • 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**. For more information, see Reviewing Intrusion Events on page 278.
download a local copy of the packet (a packet capture file in libpcap format) that triggered each selected event	use one of the following methods: • 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**. Captured packets are saved in libpcap format. This format is used by several popular protocol analyzers.
navigate to other event views to view associated events	find more information in Navigating between Workflows on page 1614. This is available on the Defense Center only.
temporarily use a different workflow	click **Workflows**. For more information, see Selecting Workflows on page 1590.

Keysight, EX-1004
Page 286

Intrusion Event Common Features (Continued)

To...	You can...
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.

The number of intrusion events that appear on the event views can 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 the events generated by IPS, 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 1597 and Navigating to Other Pages in the Workflow on page 1613.

Keysight, EX-1004
Page 287

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	click the value. 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.
drill down to the next workflow page constraining on selected events	select the check boxes next to the events you want to view on the next workflow page, then click **View**. 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. **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.
drill down to the next workflow page keeping the current constraints	click **View All**.

Keysight, EX-1004
Page 288

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. Note that the column is removed from the table.
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, use 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: • 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 1551.

Keysight, EX-1004
Page 289

Using the Packet View

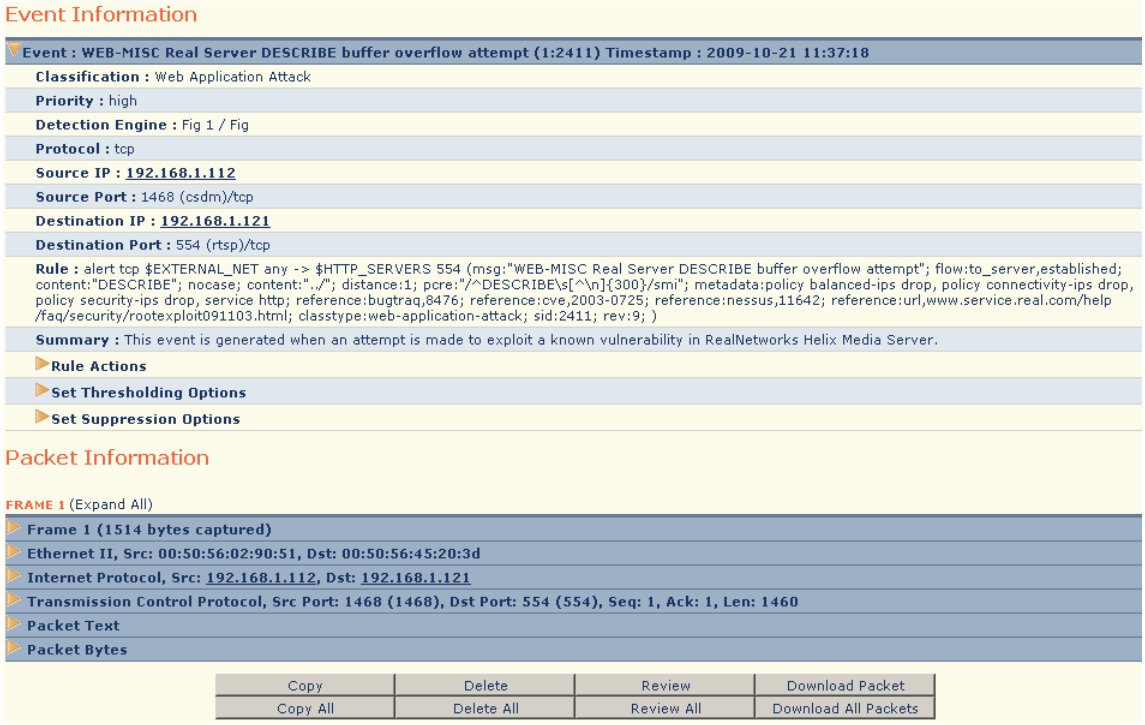
Platform: DC/MDC/Standalone Sensor

Component: IPS

A packet view provides information about the packet that triggered the rule that generated an intrusion event.

IMPORTANT! The packet view on a Defense Center or a Master Defense Center does not contain the packet information when the **Prohibit Packet Transfer to the Defense Center** option is enabled on the event detecting 3D Sensor.

The following graphic shows the Defense Center version of the packet view.



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

Keysight, EX-1004
Page 290

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 651 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 1597.
learn more about the information displayed in the packet view	find more information in: • Viewing Event Information on page 293 • Viewing Frame Information on page 299 • Viewing Data Link Layer Information on page 300 • Viewing Network Layer Information on page 301 • Viewing Transport Layer Information on page 303 • Viewing Packet Byte Information on page 306
add an event to the clipboard so you can transfer it to the incidents at a later time	click **Copy** to add the event whose packet you are viewing. The clipboard stores up to 25,000 events per user. For more information on the clipboard, see Using the Clipboard on page 316.
delete an event from the event database	either: • click **Delete** to delete the event whose packet you are viewing • click **Delete All** to delete all the events whose packets you previously selected

Keysight, EX-1004
Page 291

Packet View Actions (Continued)

To...	You can...
mark an event reviewed to remove it from event views, but not the event database.	either: • click **Review** to review the event whose packet you are viewing • click **Review All** to review all the events whose packets you previously selected For more information, see Reviewing Intrusion Events on page 278. Note that reviewed events continue to be included in the event statistics on the Intrusion Event Statistics page.
download a local copy of the packet (a packet capture file in libpcap format) that triggered the event	either: • 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 The captured packet is saved in libpcap format. This format is used by several popular protocol analyzers. 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. Note that you must have at least 15% available disk space in order to download.
expand or collapse a page section	click the arrow next to the section.

To display the packet view:

Access: Any IPS/Admin

▶ On the table view of intrusion events, you have three options:

- 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.

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.

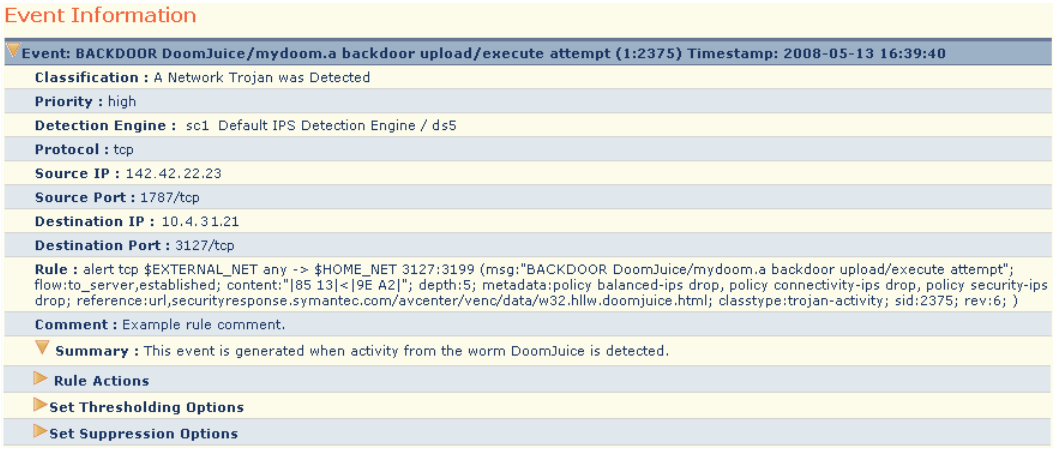
Keysight, EX-1004
Page 292

Viewing Event Information

Platform: DC/MDC/Standalone Sensor

Component: IPS

On the packet view, you can view information about the packet in the Event Information section.



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)*. *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. For more information, refer to Reading Preprocessor Generator IDs on page 488.

The Timestamp, that is, the time that the packet was captured, is also appended to the message.

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.

Keysight, EX-1004
Page 293

Detection Engine

The name of the detection engine that generated the event. On the Defense Center, the name of the sensor on which it resides is also listed.

Note that the master and slave in a clustered pair report intrusion events as if they were separate detection engines. See Using Clustered 3D Sensors on page 250 for more information.

Protocol

The protocol of the packet that triggered the event.

Source IP

The IP address of the host where the packet that triggered the event originated.

Source Port

Source port of the packet that triggered the event.

Destination IP

The IP address of the target host of the traffic that triggered the event.

Destination Port

Destination port of the packet that triggered the event.

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.

Comment

The rule comment. For rule-based events, this row appears only when you add a comment to the rule. For more information, see Adding Comments to Rules on page 867.

Summary

The rule summary. For rule-based events, this row appears when the rule documentation contains summary information.

Rule Actions

For standard text rule events, expand **Rule Actions** to edit, view documentation for, add a comment to, or change the state of the rule that triggered the event.

Keysight, EX-1004
Page 294

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Set Thresholding Options

Use this option to create a threshold for the rule that triggered this event. There are two possibilities depending on whether the intrusion policy that was applied to the detection engine that generated the event was created and applied locally. You may be able to set thresholding options:

- in the current policy
- in all locally created policies

The thresholding options are described in Setting Threshold Options within the Packet View on page 297.

Set Suppression Options

Use this option to create a suppression for the rule that triggered this event. There are two possibilities depending on whether the intrusion policy that was applied to the detection engine that generated the event was created and applied locally. You may be able to set suppression options:

- in the current policy
- in all locally created policies

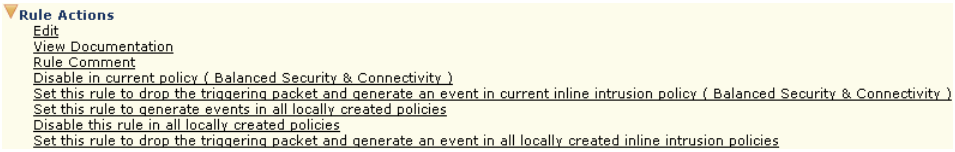
The suppression options are described in Setting Suppression Options within the Packet View on page 298.

Using Packet View Rule Actions

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.



Edit

For standard text rule events, click **Edit** to modify the rule that generated the event (if your user account has Policy & Response Administrator access).

Keysight, EX-1004
Page 295

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 864.

View Documentation

For standard text rule events, click **View Documentation** to learn more about the rule 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 867.

Disable this rule

If this event is generated by a rule, you can disable the rule, if necessary. There are two possibilities depending on whether the intrusion policy that was applied to the detection engine that generated the event was created and applied locally. You may be able to disable this rule:

- in the currently applied intrusion policy
- in all locally created intrusion policies

For more information, see Setting Rule States on page 435.

IMPORTANT! You **cannot** disable shared object rules from the packet view, nor can you disable rules in the default policies.

Keysight, EX-1004
Page 296

Set this rule to generate events

If this event is generated by a rule, you can enable the rule, that is, set it to generate events. There are two possibilities depending on whether the intrusion policy that was applied to the detection engine that generated the event was created and applied locally. You may be able to set the rule to generate events:

- in the currently applied intrusion policy; this option appears only if the rule was disabled since triggering the event
- in all locally created intrusion policies

For more information, see Setting Rule States on page 435.

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 3D Sensor is deployed inline on your network and you want to drop packets that trigger this rule, click the option in this field. There are two possibilities depending on whether the intrusion policy that was applied to the detection engine that generated the event was created and applied locally. You may be able to set the rule to drop:

- in the current intrusion policy; this option appears only in a policy where **Drop when Inline** is enabled; See Setting Drop Behavior in an Inline Deployment on page 358 for more information
- in all locally created intrusion policies

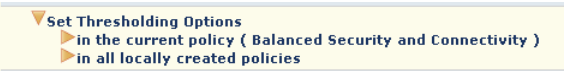
Note that if the rule is not currently set to generate events, you cannot set it to drop. For more information, see Setting Rule States on page 435.

Setting Threshold Options within the Packet View

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 the current policy, which includes the rule that triggered the event, or in all locally created policies.



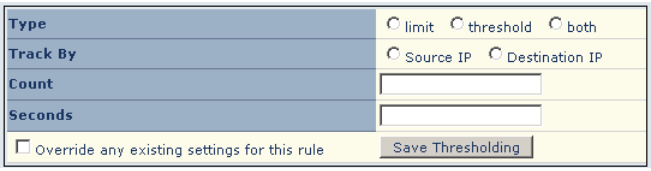
Keysight, EX-1004
Page 297

To set the threshold options within the packet view:

Access: Any IPS/Admin

1. Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Set Thresholding Options** in the Event Information section and click one of the two possible options:
   - **in the current policy**
   - **in all locally created policies**

   The thresholding options appear.



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, specify the number of event instances you want to use as your threshold.
5. In the **Seconds** field, specify the number of seconds that make up 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**.

   IPS 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 298

the current policy, which includes the rule that triggered the event, or in all locally created policies.

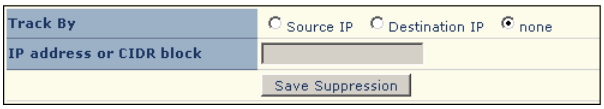


To suppress intrusion events within the packet view:

Access: Any IPS/Admin

1. Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Set Suppression Options** in the Event Information section and click one of the two possible options:
   - **in the current policy**
   - **in all locally created policies**

   The suppression options appear.



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 you want to specify as the source or destination IP address.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 299

reassembled TCP streams. For information on tagged packets, see Evaluating Post-Attack Traffic on page 850. For information on reassembled TCP streams, see Reassembling TCP Streams on page 637.

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.

Frame Number

The incremental frame number.

Frame Length

The length of the frame in bytes.

Protocols in frame

The protocols included in the frame.

Viewing Data Link Layer Information

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 300

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. On a Defense Center with both IPS and RNA, the NIC hardware vendor name is also shown.

Type

For Ethernet II packets, the type of packet that is encapsulated in the Ethernet frame; for example, IP 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

The following listing describes the protocol-specific information for an IP 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.

Keysight, EX-1004
Page 301

Differentiated Services

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)

Expand the view to display value definitions.

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.

Keysight, EX-1004
Page 302

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.

Expand the view to display value definitions.

Source

The IP address or domain name for the source host.

Note that to display the domain name, you must enable IP address resolution; for more information, see Configuring Event View Settings on page 45.

Click the address or domain name to view the context menu, then select **whois** to do a whois search on the host or **View Host Profile** to view host information. The host profile feature is available on the Master Defense Center and Defense Center only.

Destination

The IP address or domain name for the 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 45.

Click the address or domain name to view the context menu, then select **whois** to do a whois search on the host or **View Host Profile** to view host information. The host profile feature is available on the Master Defense Center and Defense Center only.

Viewing Transport Layer Information

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 303

The contents of the transport layer for each of the following protocols is described below:

- TCP Packet View on page 304
- UDP Packet View on page 305
- ICMP Packet View on page 306

IMPORTANT! Note that these examples discuss TCP, UDP, and ICMP packets; other protocols may also appear.

TCP Packet View

Platform: DC/MDC/Standalone Sensor

Component: IPS

This section describes the protocol-specific information for a TCP packet.

Source port

The number that identifies the originating application.

Destination port

The number that identifies the receiving application.

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

Keysight, EX-1004
Page 304

- `R` - reset the connection
- `S` - synchronize sequence numbers to start a new connection
- `F` - the sender has finished sending data

Expand the view to display value definitions.

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 in evasion attempt.

Expand the view to display value definitions.

Urgent Pointer

The position in the TCP segment where the urgent data ends. Used in conjunction with the U flag.

Options

The values, if present, for TCP options, where:

- `0` - the end of the options
- `1` - a padding value between options
- `2` - indicates the presence of 4 bytes specifying the maximum segment size (unlimited size if not specified)

Expand the view to display value definitions.

UDP Packet View

Platform: DC/MDC/Standalone Sensor

Component: IPS

This section describes the protocol-specific information for a UDP packet.

Source port

The number that identifies the originating application.

Destination port

The number that identifies the receiving application.

Length

The combined length of the UDP header and data.

Keysight, EX-1004
Page 305

Checksum

The indicator for whether the UDP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

Expand the view to display value definitions.

ICMP Packet View

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Viewing Packet Byte Information

Platform: DC/MDC/Standalone Sensor

Component: IPS

On the packet view, click the arrow next to **Packet Bytes** to view hexadecimal and ASCII versions of the bytes that comprise the packet.

Keysight, EX-1004
Page 306

Using Impact Flags to Evaluate Events

Platform: DC/MDC/Standalone Sensor

Component: IPS + RNA

When you install 3D Sensors with IPS and RNA on the same network segment and then view the intrusion events on a Defense Center that manages both sensors, you can gain more insight into the risk, or impact, that the events have on your network.

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 IPS data, RNA 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 red (Vulnerable) 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

Icon	Impact	Description
	gray	Unknown, indicating that neither the source nor the destination host is on a network that is monitored by RNA.
	blue	Unknown Target, indicating that either the source or destination host is on a monitored network, but there is no entry for the host in RNA's network map.
	yellow	Currently Not Vulnerable, indicating that either the source or the destination host is in RNA's network map and one of the following is true: • 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

Keysight, EX-1004
Page 307

Impact Flags (Continued)

Icon	Impact	Description
	orange	Potentially Vulnerable, indicating that either the source or the destination host is in RNA's network map and one of the following is true: • for port-oriented traffic, the port is running a service • for non-port-oriented traffic, the host uses the protocol
	red	Vulnerable, indicating that either: • the source or the destination host is in the RNA network map, and a vulnerability is mapped to the host • the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software; see Setting a Red Impact Flag on page 812 for more information

To use the impact flag on the table view to evaluate events:

Access: Any IPS/ Admin

1. Select **Analysis & Reporting > IPS > Intrusion Events**.

   The first page of the default intrusion events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

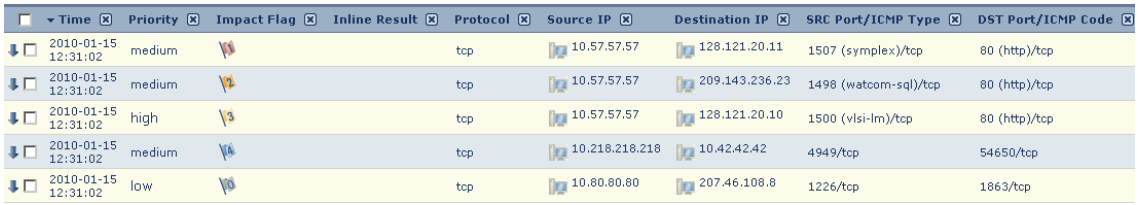
2. Constrain the event view to view only those events that you want to evaluate.

   For more information, see Using Drill-Down and Table View Pages on page 284.

Keysight, EX-1004
Page 308

3. At the top of the page, click **Table View of Events**.

   The table view of events appears. On the Defense Center, **Impact Flag** can have any of the values described in the Impact Flags table on page 307. The following graphic shows the Defense Center version of the page. Note that several of the columns were removed from the table view to simplify the graphic.



Time	Priority	Impact Flag	Inline Result	Protocol	Source IP	Destination IP	SRC Port/ICMP Type	DST Port/ICMP Code
2010-01-15 12:31:02	medium	1		tcp	10.57.57.57	128.121.20.11	1507 (symplex)/tcp	80 (http)/tcp
2010-01-15 12:31:02	medium	1		tcp	10.57.57.57	209.143.236.23	1498 (watcom-sql)/tcp	80 (http)/tcp
2010-01-15 12:31:02	high	3		tcp	10.57.57.57	128.121.20.10	1500 (vlsi-lm)/tcp	80 (http)/tcp
2010-01-15 12:31:02	medium	4		tcp	10.218.218.218	10.42.42.42	4949/tcp	54650/tcp
2010-01-15 12:31:02	low	10		tcp	10.80.80.80	207.46.108.8	1226/tcp	1863/tcp

4. To sort the table by the impact flag value, click **Impact Flag**.

   The events are sorted by impact flag in descending vulnerability order: from red to gray. High-priority events that are marked with a red icon are the most likely to have an effect on your network. The sort order is red, orange, yellow, blue, and gray.

   TIP! To sort the table in ascending vulnerability order, click **Impact Flag** again.

Searching for Intrusion Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can search for specific intrusion events by using one of the predefined searches 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 re-use later. The search criteria you can use are described in the following listing.

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 1557 and Specifying Ports in Searches on page 1558.

Keysight, EX-1004
Page 309

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`.

Protocol

Type the name or number of the transport protocol used in the event, as listed in http://www.iana.org/assignments/protocol-numbers/.

Detection Engine

Specify the name of the detection engine or detection engine group that generated the intrusion events you want to view. For Defense Centers, you can also include the name of the sensor or sensor group. For Master Defense Centers you can also include the Defense Center or Defense Center group name. For more information, see Specifying Detection Engine/Appliance Names on page 1559.

Note that the master and slave in a clustered pair report intrusion events as if they were separate detection engines. See Identifying Detection Engines and Interface Sets on a Clustered Pair on page 251 for more information.

Source IP

Specify the IP address of the source host involved in the intrusion events.

Destination IP

Specify the IP address of the destination host involved in the intrusion events.

Original Client IP

Specify the original client IP address extracted from the X-Forwarded-For (XFF) or True-Client-IP HTTP headers. See Understanding Intrusion Events on page 276 for more information.

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 IP

Specify the source or destination IP address of the host whose intrusion events you want to view.

Keysight, EX-1004
Page 310

Source/Destination User

Specify the User ID for a user logged in to the source or destination host

SRC 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.

DST 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.

Message

Specify all or part of the event message for the events you want to view.

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 a single SID, a single GID:SID combination, or a comma-separated list of either or both. For more information, see Reading Preprocessor Generator IDs on page 488.

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.

Impact Flag

For Defense Centers, specify the impact flag that the Defense Center has assigned to the intrusion event based on the correlation between IPS and RNA data. Valid values are `red`, `orange`, `yellow`, `blue`, and `gray`.

For more information, see Using Impact Flags to Evaluate Events on page 307.

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

Keysight, EX-1004
Page 311

Note that IPS 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 435, Setting Drop Behavior in an Inline Deployment on page 358, and Tap Mode on page 227.

Reviewed By

The name of the user who reviewed the event. Note that if an event was marked reviewed prior to an upgrade to version 4.10, the Reviewed By name is `unknown`.

Generator

Specify the component that generated the events you want to view, as listed in the Generator IDs table on page 489.

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-delimited 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 314 for a list of classification numbers, names, and descriptions.

VLAN ID

Specify the innermost VLAN ID associated with the packet that triggered the intrusion event.

For more information on searching, including how to load and delete saved searches, see the Searching for Events table on page 1550.

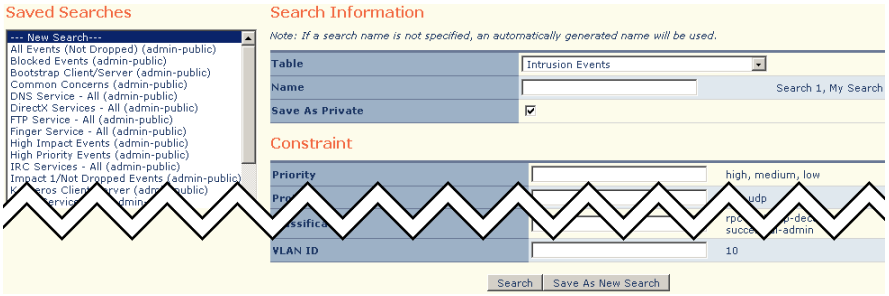
Keysight, EX-1004
Page 312

To search for intrusion events:

Access: Any IPS/ Admin

1. Select **Analysis & Reporting > Searches > Intrusion Events**.

   The Intrusion Events search page appears.



> **TIP!** To search 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 Searching for Intrusion Events on page 309. 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 data 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 intrusion events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

Keysight, EX-1004
Page 313

- 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.

Specifying Rule Classifications in Event Searches

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Rule Classifications table lists the name and number for each rule classification. In an event search, you can use the classification number, or all or part of either the classification name or description.

Rule Classifications

Number	Classification Name	Description
1	not-suspicious	Not Suspicious Traffic
2	unknown	Unknown Traffic
3	bad-unknown	Potentially Bad Traffic
4	attempted-recon	Attempted Information Leak
5	successful-recon-limited	Information Leak
6	successful-recon-largescale	Large Scale Information Leak
7	attempted-dos	Attempted Denial of Service
8	successful-dos	Denial of Service
9	attempted-user	Attempted User Privilege Gain
10	unsuccessful-user	Unsuccessful User Privilege Gain
11	successful-user	Successful User Privilege Gain
12	attempted-admin	Attempted Administrator Privilege Gain
13	successful-admin	Successful Administrator Privilege Gain
14	rpc-portmap-decode	Decode of an RPC Query

Keysight, EX-1004
Page 314

Rule Classifications (Continued)

Number	Classification Name	Description
15	shellcode-detect	Executable Code was Detected
16	string-detect	A Suspicious String was Detected
17	suspicious-filename-detect	A Suspicious Filename was Detected
18	suspicious-login	An Attempted Login Using a Suspicious Username was Detected
19	system-call-detect	A System Call was Detected
20	tcp-connection	A TCP Connection was Detected
21	trojan-activity	A Network Trojan was Detected
22	unusual-client-port-connection	A Client was Using an Unusual Port
23	network-scan	Detection of a Network Scan
24	denial-of-service	Detection of a Denial of Service Attack
25	non-standard-protocol	Detection of a Non-Standard Protocol or Event
26	protocol-command-decode	Generic Protocol Command Decode
27	web-application-activity	Access to a Potentially Vulnerable Web Application
28	web-application-attack	Web Application Attack
29	misc-activity	Misc Activity
30	misc-attack	Misc Attack
31	icmp-event	Generic ICMP Event
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

Keysight, EX-1004
Page 315

Using the Clipboard

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 284 and Using the Packet View on page 290.

The contents of the clipboard are sorted by the date and time that the events were generated. Once you add intrusion events to the clipboard, you can delete them from the clipboard as well as generate reports on the contents of 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 324.

See the following sections for more information:

- Generating Clipboard Reports on page 316
- Deleting Events from the Clipboard on page 318

Generating Clipboard Reports

Platform: DC/MDC/Standalone Sensor

Component: IPS

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:

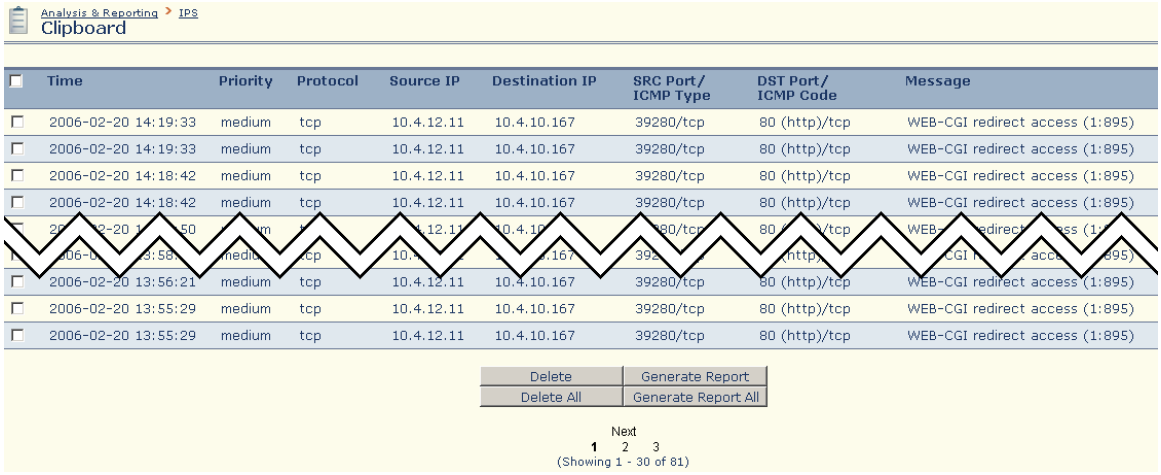
Access: Any IPS/Admin

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 284.
   - For information on how to add events to the clipboard from the packet view, see Using the Packet View on page 290.

Keysight, EX-1004
Page 316

2. Select **Analysis & Reporting > IPS > Clipboard**.

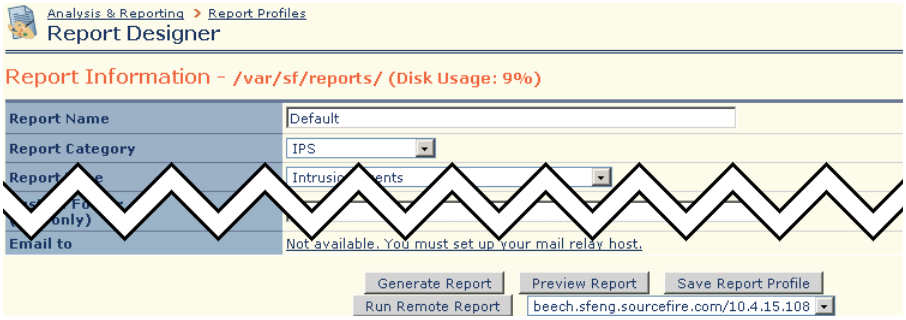
   The clipboard appears.



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 Designer page appears.



4. Specify how you want your report to look, then click **Generate Report**.

> **TIP!** For more information about using the Report Designer, see Working with Event Reports on page 1531.

5. To view the report, click **Reports** on the toolbar.

   The Reporting page appears.

6. Click **View** next to the report you generated.

   The clipboard report appears in your browser window.

Keysight, EX-1004
Page 317

Deleting Events from the Clipboard

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: IPS/Admin

1. Select **Analysis & Reporting > IPS > 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.

Keysight, EX-1004
Page 318

Chapter 8
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 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 320
- Creating an Incident on page 324
- Editing an Incident on page 326
- Generating Incident Reports on page 327
- Creating Custom Incident Types on page 329

Keysight, EX-1004
Page 319

Incident Handling Basics

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 320
- Common Incident Handling Processes on page 321
- Incident Types in the Sourcefire 3D System on page 324

Definition of an Incident

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 Sourcefire 3D System to track your response to an incident.

As explained in Working with Intrusion Events on page 267, 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 IPS 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.

Keysight, EX-1004
Page 320

Common Incident Handling Processes

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 321
- Detection and Notification on page 321
- Investigation and Qualification on page 322
- Communication on page 322
- Containment and Recovery on page 323
- Lessons Learned on page 324

Each of these phases is described in the sections that follow. The descriptions also explain how 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 on each detection engine, 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 3D Sensors with IPS 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 intrusion policy you apply to each of the IPS detection engines 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

Keysight, EX-1004
Page 321

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. With standalone 3D Sensors with IPS, 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, it is automatically prioritized for you. If you also deploy a 3D Sensor with RNA on the same network segment and use a Defense Center to correlate the information from both sensors, then events are automatically 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?

Keysight, EX-1004
Page 322

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 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.

Keysight, EX-1004
Page 323

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 329.

Creating an Incident

Platform: DC/MDC/Standalone Sensor

Component: IPS

This section explains how you create an incident.

To create an incident:

Access: Any IPS/Admin

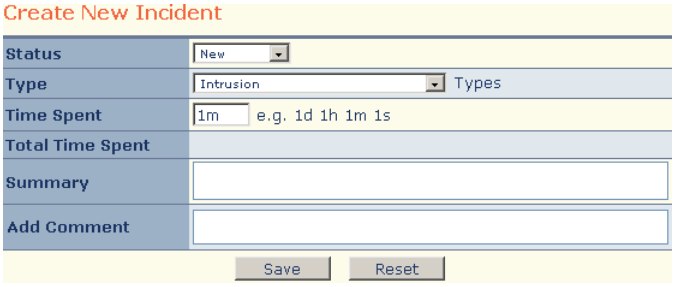
1. Select **Analysis & Reporting > IPS > Incidents**.

   The Incident page appears.

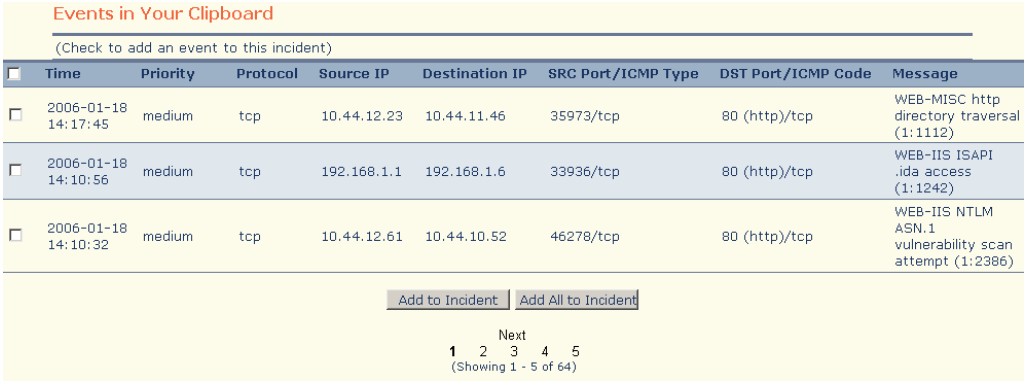
Keysight, EX-1004
Page 324

2. Click **Create Incident**.

   The Create New Incident page appears.



   If you previously copied intrusion events to the clipboard, they are displayed at the bottom of the page. See Using the Clipboard on page 316 for information about using the clipboard.



3. From the **Type** drop-down menu, select the option that best describes the incident. The following incident types are available:
   - Intrusion
   - Denial of Service
   - Unauthorized Admin Access
   - Web Site Defacement
   - Compromise of System Integrity
   - Hoax
   - Theft
   - Damage
   - Unknown
4. 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.

Keysight, EX-1004
Page 325

5. In the **Summary** text box, type a short description (up to 255 alphanumeric characters spaces, and symbols) of the incident.
6. 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can update an incident as you collect more information, including any of the following aspects:

- change the status
- change the type
- add events from the clipboard
- delete events

You can also add or delete events from the incident as your investigation progresses.

To edit an incident:

Access: Any IPS/Admin

1. Select **Analysis & Reporting > IPS > Incidents**.

   The Incident page appears.
2. Click **Edit** 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

Keysight, EX-1004
Page 326

- 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: Any IPS/Admin

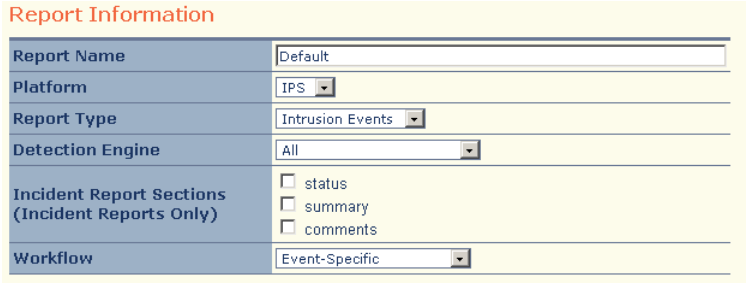
1. Select **Analysis & Reporting > IPS > Incidents**.

   The Incident page appears.
2. Click **Edit** next to the incident you want to include in your report.

Keysight, EX-1004
Page 327

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 Report Designer 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 the 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.

Keysight, EX-1004
Page 328

Creating Custom Incident Types

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: IPS/Admin

1. Select **Analysis & Reporting > IPS > Incidents**.

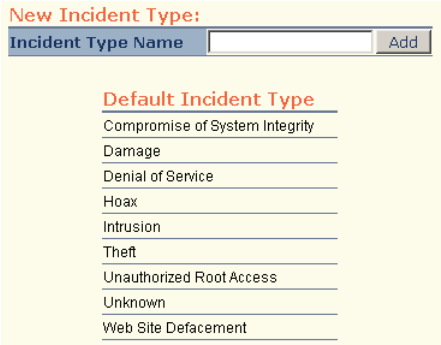
   The Incident page appears.

2. Click **Create Incident**.

   The Create New Incident page appears.

3. In the Type area, click **Types**.

   The Incident Management page appears. The default incident types are listed at the bottom of the page.



Keysight, EX-1004
Page 329

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 Incident page.

   You can use the new incident type the next time you create or edit an incident.

Keysight, EX-1004
Page 330

Chapter 9
Creating Intrusion Policies

An *intrusion policy* is a defined set of intrusion detection and prevention configurations that you can apply to an IPS detection engine on a 3D Sensor that is licensed for the IPS component.

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.

When you install a standalone or managed Sourcefire 3D Sensor that is licensed for the IPS component, and you have completed the initial setup procedure, the detection engine on your sensor immediately begins monitoring your network for malicious traffic and policy violations. You can begin reviewing the events that IPS generates and evaluating whether they are important in the context of your network environment and your security policies.

However, although you can begin monitoring your network immediately, you should modify and reapply your intrusion policy to improve the performance of IPS in your environment and to provide a focused view of the traffic on your

Keysight, EX-1004
Page 331

network. At a minimum, you should consciously choose whether to configure the following settings in your intrusion policy:

- Specify whether you want IPS 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 358 for more information.
- Identify the detection engines where you will apply your policy so you can more easily apply and reapply your policy after modifying it. See Managing Targets on page 368 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 371 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 435 for more information.
- If your deployment includes the RNA component, use RNA data to associate hosts and services on your network with rules written to protect those hosts and services and recommend rule state changes. See Managing RNA Rule State Recommendations on page 468 for more information.

You can configure multiple policies to apply to different detection engines. You must apply one policy to each detection engine except in the case of policies filtered to detect VLAN or subnetwork traffic, in which case you apply one non-filtered policy and one or more VLAN- or subnetwork-filtered policies to each detection engine.

See the following sections for more information:

- Planning and Implementing an Intrusion Policy on page 333 describes, at a high level, the process you use to create an intrusion policy.
- Managing Intrusion Policies on page 335 explains how to view a listing of your intrusion policies, create and edit policies, apply them to your detection engines, and view intrusion policy reports.
- Setting Drop Behavior in an Inline Deployment on page 358 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 361 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 Targets on page 368 explains how to specify detection engines where you will apply your intrusion policies.

Keysight, EX-1004
Page 332

- Managing Variables on page 371 explains how you can use detection-engine-specific, policy-specific, and system default variables to tailor your policy to your network.
- Managing Rules in an Intrusion Policy on page 403 explains how you can enable and disable rules and configure other rule attributes such as thresholds, suppression, and so on.
- Defining VLANs and Subnetworks on page 393 explains how to apply policies filtered for specific VLANs and subnetworks to your detection engines.
- Managing RNA Rule State Recommendations on page 468 explains how you can generate rule state recommendations for intrusion rules based on RNA data that identifies the hosts and services on your network.
- Using Advanced Settings in an Intrusion Policy on page 477 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 398 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 494 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

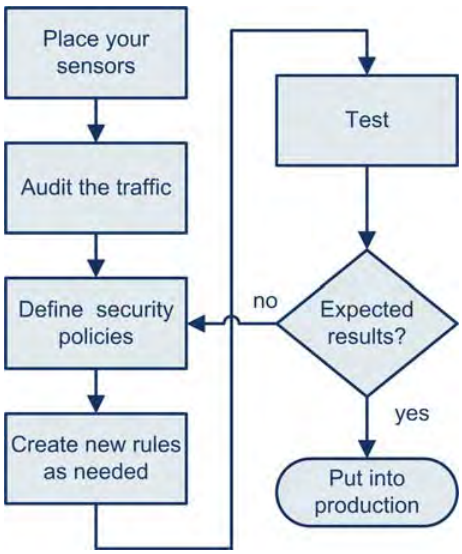
Platform: DC/MDC/Standalone Sensor

Component: IPS

Building custom intrusion policies can improve the performance of IPS in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network.

Keysight, EX-1004
Page 333

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 3D Sensors with IPS on the network.

   There are a variety of deployment options in tuning your sensor. For details on deciding where to place your 3D Sensors to best monitor the traffic that matters to you, see the *Installation Guide* for your sensor.

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 configuration.

   - 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.

Keysight, EX-1004
Page 334

- 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 477.
- 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, detection or system level allows you to tune your policy, detection engines, 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 371 and Using Variables within Detection Engines on page 213.
- 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 sensors with inline detection engines, 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 435.

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 761.
5. Test your configuration.

Managing Intrusion Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

On the Intrusion Policy page (**Policy & Response > IPS > Intrusion Policy**) you can view all your current intrusion policies by name with optional description as well as the time and date the policy was last modified and the user who modified it. 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

Keysight, EX-1004
Page 335

saved settings in each policy, apply a policy to one or more detection engines, and edit, delete, or export a policy.

TIP! You can import intrusion policies from other Defense Centers and 3D Sensors in your deployment. See Importing and Exporting Objects on page 1997 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).

Status information provides the following information:

- when dropping packets in an inline deployment is enabled in the policy
- when filtering traffic by VLAN or network is configured in the policy
- whether you have applied modified policies to detection engines targeted in individual policies
- when a policy has unsaved changes

The text style of policy-apply and saved-changes messages varies as follows:

- Plain black text indicates the number of targeted detection engines where you have applied the most recently saved changes to the policy.
- Plain red text indicates the number of targeted detection engines where you have not applied the most recently saved changes in the policy.
- *Italicized* black text indicates that the policy has unsaved changes.

  Note that modifying a rule in the rule editor (**Policy & Response > IPS > Rule Editor**) when the rule is enabled in your intrusion policy does not cause the status to indicate that your policy is out of date on a targeted detection engine.

You can take the following actions on the Intrusion Policy page:

Compare Policies

Compare the settings of two intrusion policies or two revisions of the same policy. See Comparing Two Intrusion Policies on page 354 for more information.

Create Policy

Create a new intrusion policy. See Creating an Intrusion Policy on page 338 for more information.

Report

View a PDF report that lists the current configuration settings in an intrusion policy. See Viewing an Intrusion Policy Report on page 350 for more information.

Keysight, EX-1004
Page 336

Apply

Apply an intrusion policy to one or more detection engines. See Applying an Intrusion Policy on page 347 for more information.

Edit

Modify the settings in an intrusion policy; alternatively, click the policy name. See Editing an Intrusion Policy on page 341 for information.

Delete

Delete an intrusion policy; click **Delete**, 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.

Export

Export an intrusion policy for use on another appliance of the same type. See Exporting an Intrusion Policy on page 2001 for more information.

See the following sections for more information:

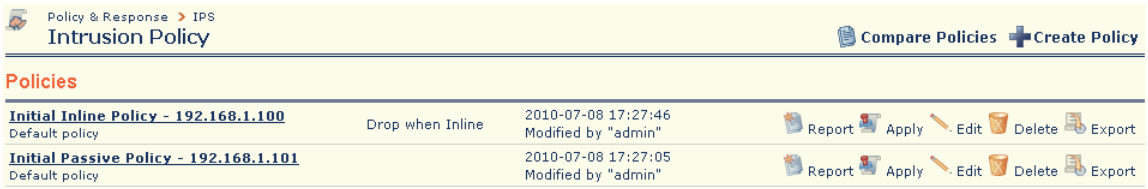
- Creating an Intrusion Policy on page 338
- Editing an Intrusion Policy on page 341
- Using the Navigation Panel on page 344
- Committing Intrusion Policy Changes on page 345
- Applying an Intrusion Policy on page 347
- Viewing an Intrusion Policy Report on page 350
- Comparing Two Intrusion Policies on page 354

To manage your intrusion policies:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.



2. Take any of the actions described in this section.

Keysight, EX-1004
Page 337

Creating an Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can create one or more intrusion policies. For example, you can create policies that you apply to your detection engines to 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 RNA 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 and 3D Sensors in your deployment. See Importing and Exporting Objects on page 1997 for more information.

To create an intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

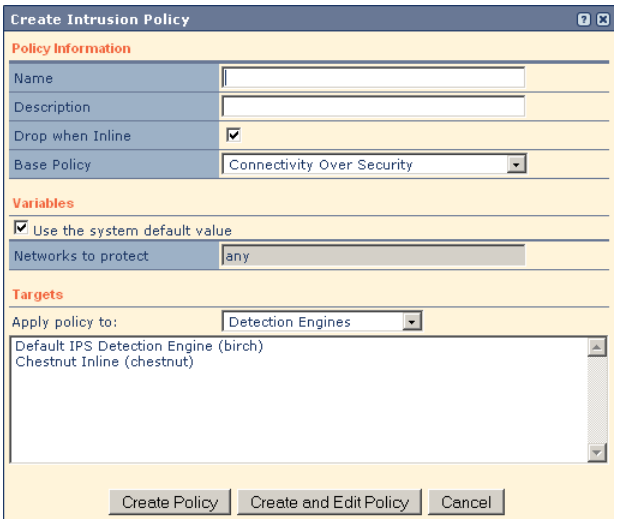
   The Intrusion Policy page appears.

Keysight, EX-1004
Page 338

2. Click **Create Policy**.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Create Intrusion Policy pop-up window appears.



3. Type a unique name that identifies your policy and, optionally, a description that differentiates it from other policies.

4. Specify whether you want IPS 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 IPS does not drop packets in a passive deployment.

   See Setting Drop Behavior in an Inline Deployment on page 358 for more information.

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 361 for more information.

Keysight, EX-1004
Page 339

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 372 and Modifying Variables on page 374 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 (any IP address) or one or more IP addresses, CIDR blocks, or variables, including negation (!), separating multiple entries with commas. You must enclose a negated list of addresses in brackets, for example: `![192.168.1.1,192.168.1.5]`

     For more information, see Managing Variables on page 371, Defining IP Addresses and Ports for Your Network on page 398, and IP Address Conventions on page 59.
7. Optionally, identify specific IPS detection engines in your intrusion policy so you can quickly apply your policy to the targeted detection engines each time you apply the policy.

   Select **Detection Engines**, **Detection Engine Groups** or **Sensors** from the **Apply policy to:** drop-down list, then select one or more entries in the field below the drop-down list. Use Ctrl or Shift while clicking to select multiple items or to clear items.

   Sourcefire recommends for your convenience that you target the detection engines on the network where you will apply the policy. See Managing Targets on page 368 and Applying an Intrusion Policy on page 347 or more information.
8. You have the following options:
   - You can to exit the pop-up window without creating a policy. Click **Cancel**.

     The Intrusion Policy page appears.
   - You can save your changes. Click **Create Policy**.

     The Intrusion Policy page appears.

     You must apply the policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347 for more information.
   - You can open the advanced intrusion policy editor. Click **Create and Edit Policy**. See Editing an Intrusion Policy on page 341 for more information.

Keysight, EX-1004
Page 340

Editing an Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 477 and Using Layers in an Intrusion Policy on page 494.

You can configure the following commonly used settings in the advanced editor:

Name

Type a unique name for the policy.

Description

Type a description to differentiate the policy.

Drop when Inline

Select (enable) or clear (disable) to specify whether rules set to Drop and Generate Events drop offending packets in an inline deployment. See Setting Drop Behavior in an Inline Deployment on page 358 for more information.

Base Policy

Select a Sourcefire-authored or user-created policy from the drop-down list to determine the source of the default settings for all rules and advanced settings within the intrusion policy. See Understanding the Base Policy on page 361 for more information.

You can also click **Manage Base Policy** to view the advanced settings that are enabled by default in your base policy. See Using Advanced Settings in an Intrusion Policy on page 477 for more information.

Manage Targets

Identify one or more detection engines where you want to apply your intrusion policy.

Note that Sourcefire recommends that you target the detection engines on the network where you will apply the policy. See Managing Targets on page 368 and Applying an Intrusion Policy on page 347 or more information.

Manage Variables

Tailor variables for your specific network environment. See Managing Variables on page 371 for more information.

Keysight, EX-1004
Page 341

Make sure you accurately specify the values for $HOME_NET, which represents the IP addresses on the network segment protected by IPS, 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.

Manage Rules

Display or modify configured rule attributes for the rules in your intrusion policy. You can also display a filtered view of the intrusion policy Rules page showing rules enabled in your policy by current rule state and, optionally, you can set rule attributes for specified rules: click **View** next to the number of rules that are set to Generate Events or to Drop and Generate Events.

See Managing Rules in an Intrusion Policy on page 403 for more information.

RNA Recommendations

Click this option in the navigation panel to display the RNA Recommended Rules Configuration page where you can elect to use rule state recommendations for standard text rules and shared object rules based on RNA data.

Alternately, click **Click here to set up RNA recommendations** on the Policy Information page if you have not generated recommendations, or **Click to change recommendations** if you have generated recommendations.

You can also display a filtered view of the Rules page showing rules with recommended rule states and, optionally, you 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.

See Managing RNA Rule State Recommendations on page 468 for more information.

Policy by VLAN or Network

Click this option in the navigation panel to the left of the Policy Information page to configure a VLAN- or subnetwork-filtered intrusion policy. See Defining VLANs and Subnetworks on page 393 for more information.

To edit an intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

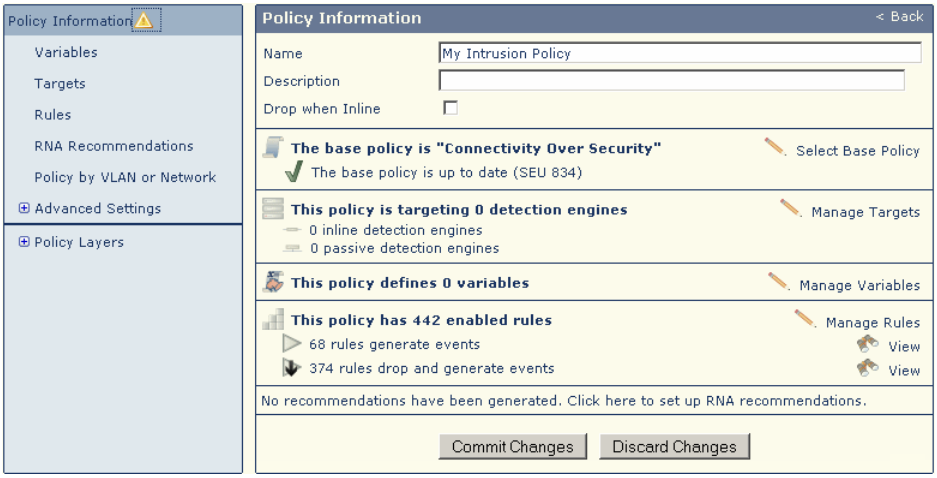
Keysight, EX-1004
Page 342

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

   The following graphic shows the Defense Center version of the Policy Information page.



3. You can take any of the actions described in Editing an Intrusion Policy on page 341.

4. You have three options:

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 343

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

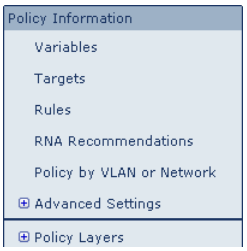
  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Using the Navigation Panel

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 with the exception of the Policy by VLAN or Network page, which you can access only from the navigation panel.

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

Keysight, EX-1004
Page 344

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 477 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 494 and Managing Rules in an Intrusion Policy on page 403 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You must save (that is, commit) changes to your intrusion policy before you can apply the changes to a detection engine. When you apply a policy, the system applies the most recently saved configuration. See Applying an Intrusion Policy on page 347 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 only when you edit another policy as the same user without saving your changes, or when you import an SEU that updates the cached policy. See Importing SEUs and Rule Files on page 1848 for more information.

Note that the policy status on the Intrusion Policy page (**Policy & Response > IPS > Intrusion Policy**) does not indicate that your policy is out of date on a targeted detection engine when you modify a rule in the rule editor (**Policy & Response > IPS > Rule Editor**) and enable the rule in your intrusion policy.

Keysight, EX-1004
Page 345

The following might also occur when you commit your changes:

- If you have not targeted detection engines in your policy, you are prompted whether to continue. Click **OK** to commit your changes or click **Cancel** and the detection engines Targets management page appears.
- 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 1735 and Viewing Audit Records on page 1969 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 Editing a System Policy on page 1735 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, or click **Cancel** to return to the Policy Information page. See Automatically Enabling Advanced Settings on page 490 for more information.
- 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.

To commit unsaved changes to an intrusion policy:

Access: P&R Admin/Admin

1. You have the following choices, depending on whether you are currently editing your policy:
   - If you are currently editing your policy, skip to step 4.
   - If you are not currently editing your policy, continue with the next step.
2. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.
3. Click **Edit** next to the policy with unsaved changes that you want to commit.

   The Policy Information page appears.

Keysight, EX-1004
Page 346

4. You have the following options:
   - You can continue editing your policy.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

Applying an Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

After making any changes to an intrusion policy, you must apply the policy to one or more detection engines to implement the configuration changes on the networks monitored by the detection engines.

You can quickly apply or reapply a policy to detection engines that you have targeted in the intrusion policy. If you have not targeted detection engines in the policy, you can identify the detection engines where you want to apply the policy when you apply the policy. Sourcefire recommends that you associate each intrusion policy with targeted detection engines on your network.

Keep the following points in mind when applying intrusion policies:

- Unless **Inspect Traffic During Policy Apply** is enabled for the detection engine for a 3D Sensor other than a 3D500, 3D1000, or 3D3800, applying an intrusion policy causes the detection engines on the sensor to restart, which can cause a short pause in processing and, for detection engines with inline interface sets, may cause a few packets to be dropped. Make sure you plan these actions for times when they will have the least impact on your deployment. If **Inspect Traffic During Policy Apply** is enabled, the

Keysight, EX-1004
Page 347

detection engine does not restart and interrupt traffic inspection when the policy is applied. This option may degrade performance when you apply a policy and may result in longer policy-apply periods. See Creating a Detection Engine on page 206 for more information.

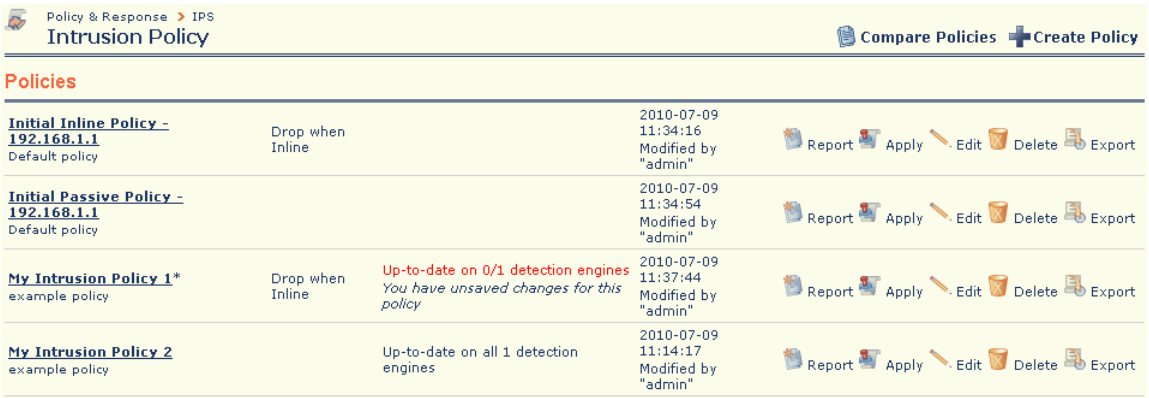
- For Crossbeam-based IPS software deployed inline, you may want to remove the affected VAP from the load-balanced list until the detection engine restarts, then reinstate it. For more information, refer to the *Installation and Configuration Guide* for your software sensor.
- Rules set to Drop and Generate Events in a policy where **Drop when Inline** is selected will generate events but will **not** drop any packets or block any attacks when you apply the policy to a detection engine that uses a passive interface set or an inline interface set in tap mode.
- You cannot apply intrusion policies to a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). See Using Clustered 3D Sensors on page 250 for more information.
- You can use the **Reapply All** or **Reapply Intrusion Policies** option on the Detection Engines page (**Operations > Configuration > Detection Engines**) to reapply your intrusion policies. These options can be convenient, for example, if you import an SEU and do not automatically reapply your intrusion policies after the import completes. See Understanding Detection Engines on page 199 and Importing SEUs and Rule Files on page 1848 for more information.

To apply an intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears. The following graphic shows the Defense Center version of the page.



Keysight, EX-1004
Page 348

2. Click **Apply** next to the policy you want to apply. You have the following choices, depending on whether your most recently saved intrusion policy targets one or more detection engines:
   - If your intrusion policy targets one or more detection engines, the Apply to Detection Engines pop-up window appears; go to step 3.
   - If your intrusion policy does not target detection engines, the Apply Policy page appears; go to step 4.

   IMPORTANT! For 3Dx800 sensors, applying an intrusion policy can take a significant amount of time (often greater than 10 minutes), especially the first time you apply a policy after updating the SEU on the Defense Center. Subsequent policy applications can be considerably faster.

3. You have the following choices:
   - To apply your policy to the targeted detection engines, click **Quick Apply**.

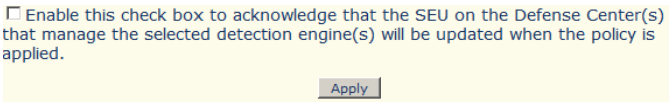
     The page refreshes and the apply task is queued. You can exit this procedure.
   - To return to the Intrusion Policy page without applying the policy, click **Cancel**. To apply your policy, repeat this procedure at a later time.
   - To manually select the detection engines where you want to apply your policy, click **Custom Apply**.

     The Apply Policy page appears.
4. Optionally, select a sort order from the drop-down list. You can choose from the following options:
   - **By Group** (that is, by detection engine group)
   - **By Sensor** (for the Defense Center only)
   - **By Policy** (that is, by the currently applied policy)
   - **By DE Type** (that is, by detection engine type: IPS or RNA)
   - **By Set Type** (that is, by interface set type: passive, inline, or inline with fail open)

Keysight, EX-1004
Page 349

5. In the **Available Detection Engines** list, select the detection engines where you want to apply the policy. You have two options:
   - To apply the policy to all detection engines within a specific category, select the check box for the category. For example, if you want to apply a policy to all the inline detection engines, sort by Set Type, then select the check box next to its label.
   - To apply the policy to unrelated detection engines, select the check boxes next to those detection engines.

   Requires: MDC If you apply an intrusion policy from a Master Defense Center and there are any SEUs older than the Master Defense Center's residing on Defense Centers managing those detection engines, those SEUs are updated. Therefore, a warning message with check box appears. Acknowledge the message by clicking its check box. That activates the **Apply** button.



6. Click **Apply**.

   If you enabled the **Require confirmation before applying policy?** option in the system policy, you must confirm that you want to apply the intrusion policy.

   It can take several minutes for the policy to be updated and applied to your detection engines.

WARNING! For 3Dx800 sensors, applying an intrusion policy can take a significant amount of time (often greater than 10 minutes), especially the first time you apply a policy after updating the SEU on the Defense Center. Subsequent policy applications can be considerably faster.

Viewing an Intrusion Policy Report

Platform: DC/MDC/Standalone Sensor

Component: IPS

An intrusion policy report is a record of all enabled intrusion policy features and settings at a specific point in time. IPS 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.

For any intrusion policy you can access, you can generate an Intrusion Policy report from either the Intrusion Policy page (**Policy & Response > IPS > Intrusion Report**) or the Detection Engine page (**Operations > Configuration > Detection**

Keysight, EX-1004
Page 350

Engines). Both reports provide the same information, but from different points of view.

- Generate the report from the intrusion policy page to provide a list of all detection engines, and detection engine groups to which the intrusion policy is applied.
- Generate the report from the detection engine page to identify which intrusion policy is currently applied on that detection engine.

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 354.

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 Virtual IPS 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 SEU version, whether the base policy is locked to the current SEU, 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 341.
Targets	Identifies sensors, detection engines, and detection engine groups that are associated with the intrusion policy. See Managing Targets on page 368.
Variables	Lists any policy-specific variables. See Managing Variables on page 371.

Keysight, EX-1004
Page 351

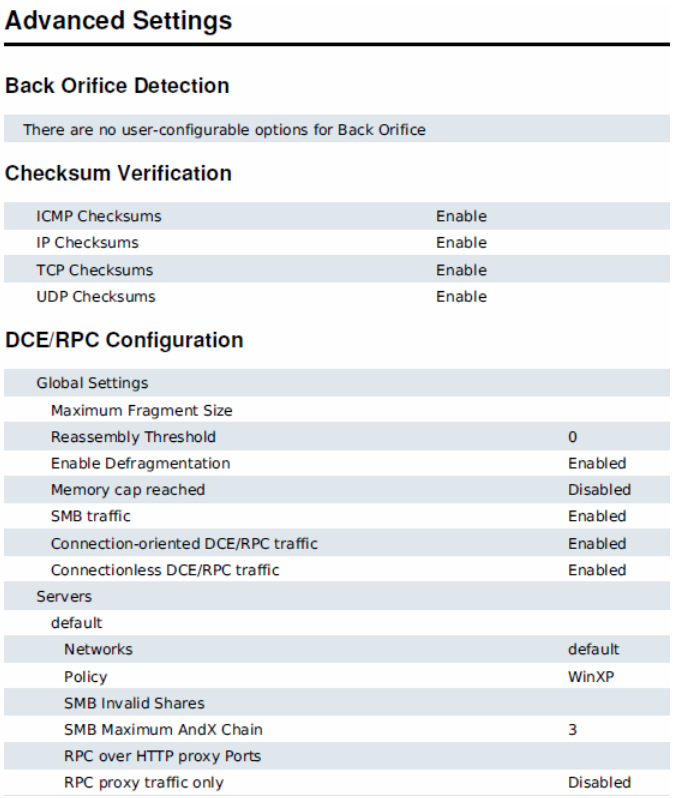
Intrusion Policy Report Sections (Continued)

Section	Description
RNA Recommendations	Provides information on any recommended rule states based on RNA data. Optionally, you can set your intrusion policy to **Include all differences between recommendations and rule states in policy reports**. See Managing RNA Rule State Recommendations on page 468.
Virtual IPS	Provides network or VLAN configuration details for filtered policies.See Defining VLANs and Subnetworks on page 393.
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 477.
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 403.

The following sample graphic displays the Advanced Settings section of an example intrusion policy report, and lists the configuration for each advanced

Keysight, EX-1004
Page 352

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 353

2. Click **Report** 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.

   IPS 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.

TIP! You can also create an intrusion policy report for a specific detection engine from the Detection Engines page (**Operations > Configuration > Detection Engines > Detection Engines**). For more information, see Understanding Detection Engines on page 199.

Comparing Two Intrusion Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can compare any two intrusion policies or two revisions of the same intrusion policy, for the intrusion policies you can access. You can then generate a 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 two intrusion policies or intrusion policy revisions in a side-by-side format; each policy or policy revision is identified by name in the title bar on the left and right sides of the comparison view.

  This allows you to view and navigate both policies or policy revisions in their entirety, 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.

  This allows you to identify and examine only the differences between two intrusion policy configurations.

For more information on understanding and using the intrusion policy comparison tools, see:

- Using the Intrusion Policy Comparison View on page 355
- Using the Intrusion Policy Comparison Report on page 357

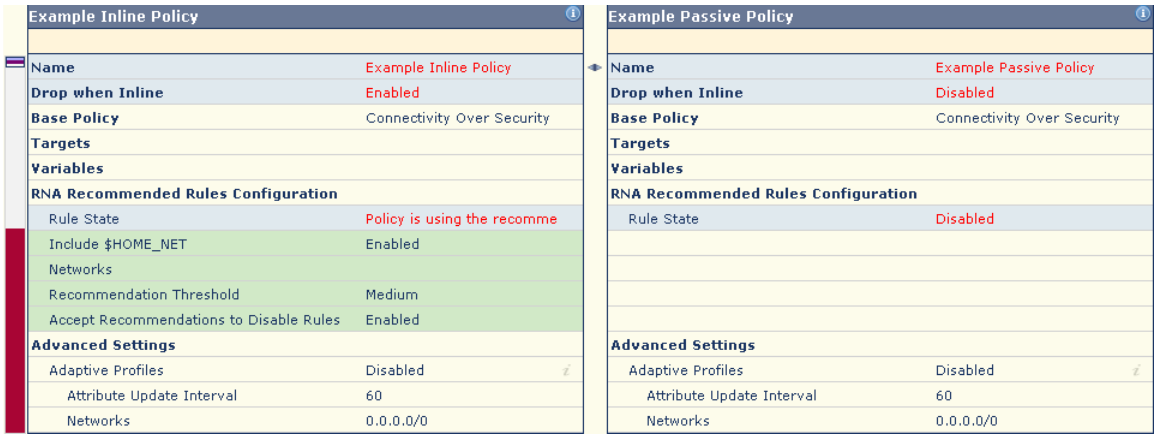
Keysight, EX-1004
Page 354

Using the Intrusion Policy Comparison View

Platform: DC/MDC/Standalone Sensor

Component: IPS

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. 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 Inline Policy	
Name	Example Inline Policy
Drop when Inline	Enabled
Base Policy	Connectivity Over Security
Targets	
Variables	
RNA Recommended Rules Configuration	
Rule State	Policy is using the recomme
Include $HOME_NET	Enabled
Networks	
Recommendation Threshold	Medium
Accept Recommendations to Disable Rules	Enabled
Advanced Settings	
Adaptive Profiles	Disabled
Attribute Update Interval	60
Networks	0.0.0.0/0

Example Passive Policy	
Name	Example Passive Policy
Drop when Inline	Disabled
Base Policy	Connectivity Over Security
Targets	
Variables	
RNA Recommended Rules Configuration	
Rule State	Disabled
Advanced Settings	
Adaptive Profiles	Disabled
Attribute Update Interval	60
Networks	0.0.0.0/0

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.

Keysight, EX-1004
Page 355

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.
navigate to an identified change	click a red line on the left scroll bar. The scope icon ( ) on the left scroll bar moves to the located change between two policies or policy revisions. A wide red line indicates contiguous changes.
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 494 for more information.
view information about the most recent modification	hover over the information icon ( ) on the title bar. The window displays 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 page displays the time a policy was last modified in local time, but the intrusion policy report lists the time modified in UTC.
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 document that lists only the differences between the two intrusion policies or intrusion policy versions.

Keysight, EX-1004
Page 356

Using the Intrusion Policy Comparison Report

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 in .pdf format. You can use this report to 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 Rules section of an intrusion policy comparison report, and lists the configuration for each rule for both intrusion policy configurations. Each section uses the same format and provides the same level of detail. Note that the Left Value and Right Value columns represent the policies or policy revisions you configured in the comparison view.

Rules

Policy Component	Left Value
ATTACK-RESPONSES command completed (1:494)	Drop and generate events
ATTACK-RESPONSES file copied ok (1:497)	Drop and generate events
ATTACK-RESPONSES id check returned root (1:498)	Drop and generate events
ATTACK-RESPONSES Microsoft cmd.exe banner (1:2123)	Drop and generate events
ATTACK-RESPONSES successful cross site scripting forced download attempt (1:2412)	Drop and generate events
ATTACK-RESPONSES successful gobbles ssh exploit GOBBLE (1:1810)	Drop and generate events
ATTACK-RESPONSES successful gobbles ssh exploit uname (1:1811)	Drop and generate events
ATTACK-RESPONSES successful kadmind buffer overflow attempt (1:1900)	Drop and generate events
ATTACK-RESPONSES successful kadmind buffer overflow attempt (1:1901)	Drop and generate events
BACKDOOR aol admin runtime detection (1:7104)	Drop and generate events
BACKDOOR aol admin runtime detection (1:7105)	Drop and generate events

Keysight, EX-1004
Page 357

To compare two intrusion policies or two revisions of the same policy:

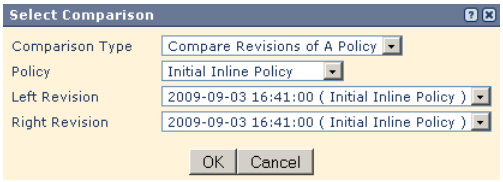
Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Compare Policies**.

   The **Select Comparison** window appears.



3. From the **Comparison Type**drop-down list, select whether you want to compare two different policies, or two revisions of the same policy. Remember to commit any potential 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 options:
   - If you are comparing two different policies, select the policies you want to compare from the **Left Policy** and **Right Policy** drop-down lists.
   - If you are comparing two revisions of the same policy, select the policies you want to compare from the **Left Revision** and **Right Revision** 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.

   IPS 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.

Setting Drop Behavior in an Inline Deployment

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 IPS handles drop rules in an inline deployment; see Setting Rule States on page 435 for information on setting rule states in your intrusion policy.

Keysight, EX-1004
Page 358

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, IPS 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 and reapply your policy.

When you set your intrusion policy to drop packets in an inline deployment, and you apply your policy to a detection engine on a Sourcefire 3D Sensor with an inline or inline with failopen interface set, IPS drops packets that trigger enabled drop rules and generates events for the triggered rules.

Note that in a passive deployment, including when an inline interface is in tap mode, IPS treats rules set to Drop and Generate Events the same as rules set to Generate Events; that is, IPS generates events but does not drop packets that trigger the rules regardless of the drop behavior of your policy. See Tap Mode on page 227 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 276 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 765.

Keysight, EX-1004
Page 359

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 792.

To set the drop behavior of your intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

3. Specify whether you want IPS 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 IPS does not drop packets in a passive deployment.

   TIP! On a 3Dx800 and 3D9900 sensor, IPS detection engines that use an inline or inline with fail open interface set can use *tap mode*, which allows you to use interface sets to passively monitor traffic.

4. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 360

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding the Base Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

The base policy in an intrusion policy defines the default settings for all rules and advanced settings in the policy. Any default or custom policy can be used as the base policy. Modifying a rule or advanced setting in your policy overrides the corresponding default setting.

Note that 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 494.

You can choose whether importing an SEU updates your base policy with modified rule or advanced settings in the SEU. However, modifications from imported SEUs do not override changes that you make to rules or advanced settings in your policy. See Importing SEUs and Rule Files on page 1848 and Allowing SEUs to Modify the Base Policy on page 363 for more information.

By default, the base policy for an intrusion policy is one of the default intrusion policies provided by the Sourcefire Vulnerability Research Team (VRT). The VRT determines all rule, preprocessor, and other advanced settings in these policies.You can also use a custom policy as your base policy.

See the following sections for more information:

- Using Default Intrusion Policies on page 362
- Using a Custom Base Policy on page 363
- Allowing SEUs to Modify the Base Policy on page 363
- Selecting the Base Policy on page 364
- Accepting Rule Setting Changes from a Custom Base Policy on page 366

Keysight, EX-1004
Page 361

Using Default Intrusion Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

Four default intrusion policies are delivered with the Sourcefire 3D System. The Sourcefire Vulnerability Research Team (VRT) sets the default state of all IPS features within each default policy provided by Sourcefire. The VRT sets the state of each intrusion, packet decoder, 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. IPS 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 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 a 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 policy with a tuned rule set and advanced settings configurations to inspect traffic in the way

Keysight, EX-1004
Page 362

that matters most to you. By doing this, you can improve both the performance of your sensor and your ability to respond effectively to the events it generates.

IMPORTANT! Because 3Dx800 sensors do not have a web interface, you must use the Defense Center to create and apply intrusion policies for these models.

Using a Custom Base Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use a custom policy, that is, a policy you have created, 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 the custom policy; however, your base policy does not include other settings such as variables, targets, or policies by VLAN or network. For example, consider the case where you create two intrusion policies, Policy A and Policy B, and you select Policy A as the base policy for Policy B. If you subsequently change the rule state for a rule in Policy A from Disable to Generate Events, the rule is set to Generate Events in the base policy for Policy B. However, if you create a policy-specific variable in Policy A, that policy-specific variable is not added to the base policy for Policy B; to use the variable in Policy B, you would have to recreate it in Policy B.

A modified rule or advanced setting in a custom policy becomes the corresponding default setting in a policy that uses the custom policy as its base policy. In this case, importing an SEU will not change the default in the policy using the custom policy even when the custom policy or the policy using it allows SEUs to modify its base policy. You can, however, override the default setting by modifying it in the policy using the custom policy.

Allowing SEUs to Modify the Base Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

SEUs that you import can contain new and updated packet decoder, preprocessor, and intrusion rules, modified states for existing rules, new IPS features such as new preprocessors and decoders, new options for existing IPS features, and modified option settings. SEUs can also delete rules, features, and feature options. See Importing SEUs and Rule Files on page 1848 for more information.

SEUs always modify the default policies provided by Sourcefire. You can choose whether to allow SEUs to update your base policy when a new SEU is installed to match the settings in the updated default policy. If you allow SEUs to update your base policy, a new SEU will not change settings you make that override settings in

Keysight, EX-1004
Page 363

your base policy. If you do not allow SEUs to update your base policy, you can manually update your base policy after importing one or more SEUs.

Consider the case of an intrusion policy in an inline deployment where a rule is disabled in your base policy and a new SEU sets the state of the rule to Generate Events. If you allow SEUs to change your base policy, the rule will be set to Generate Events in your base policy. However, if you do not allow SEUs to change your base policy, the rule continues to be disabled in your base policy. If you have manually set the rule to Drop and Generate Events and the policy is set to drop packets in an inline deployment, IPS will continue to drop packets that trigger the rule and generate events when the rule triggers, regardless of whether you allow SEUs to change your base policy.

Note that SEUs always delete rules that VRT deletes from default policies. If you do not allow SEUs to update your policy and a new SEU deletes a rule, the rule will be deleted the next time you apply your policy, regardless of whether you have set a rule state that overrides the base policy. The rule will remain in your policy and, depending the state you set, generate events if it is triggered until you reapply your policy.

New and updated advanced features such as preprocessors and their options are handled in the same way as rules. Consider the following cases:

- When a new option is added to a preprocessor, the option is enabled by default, and you do not allow SEUs to update your base policy, the new option would be added to your based policy and the option would be disabled.
- When a preprocessor option is removed if, for example, Sourcefire determines that the option no longer provides useful functionality, and you do not allow SEUs to update your base policy, the SEU would delete the option.

Selecting the Base Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can select the base policy for your intrusion policy and choose whether to allow SEUs 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.

Keysight, EX-1004
Page 364

To select the base policy in your intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 361 for more information.

5. Optionally, select or clear the **Update when a new SEU is installed** check box to specify whether you want new SEUs to update your base policy.

   When you save your changes with the check box cleared and then import an SEU, 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 SEU.

   See Allowing SEUs to Modify the Base Policy on page 363 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 403 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 477 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 477.

Keysight, EX-1004
Page 365

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Accepting Rule Setting Changes from a Custom Base Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 499 to accept settings for these rules in a policy where you have added layers.

Keysight, EX-1004
Page 366

To accept rule setting changes in a policy where you have not added layers:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** 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. To save the changes, click **Cancel**, open the other policy and commit changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon () appears when you add the first keyword or type the first character and disappears if you delete all characters.

     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.

   Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

8. Select **Inherit** from the Rule State drop-down.

Keysight, EX-1004
Page 367

9. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Managing Targets

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can target specific IPS detection engines in your intrusion policy so you can quickly apply your policy to the targeted detection engines each time you apply the policy. See Managing Detection Engines on page 206 and Using Detection Engine Groups on page 211 for more information.

You can specify any or all of the following types of targets in your policy:

- one or more detection engines
- one or more detection engine groups
- for a standalone sensor, all of the detection engines on the sensor
- for managed sensors, all of the detection engines on one or more sensors

You can then apply your policy to all targeted detection engines at once, or you can manually select the detection engines where you want to apply your policy.

Keysight, EX-1004
Page 368

You can also choose not to target detection engines in your policy, in which case you must manually select the detection engines where you want to apply the policy when you apply the policy.

You can later modify your policy by adding detection engine targets and deleting targeted detection engines.

IPS applies the intrusion policy to each detection engine only once if you target a detection engine more than once in the policy. For example, if you add an individual detection engine and you also add a group of detection engines that includes the individual detection engine, and then apply the policy, IPS will apply the policy to the detection engine only once.

To target detection engines in your intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

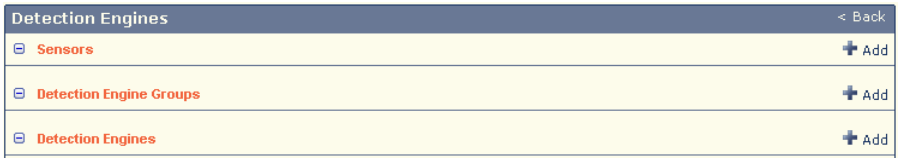
2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

3. Click **Manage Targets** on the Policy Information page.

   The **Detection Engines** page appears.



4. You have the following options:
   - Go to step 5 to add one or more of a type of item (sensors, detection engine groups, or detection engines).
   - Click the **+** or **-** sign to the left of an item type (sensors, detection engine groups, or detection engines) to expand or collapse a list of previously added targets of that type.
   - Click the delete icon ( ) next to a previously added target to delete the target.

Keysight, EX-1004
Page 369

5. Optionally, click **Add** next to the item (sensors, detection engine groups, or detection engines) you want to add.

   If there are no items (sensors, detection engine groups, or detection engines) configured on your system, a pop-up message appears.

   The figure below shows the Add Detection Engine(s) pop-up message.

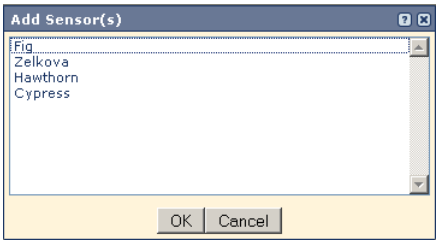


   You have the following choices:

   - To return to the Detection Engines page, click **OK** or **Cancel**.
   - To add managed sensors to a Defense Center, see, Adding Sensors to the Defense Center on page 122. After adding one or more sensors that you want to target in your policy, return to the beginning of this procedure.
   - To create detection engine groups on your system, see Creating Detection Engine Groups on page 211. After creating one or more detection engine groups that you want to target in your policy, return to the beginning of this procedure.
   - To create detection engines for your system, see Creating a Detection Engine on page 206. After creating one or more detection engines that you want to target in your policy, return to the beginning of this procedure.

   If the type of item (sensors, detection engine groups, or detection engines) that you want to target exists on your system, the pop-up window for the item appears.

   The figure below shows an example of the Add Sensor(s) pop-up window.



6. Click the item or items (sensors, detection engine groups, or detection engines) that you want to add and click **OK**, or click **Cancel**. Use Ctrl or Shift while clicking items to select multiple items or to clear individual items.

   The Detection Engines page appears.

7. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 370

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Managing Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 371

- 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 372 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 374 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 377 describes how to create new variables for use in rules and system configuration.
- Deleting Unused Variables on page 389 describes how to delete unused variables.
- Understanding Custom Variables on page 391 describes custom variables you can use to configure IPS capabilities not otherwise configurable via the web interface.
- Using Variables within Detection Engines on page 213 explains how you can use detection-engine-based variables to tailor your detection capabilities to more closely match your infrastructure.

Understanding Existing Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 398 for more information about acceptable syntax for IP addresses and ports in variables)

Keysight, EX-1004
Page 372

- 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 IPS 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 391.

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
$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

Keysight, EX-1004
Page 373

Variable Descriptions (Continued)

Variable Name	Description	Modify?
$SMTP_SERVERS	Defines SMTP servers on your network, and is used in rules that address exploits that target mail servers.	Yes, if you run SMTP servers
$SNMP_SERVERS	Defines SNMP servers on your network, and is used in rules that scan for attacks on SNMP servers.	Yes, if you run SNMP servers
$SQL_SERVERS	Defines database servers on your network, and is used in rules that address database-targeted exploits.	Yes, if you run SQL 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 377 and Creating New Variables for Detection Engines on page 217.

IMPORTANT! For more information about pre-configured variables and what they are used for, see Understanding Existing Variables on page 372.

To modify an existing variable:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 374

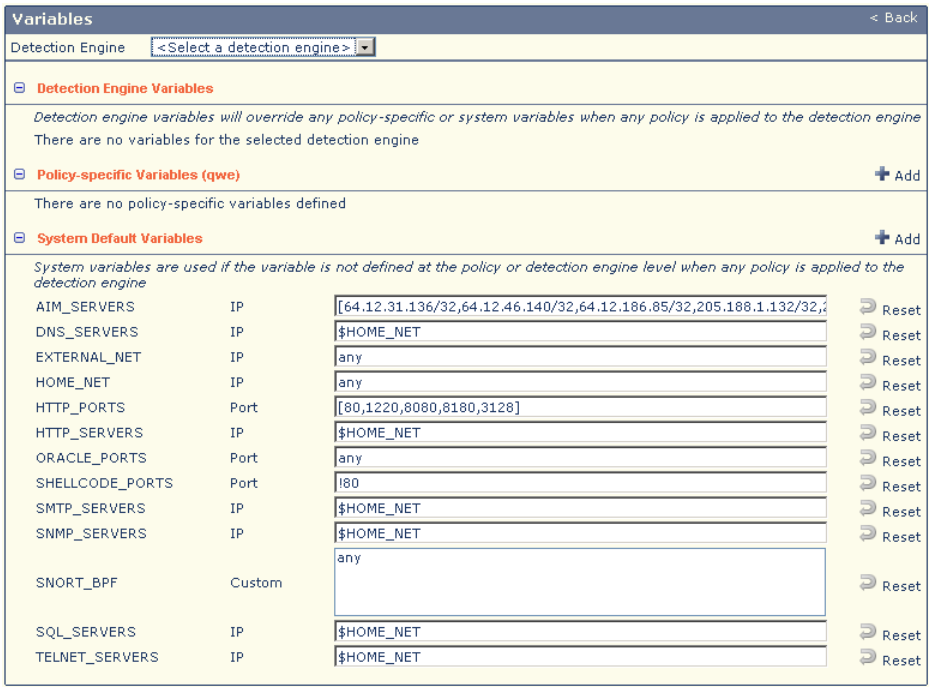
2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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. Optionally, click the **+** sign to the left of **System Default Variables** to display your system default variables.



4. You have the following options:

   - Select a detection engine from the Detection Engine drop-down list to specify the detection engine where you want to modify a variable and to display any variables already configured for that detection engine.
   - 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.

Keysight, EX-1004
Page 375

- 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 377 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 detection engine variable or 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 399.
   - If you are modifying a port-based variable, use the syntax described in Defining Ports in Variables and Rules on page 401.
   - If you are modifying a custom variable, use the syntax described in the Custom Variables table on page 391.

   TIP! For a description of predefined system default variables, see the Variable Descriptions table on page 373.

6. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 376

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Creating New Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 3D Sensor or Defense Center where you create your intrusion policy

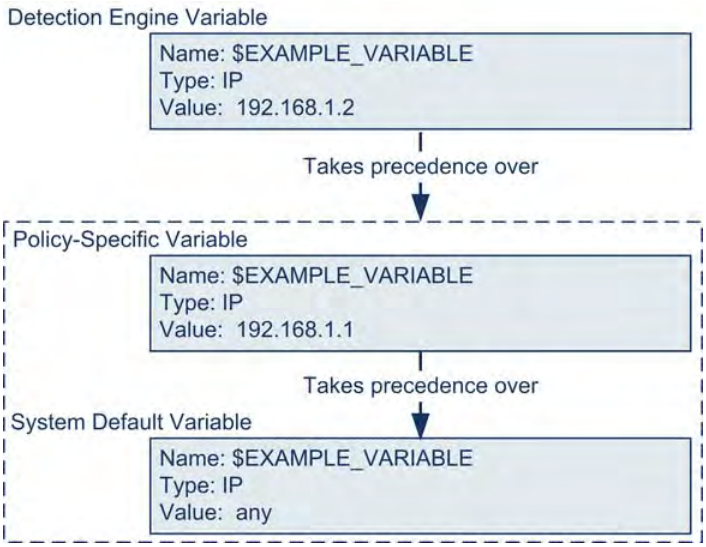
  Note that the system exports custom system default variables and their values when you export an intrusion policy. See Importing and Exporting Objects on page 1997 for more information.
- policy-specific variables, which are specific to the policy in which you create them
- detection-engine-specific variables, which are specific to the detection engine where you apply the policy

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 detection-engine-specific variables, which reside on the detection engine where you create them, and system default variables, which reside on the appliance where you create the policy. You can create and modify all three types of variables from within an intrusion policy.

In the case of detection-engine-specific and policy-specific variables, you can also use an existing variable of any other type to specify a new detection-engine-specific or policy-specific value for the same variable. As shown in the following figure, a detection-engine-specific value takes precedence over both policy-specific and system default values for the same variable in an

Keysight, EX-1004
Page 377

intrusion policy, and a policy-specific value takes precedence over the value for the system default variable.



All detection-engine-specific, policy-specific, and system default variables use the same syntax and must follow the same guidelines.

TIP! You can limit a variable definition to VLANs or subnetworks by defining the variable in your filtered or non-filtered policy, so long as you do not define a detection engine variable with the same name on the detection engine where you apply your non-filtered and filtered policies. See Limiting Variables to VLANs and Subnetworks on page 398 for an example variable configuration.

You can create nested variables so long as the nesting is not circular. 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:

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.

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]`.

Keysight, EX-1004
Page 378

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 Detection Engine Variables on page 379
- Creating New Policy-Specific Variables on page 383
- Creating New System Default Variables on page 386

Creating New Detection Engine Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

Detection engine variables set explicit variable values for individual detection engines. You can add a new detection-engine-specific variable from the Detection Engine Variables section on the Variables page. You can also set a detection-engine-specific value using an existing system default or policy-specific variable.

Keysight, EX-1004
Page 379

When you create a new detection engine variable from within an intrusion policy, IPS also creates a system default variable with the same name and a value of any and makes the new system default variable available by listing it in all intrusion policies on the system.

You can also create new detection engine variables and specify explicit detection engine values for system default variables on the detection engine Variable page (select **Operations > Detection Engines > Detection Engines** and click **Variables**). See Using Variables within Detection Engines on page 213 and Creating New Variables for Detection Engines on page 217 for more information.

Note that you can display detection engine variables using the web interface on a managed sensor, because detection engine variables are specific to the detection engine on the sensor. However, you cannot display policy-specific or system default variables used by an intrusion policy that you apply to a managed sensor.

To create a new detection engine variable:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

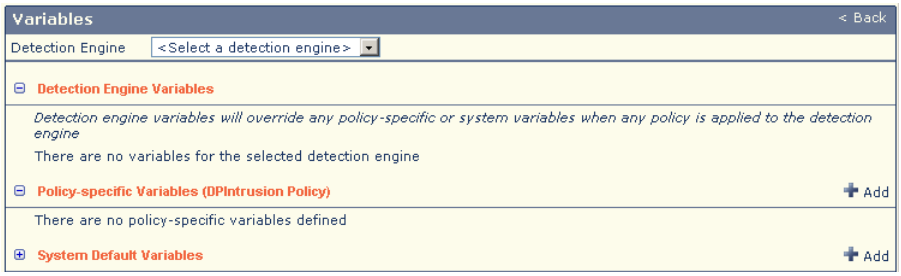
2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.



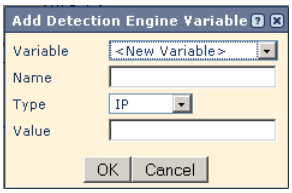
4. From the Detection Engine drop-down list, select the detection engine where you want to add a variable.

   Any existing variables for that detection engine appear in the Detection Engine Variables section.

Keysight, EX-1004
Page 380

5. Click **Add** in the Detection Engine Variables section.

   The Add Detection Engine Variable pop-up window appears.



6. You have two choices:

   - To create a new detection-engine-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 specify a detection-engine-specific value using an existing system default or policy-specific variable, select the name of the existing variable from the Variable drop-down list.

     The **Type** field is constrained. Go to step 8.

7. If you created a new variable in step 6, select `IP`, `Port`, or `Custom` from the **Type** drop-down list:

   - See Defining IP Addresses in Variables and Rules on page 399 for more information if you are defining a IP address-based variable.
   - See Defining Ports in Variables and Rules on page 401 for more information if you are defining a port-based variable.
   - See Understanding Custom Variables on page 391 if you are defining a special-purpose custom variable with one of the reserved variable names described in the Custom Variables table on page 391.

     When the new variable is a custom variable, the **Value** field expands.

8. Enter the new variable **Value**:

   - If the value describes IP addresses, see Defining IP Addresses in Variables and Rules on page 399 for information about the syntax you can use.
   - If the value describes port numbers, see Defining Ports in Variables and Rules on page 401 for information about the syntax you can use.

Keysight, EX-1004
Page 381

- 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 391.

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.

9. Click **OK**.

   The new variable is created. If a corresponding system default variable did not exist, IPS adds a system default variable with the same name and a value of any to the appliance and lists the new system default variable in all intrusion policies on the appliance. The new system default variable name appears in italicized text in the current policy to indicate that it is overridden by the new detection engine variable.

10. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 382

Creating New Policy-Specific Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 or detection engine variable. Note that if you apply your policy to the same detection engine whose variable you used to create the policy-specific variable, the detection engine value for the variable will override the policy-specific value.

When you add a new policy-specific variable to an intrusion policy, IPS 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

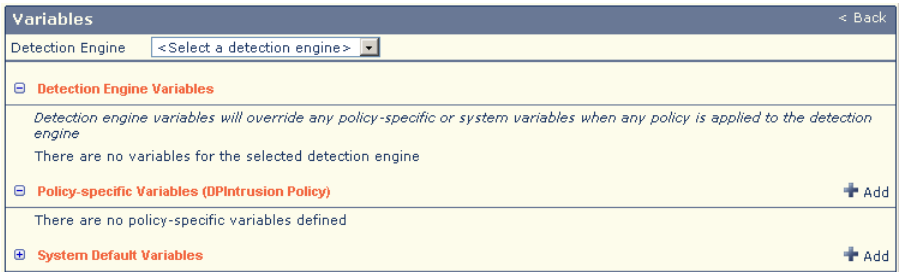
   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 383

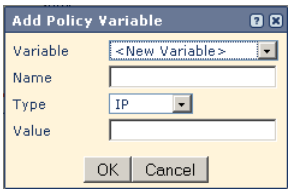
3. Click **Manage Variables** on the Policy Information page.

   The Variables page appears.



4. Click **Add** in the Policy-Specific Variables section.

   The Add Policy Variable pop-up window appears.



5. 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 or detection engine variable, select the name of the existing variable from the Variable drop-down list.

     The **Type** field is constrained. Go to step 7.

Keysight, EX-1004
Page 384

6. 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 399 for more information if you are defining a IP address-based variable.
   - See Defining Ports in Variables and Rules on page 401 for more information if you are defining a port-based variable.
   - See Understanding Custom Variables on page 391 if you are defining a special-purpose custom variable with one of the reserved variable names described in the Custom Variables table on page 391.

     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 399 for information about the syntax you can use.
   - If the value describes port numbers, see Defining Ports in Variables and Rules on page 401 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 391.

   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. The system default variable name appears in italicized text in the current policy to indicate that it is overridden by the new detection engine variable.
9. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 385

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Creating New System Default Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

Each intrusion policy lists the existing, pre-configured system default variables described in the Variable Descriptions table on page 373. 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 3D Sensor or Defense Center where you added the variables. When you enable an intrusion rule that uses a particular variable, IPS uses the system default value for the variable unless it is overridden in the policy by a policy-specific value, or the policy is applied to a detection engine with an explicit detection-engine-specific value for the variable.

TIP! IPS includes system default variables that you create when you export an intrusion policy that uses them.

To create a new system default variable:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 386

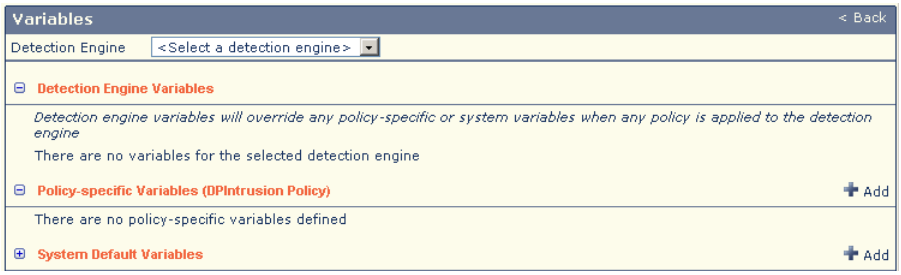
2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

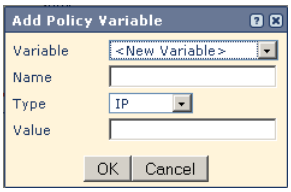
3. Click **Manage Variables** on the Policy Information page.

   The Variables page appears.



4. Click **Add** in the System Default Variables section.

   The 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 or detection-engine-specific variable to set a value for a system default variable, because IPS automatically creates a new system default variable with a value of any when you create a new policy-specific or detection-engine-specific variable. See Modifying Variables on page 374 for information on modifying the value of a system default variable.

Keysight, EX-1004
Page 387

6. Select `IP`, `Port`, or `Custom` from the **Type** drop-down list.
   - See Defining IP Addresses in Variables and Rules on page 399 for more information if you are defining a IP address-based variable.
   - See Defining Ports in Variables and Rules on page 401 for more information if you are defining a port-based variable.
   - See Understanding Custom Variables on page 391 if you are defining a special-purpose custom variable with one of the reserved variable names described in the Custom Variables table on page 391.

     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 399 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 399 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 391.

   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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 388

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Deleting Unused Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can delete detection-engine-specific and policy-specific variables from within an intrusion policy. You cannot delete a variable within a policy when the variable 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 from within an intrusion policy. You can delete system default variables with predefined default values on the detection engine Variable List page (select **Operations > Detection Engines > Detection Engines** and click **Variables**), but only if they are not used in any active or inactive rule in any policy on the system. See Deleting and Resetting Variables on page 218 for more information.

To delete a variable:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 389

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 Detection Engine Variables, 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. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 390

Understanding Custom Variables

Platform: DC/MDC/Standalone Sensor

Component: IPS

Custom variables allow you to configure features that you cannot otherwise configure via the web interface. The Custom Variables table describes the custom variables that are available.

Custom Variables

Variable	Description	Value for Variable Binding
SNORT_BPF	Allows you to apply a Berkeley Packet Filter (BPF) to filter traffic before it reaches IPS. Note that with or without BPF, IPS filters traffic internally so that only traffic covered by intrusion rules and preprocessors is processed.	Use the same standard BPF instruction set as accepted by tcpdump.
USER_CONF	Provides a general tool that allows you to configure one or more features not otherwise available via the web interface. When you apply an intrusion policy, IPS stores the configuration settings in the policy on the sensor in a file named `snort.conf`. IPS then implements the settings in `snort.conf` and, if they exist, the settings in `USER_CONF`. If `USER_CONF` does not exist, IPS implements any settings in the file `user.conf`. **WARNING! Do not** use the reserved variable `USER_CONF` to configure an IPS feature unless you are instructed to do so in the feature description or by Sourcefire Support. Conflicting or duplicate configurations will halt IPS.	Use the feature-specific directives provided by Sourcefire Support or in the feature description. You can type up to 4096 total characters on a single line; the line wraps automatically. You can configure any number of features up to the maximum 8192 characters in a variable, and can use the backslash (\) line continuation character after any complete argument in a command directive.

You create a custom variable using a specific assigned name and a set of instructions appropriate to the function the variable provides. IPS 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 system default, policy-specific, or detection-engine-specific variable. See Creating New Variables on page 377 for more information.

Keysight, EX-1004
Page 391

You can also create a detection-engine-specific custom variable from the detection engine Variable List page. See Configuring Custom Variables in Detection Engines on page 219 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 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. You can also modify a custom variable and remove or add instructions as needed. See the Value for Variable Binding column in the Custom Variables table on page 391 for the syntax to use when defining each type of custom variable.

Configuring Custom Variables

You can set an explicit value for the predefined system default SNORT_BPF custom variable, or you can create a detection-engine-specific or policy-specific variable based on the existing system default variable.

You can create a new USER_CONF variable as a detection engine, policy-specific, or system default variable using the reserved name USER_CONF.

Keysight, EX-1004
Page 392

To configure the SNORT_BPF custom variable in an intrusion policy:

Access: P&R Admin/Admin

▶ You have the following options:

- To set an explicit value for SNORT_BPF, type 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 374 for more information.
- To specify a detection-engine-specific or policy-specific value for SNORT_BPF using the existing system default variable, see Creating New Detection Engine Variables on page 379 or Creating New Policy-Specific Variables on page 383.

To configure the USER_CONF custom variable in an intrusion policy:

Access: P&R Admin/Admin

▶ You have the following options:

- To create USER_CONF as a new detection-engine-specific variable using the reserved name USER_CONF, see Creating New Detection Engine Variables on page 379.
- To create USER_CONF as a new policy-specific variable using the reserved name USER_CONF, see Creating New Policy-Specific Variables on page 383.
- To create USER_CONF as a new system default variable using the reserved name USER_CONF, see Creating New System Default Variables on page 386.

Defining VLANs and Subnetworks

Platform: DC/MDC/Standalone Sensor

Component: IPS

On most Sourcefire 3D Sensors, you can apply one or more custom intrusion policies that are filtered to monitor VLAN or subnetwork traffic on the network monitored by the detection engine where you apply the policy.

This capability enables you to customize intrusion policies to inspect traffic differently for distinct network segments or VLANs. For example, traffic on one network segment (for example, VIP/SIP) may be very different from traffic on another network segment (HTTP, FTP). VLAN and subnetwork filtering allows you

Keysight, EX-1004
Page 393

to apply multiple intrusion policies to a single detection engine, and prevents the mixing of intrusion events from multiple network segments.

TIP! You can limit a variable definition to the VLANs or subnetworks defined in your filtered intrusion policy. For more information, see Limiting Variables to VLANs and Subnetworks on page 398.

Before you apply a filtered policy, you must apply a non-filtered policy to the detection engine. After you apply both policies, IPS determines which policy to use to examine each packet in the following order:

- First, IPS determines whether a VLAN ID is present in each packet, and if it is, uses the innermost VLAN ID to determine which intrusion policy to use to examine the packet.
- If the packet does not have an associated VLAN ID, IPS determines whether to use a filtered policy based on first the destination IP address in the packet, and then the source IP address. Note that this means that IPS may use different policies to examine traffic in a single session, depending on how you configured the filtered policies and on the direction of the traffic.
- All traffic not filtered by a VLAN- or subnetwork-filtered policy is examined by the non-filtered policy.

Just as you must create and apply a non-filtered policy before you apply a filtered policy, you must delete any filtered policies applied to that detection engine before you delete a non-filtered policy from a detection engine. See Understanding Detection Engines on page 199 for information on deleting intrusion policies.

VLAN and subnetwork filtering has the following limitations:

- Although you can apply multiple VLAN-filtered policies or multiple subnetwork-filtered policies to a detection engine, you cannot mix the two policy types.
- You cannot apply multiple overlapping filtered policies. That is, you cannot apply filtered policies with overlapping networks or overlapping VLAN ID ranges to a detection engine.

Note that you can use the **Reapply All** or **Reapply Intrusion Policies** option on the Detection Engines page (**Operations > Configuration > Detection Engines**) to reapply your intrusion policies. These options can be convenient, for example, if you import an SEU and do not automatically reapply your intrusion policies after the import completes. These options conveniently reapply filtered and non-filtered intrusion policies simultaneously for each detection engine. See Understanding Detection Engines on page 199 and Importing SEUs and Rule Files on page 1848 for more information.

The Filtered Policies by Model table lists the maximum number of VLAN- or subnetwork-filtered policies you can apply to detection engines on different sensors in addition to the non-filtered intrusion policy that you must apply. In the

Keysight, EX-1004
Page 394

table, a setting of `0` indicates that you cannot apply VLAN-filtered policies to that type of sensor because of physical hardware limitations.

Filtered Policies by Model

Model	Maximum per Sensor
3D500	0
3D1000	1
3D2000	3
3D2100	3
3D2500	9
3D3000	9
3D3500	19
3D3800	0
3D4500	49
3D5800	9
3D6500	99
3D8140	66
3D8250	99
3D9800 (800MB RAM)	3
3D9800 (1.2GB RAM)	7
3D9900	99
Crossbeam X-Series APM (2 GB RAM)	16
Crossbeam X-Series APM (4 GB RAM)	33
Crossbeam X-Series APM (8 GB RAM)	67

Keysight, EX-1004
Page 395

Filtered Policies by Model (Continued)

Model	Maximum per Sensor
Crossbeam X-Series APM (16 GB or more RAM)	99
Virtual 3D Sensor	3

To configure the VLAN or subnetwork for a filtered custom intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

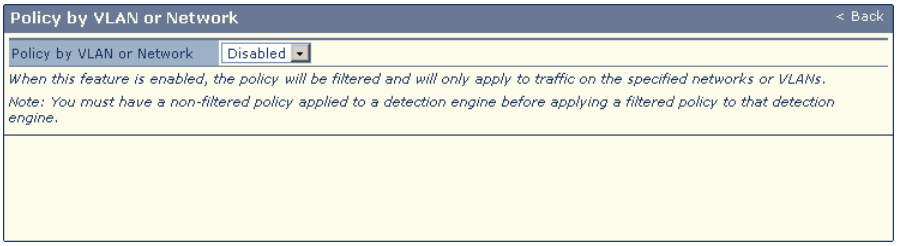
2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

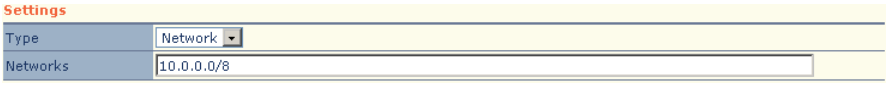
3. Click **Policy by VLAN or Network** under Policy Information in the navigation panel on the left.

   The Policy by VLAN or Network page appears.



4. From the **Policy by VLAN or Network** drop-down list, select **Enabled**.

   The Settings section appears. By default, the web interface displays the settings for filtering by subnetwork.



> **TIP!** To disable VLAN or subnetwork filtering, select **Disabled**.

5. From the **Type** drop-down list, select whether you want to filter the intrusion policy by network or by VLAN.

Keysight, EX-1004
Page 396

6. You have two options, depending on how you chose to filter the policy in the previous step:

   - If you are filtering by network, specify the network or networks you want to monitor with your filtered policy.

     You can enter a single IP address or CIDR block, or a comma-separated list comprised of either or both, including negation. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     Except as limited by the maximum number of policies allowed on your Sourcefire 3D Sensor, you can specify up to 512 IP addresses, CIDR blocks, or a combination of both per network policy, and you can specify up to 4096 of these in all network-filtered policies applied to the same detection engine. See the Filtered Policies by Model table on page 395 for the maximum number of policies allowed for each 3D Sensor model.

   - If you are filtering by VLAN, specify the VLAN IDs.

     You can enter a single VLAN ID, a VLAN ID range (using a hyphen), or a a comma-separated list comprised of either or both, including negation.

   TIP! Ensure that the IP addresses in your VLAN or subnetwork do not overlap, that is, that IP addresses in your VLAN or subnetwork are unique, before enabling adaptive profiles or RNA recommended rules in a filtered or non-filtered policy. See Managing RNA Rule State Recommendations on page 468 and Using Adaptive Profiles on page 697 for more information.

7. You have three options:

   - To leave your changes in the system cache, exit your intrusion policy.

     Note that you must commit or discard your changes before editing another policy when you are logged in as the same user.

   - To save your changes, click the policy name to return to the Policy Information page, then click **Commit Changes** and respond to any prompts that appear, clicking **OK** as needed.

     Your changes are saved and the Intrusion Policy page appears.

     You must apply the policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

   - To discard your changes, click the policy name to return to the Policy Information page, then click **Discard Changes**. At the prompt, click **OK.**

     Your changes are discarded and the Intrusion Policy page is displayed.

   See Committing Intrusion Policy Changes on page 345 for information on caching, committing, or discarding changes to your policy.

Keysight, EX-1004
Page 397

Limiting Variables to VLANs and Subnetworks

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can limit a variable definition to the VLANs or subnetworks defined in your filtered intrusion policy by defining the variable in your filtered or non-filtered policy; in doing so, you cannot define a detection engine variable with the same name on the detection engine where you apply your non-filtered and filtered policies.

Note that the value for a detection engine variable always overrides the value of a policy-specific or system default variable with the same name. See Modifying Variables on page 374 and Creating New Variables on page 377 for more information.

For example, if you create a $HOME_NET variable for your detection engine, it will always override any value for $HOME_NET that you set in a policy-specific variable or a system default variable.

However, consider the following case, where:

- you have not set a detection-engine-specific value for $HOME_NET
- the policy-specific value for $HOME_NET in your non-flitered policy is 10.4.0.0/16
- the policy-specific value for $HOME_NET in your filtered policy is 10.5.0.0/16

In this example case, the value for $HOME_NET is 10.5.0.0./16 for your VLAN traffic and 10.4.0.0./16 for all other traffic.

Defining IP Addresses and Ports for Your Network

Platform: DC/MDC/Standalone Sensor

Component: IPS

When tuning your system, you can define IP addresses and ports in both 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)

Keysight, EX-1004
Page 398

- SQL database servers
- SMTP servers
- SNMP 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

For details on defining variables, see Managing Variables on page 371. For details on existing variables. see Understanding Existing Variables on page 372.

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 764.

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 399 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 401 describes the syntax you can use to define single and multiple ports in variables and rules.

Defining IP Addresses in Variables and Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

When configuring variables and writing standard text rules, you can specify both IP addresses and port numbers. You can specify values in a variety of ways, depending on your needs. You can enter a single IP address directly, or you can use any, IP address lists, or CIDR notation. Additionally, you can indicate that you want to exclude a specific IP address or set of IP addresses.

See the following sections for more information on how to specify IP addresses:

- Specifying Any IP Address on page 400
- Specifying Multiple IP Addresses on page 400
- Excluding IP Addresses on page 400

Keysight, EX-1004
Page 399

Specifying Any IP Address

You can specify that the rule evaluate packets with any source IP address and any destination IP address using the argument **any**. For example:

```
alert tcp any any -> any any
```

Within the Sourcefire 3D System web interface, you can specify this by typing the word `any` in the **Source IP** or **Destination IP** field.

Specifying Multiple IP Addresses

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]
```

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 400 for more information.

You can also use Classless Inter-Domain Routing (CIDR) notation to specify IP address ranges. For example, `192.168.1.0/24` specifies the IP 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 59.

TIP! If you need to specify a range of IP addresses but cannot express it using CIDR notation alone, you can use CIDR blocks in an IP address list.

Excluding IP Addresses

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 400

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 IP 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.

Defining Ports in Variables and Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 766.

You can list ports by separating the ports with commas, as show in the following example:

```
80, 8080, 8138, 8600-9000, !8650-8675
```

Optionally, as shown in the following example, 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]
```

The following list describes the syntax you can use.

Keysight, EX-1004
Page 401

Any Port

Specify `any` to specify any port.

For example: `any`

Port Number

Specify a port number to specifies a single port.

For example: `80`

Port Ranges

Specify *first_port-last_port* to specify a range of ports, where *`first_port`* represents the first port in the range, and *`last_port`* specifies the last port.

For example: `20-21`

Specify *`-port`* to specify all ports less than or equal to *`port`*.

For example: `-21`

Specify *`port-`* to specify all ports greater than or equal to *`port`*.

For example: `80-`

Note that you must separate port ranges with a dash (-) in version 4.9 or later. Port ranges indicated with a colon (:) are supported for backward compatibility, but you can no longer use a colon in port variables that you create.

Negation

Specify `!` to specify all ports except the specified port, port list, or range of ports. Note that you can logically use negation with all port designations except *`any`*, which if negated would indicate *no port*.

For example: `!80`

Keysight, EX-1004
Page 402

Chapter 10
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 IPS 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 358 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.

On a system using both the RNA and IPS components, you can generate rule state recommendations based on vulnerabilities associated with the hosts and services on your network and, optionally, update rules to reflect the recommended states.

Keysight, EX-1004
Page 403

See the following sections for more information:

- Understanding IPS Rule Types on page 404 describes the intrusion rules and preprocessor rules you can view and configure in an intrusion policy.
- Viewing Rules in an Intrusion Policy on page 406 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 417 describes how you can use rule filters to find the rules for which you want to apply rule settings.
- Setting Rule States on page 435 describes how to enable and disable rules from the Rules page.
- Filtering Intrusion Event Notification Per Policy on page 439 explains how to set event filtering thresholds for specific rules and set suppression on specific rules.
- Adding Dynamic Rule States on page 452 explains how to set rule states that trigger dynamically when rate anomalies are detected in matching traffic.
- Adding Alerts on page 458 describes how to associate SNMP or OPSEC alerts with specific rules.
- Automatically Enabling Advanced Settings on page 490 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 466 describes how to add comments to rules in an intrusion policy.
- Managing RNA Rule State Recommendations on page 468 describes how to generate rule state recommendations based on vulnerabilities associated with the hosts and services on your network.
- Using Layers in an Intrusion Policy on page 494 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.

Note that shared object rules, standard text rules, and preprocessor rules appear on the Rules page.

Understanding IPS Rule Types

The IPS component includes two types of rules:

- intrusion rules
- preprocessor rules

An intrusion rule is a specified set of keywords and arguments on a 3D Sensor with the IPS component; an intrusion rule detects attempts to exploit vulnerabilities on your network by analyzing network traffic to check if it matches

Keysight, EX-1004
Page 404

the criteria in the rule. IPS 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 IPS component of the Sourcefire 3D 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. See Modifying Existing Rules on page 864 for more information.

You can use the rule editor to create your own custom standard text rules. You can also import local standard text rules that you create on a local machine. See Understanding and Writing Intrusion Rules on page 761 and Importing Local Rule Files on page 1857 for more information.

The IPS component also includes preprocessor rules, which are rules associated with preprocessor and packet decoder detection options included in the IPS component. You cannot copy or edit preprocessor rules.

The Sourcefire VRT determines the default rule states of Sourcefire's shared object rules, standard text rules, and preprocessor rules for each default policy included in the IPS component. Note that 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 IPS to generate events or drop offending packets and generate events for preprocessor rules.

The Rule Types table describes each type of rule included in the IPS component of 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 864 for more information.

Keysight, EX-1004
Page 405

Rule Types (Continued)

Type	Description
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. 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 in the IPS component of the Sourcefire 3D System. You must enable preprocessor rules if you want them to generate events. Preprocessor rules have a preprocessor-specific GID (generator ID). See the Generator IDs table on page 489 for more information.

Viewing Rules in an Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 417
- the rule attribute menus - for more information, see Setting Rule States on page 435, Filtering Intrusion Event Notification Per Policy on page 439, Adding Dynamic Rule States on page 452, Adding Alerts on page 458, and Adding Rule Comments on page 466
- the rules listing - for more information, see the Rules Page Columns table on page 407
- the rule details - for more information, see Viewing Rule Details on page 410

You can also sort rules by different criteria; for more information, see Sorting the Rule Display on page 409.

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.

Keysight, EX-1004
Page 406

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 488
SID	Integer which indicates the Snort ID (SID), which acts a unique identifier for the rule.	Reading Preprocessor Generator IDs on page 488
Message	Message included in events generated by this rule, which also acts as the name of the rule.	Defining the Event Message on page 772
	The rule state for the rule, which may be one of four states: • Drop and Generate Events () • Generate Events () • Disable () • Inherit (blank) Note that you can access the Set rule state dialog box for a rule by clicking on its rule state icon.	Setting Rule States on page 435
	RNA recommended rule state for the rule.	Managing RNA Rule State Recommendations on page 468
	Event filter, including event thresholds and event suppression, applied to the rule.	Filtering Intrusion Event Notification Per Policy on page 439
	Dynamic rule state for the rule, which goes into effect if specified rate anomalies occur.	Adding Dynamic Rule States on page 452

Keysight, EX-1004
Page 407

Rules Page Columns (Continued)

Heading	Description	For more information, see...
	Alerts configured for the rule, including SNMP and OPSEC alerts.	Adding Alerts on page 458
	Comments added to the rule.	Adding Rule Comments on page 466

You can also use the policy 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 494 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 361 for information on the base policy.

To view the rules in an intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

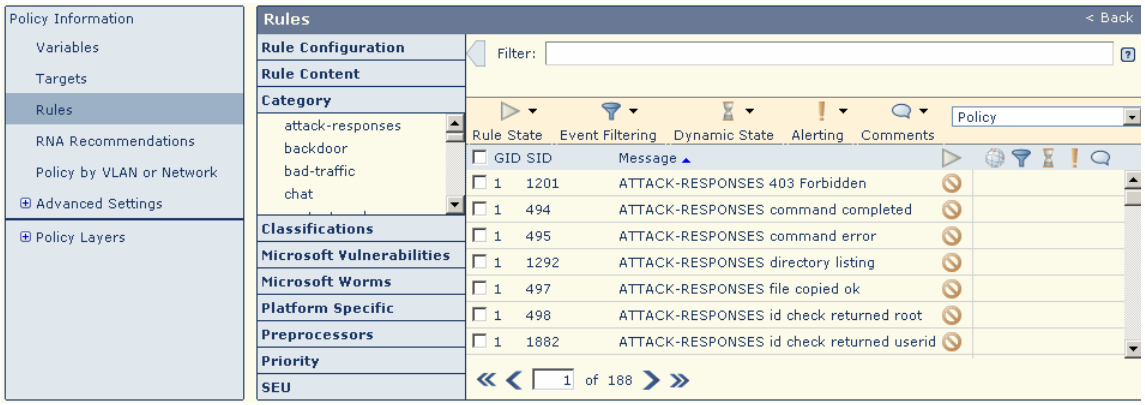
   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 408

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 409

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can view rule documentation, RNA 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 an RNA 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 293
Rule State	The current rule state for the rule. Also indicates the layer where the rule state is set.	Setting Rule States on page 435; Using Layers in an Intrusion Policy on page 494
RNA Recommendation	If RNA recommendations have been generated, the recommended rule state for the rule.	Managing RNA Rule State Recommendations on page 468
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 471
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 412

Keysight, EX-1004
Page 410

Rule Details (Continued)

Item	Description	For more information, see...
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 413
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 413
Alerts	Alerts currently set for this rule, as well as the facility to add an alert for the rule.	Setting an OPSEC Alert for a Rule on page 415 and Setting an SNMP Alert for a Rule on page 415
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 417
Documentation	The rule documentation for the rule, supplied by the Sourcefire Vulnerability Research Team (VRT).	Using Packet View Rule Actions on page 295

To view rule details:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

Keysight, EX-1004
Page 411

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can set a a single threshold for a rule from the Rule Details page. Adding a threshold overwrites any existing threshold for the rule. For more information on thresholding, see Configuring Event Thresholding on page 439.

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: P&R Admin/Admin

1. Click **Add** next to Thresholding.

   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, specify the number of event instances you want to use as your threshold.
5. In the **Seconds** field, specify the number of seconds that make up the time period for which event instances are tracked.
6. Click **OK**.

   IPS 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, IPS includes an indication over the icon of the number of event filters.

Keysight, EX-1004
Page 412

Setting Suppression for a Rule

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can set one or more suppressions for a rule from the Rule Details page. For more information on suppression, see Configuring Suppression Per Intrusion Policy on page 447.

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: P&R Admin/Admin

1. Click **Add** next to Suppression.

   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, CIDR block, or variable you want to specify as the source or destination IP address.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

4. Click **OK**.

   IPS 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, IPS includes an indication over the icon of the number of event filters.

Setting a Dynamic Rule State for a Rule

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can set one or more dynamic rule states for a rule from the Rule Details 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 453.

Keysight, EX-1004
Page 413

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: P&R Admin/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 for the detection engine.
3. Optionally, 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, CIDR block, variable, or a comma-separated list of these. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
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.

Keysight, EX-1004
Page 414

7. Click **OK**.

   IPS 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, IPS includes an indication over the icon of 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can set an SNMP alert for a rule from the Rule Details page. For more information on SNMP alerts, see Adding Alerts on page 458.

To add an SNMP alert from the rule details:

Access: P&R Admin/Admin

- Click **Add SNMP Alert** next to Alerts.

  IPS adds the alert and displays an alert icon ( ) next to the rule in the Alerting column. If you add multiple alerts to a rule, IPS includes an indication over the icon of the number of alerts.

Setting an OPSEC Alert for a Rule

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can set an OPSEC alert for a rule from the Rule Details page. For more information on OPSEC alerts, see Adding OPSEC Alerts on page 460.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

To add an OPSEC alert from the rule details:

Access: P&R Admin/Admin

1. Click **Add OPSEC Alert** next to Alerts.

   The Add OPSEC Alert dialog box appears.

2. From the **Firewalls Object** list, select the firewall object that receives this response.

   You can use the defaults of All, Gateways, or Localhost, or can create additional firewall objects using the SAM Control Panel. For details, refer to Adding and Deleting Firewall Objects on page 745.

Keysight, EX-1004
Page 415

3. From the **Logging** list, select one of the following log types.

Log Type	Description
None	Does not log SAM responses to this rule.
Long	Performs detailed logging, but does not store the event that caused the SAM response.
Long with Alert	Keeps more detailed logs along with the event that caused the SAM response.

4. From the **Firewall Action** list, select the action you want the firewall to take. You can select one of the following actions.

Action	Description
Notify	All connection attempts are logged as specified by the logging attribute. Traffic is not impeded.
Inhibit	All connection attempts are rejected. Existing connections will continue to function.
Inhibit and Close	All connection attempts are rejected. Existing connections are closed.
Drop	All connection attempts are dropped. Existing connections will continue to function.
Drop and Close	All connection attempts are dropped. Existing connections are closed.

5. In the **Timeout** field, specify the number of seconds that the specified action continues after it is initiated by the firewall.

IMPORTANT! The maximum value is 2,147,483 seconds. Specifying a value of 0 causes the action to continue until it is manually disabled by a firewall administrator.

6. From the **Filter Type** list, select the filter type you want to apply from the drop-down list.

   Sourcefire supports 15 different filter types. For details, refer to Understanding OPSEC Filter Options on page 464.

Keysight, EX-1004
Page 416

7. Depending on the filter type you selected, the **Network Mask** field may appear. If it appears, enter the appropriate network mask value (for example, 255.255.255.0).
8. To install the filter with reversed source and destination information from the packet that triggered the rule, check **Reverse Source Destination**.
9. Click **OK**.

   IPS adds the alert and displays an alert icon ( ) next to the rule in the Alerting column. If you add multiple alerts to a rule, IPS includes an indication over the icon of the number of alerts.

Adding a Rule Comment for a Rule

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can add a rule comment for a rule from the Rule Details page. For more information on rule comments, see Adding Rule Comments on page 466.

To add a comment from the rule details:

Access: P&R Admin/Admin

1. Click **Add** next to Comments.

   The Add Comment dialog box appears.
2. Type the rule comment.
3. Click **OK**.

   IPS adds the comment and displays a comment icon ( ) next to the rule in the Comments column. If you add multiple comments to a rule, IPS includes an indication over the icon of 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can filter the rules you display on the Rules page by a single criteria, or a combination of one or more criteria.

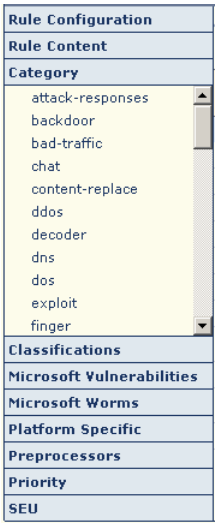
Keysight, EX-1004
Page 417

The filter you construct is shown in the Filter text box.



You can click keywords and keyword arguments in the filter panel to construct a filter. When you select multiple keywords, IPS combines them using AND logic to create a compound search filter. For example, if you select **preprocessor** under **Category** and **112** under **Rule Content > GID**, you get a filter of `Category: "preprocessor" GID:"112"` which retrieves all rules that are preprocessor rules **and** have a GID of 112.

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 **attack-responses** and **backdoor** from **Category** to produce the filter `Category:"backdoor,attack-responses"`, which retrieves any rules in the `attack-responses` category or in the `backdoor` category.



To show the filter panel, click the show icon( ).

To hide the filter panel, click the hide icon( ).

Note that you can view help on filtering by clicking the help icon ( ) next to the filter.

For more information, see the following topics:

- Understanding Rule Filtering in an Intrusion Policy on page 419
- Setting a Rule Filter in an Intrusion Policy on page 432

Keysight, EX-1004
Page 418

Understanding Rule Filtering in an Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS offers rule filter keywords to 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 419
- Understanding Rule Configuration Filters on page 422
- Understanding Rule Content Filters on page 425
- Understanding Rule Categories on page 427
- Editing a Rule Filter Directly on page 431

Guidelines for Constructing Intrusion Policy Rule Filters

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"`.

Keysight, EX-1004
Page 419

- When you select a filter type group heading that is a keyword (Category, Classifications, Microsoft Vulnerabilities, Microsoft Worms, Priority, and SEU), 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 **attack-responses** under Category in the filter panel, `Category:"attack-responses"` is added to the filter text box. If you then click **dos** under Category, the filter changes to `Category:"dos"`.
- 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 > CVS ID** in the filter panel, a pop-up window prompts you to supply the CVS 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 **attack-responses** under Category in the filter panel, `Category:"attack-responses"` is added to the filter text box. If you then click **MS00-006** under Category, the filter changes to `Category:"attack-responses" MicrosoftVulnerabilities:"MS00-006"`.
- When you select multiple keywords, IPS combines them using AND logic to create a compound search filter. For example, if you select **preprocessor** under **Category** and **112** under **Rule Content > GID**, you get a filter of `Category:“preprocessor” GID:”112”` which retrieves all rules that are preprocessor rules **and** have a GID of 112.
- 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 **attack-responses** and **backdoor** from **Category** to produce the filter `Category:"backdoor,attack-responses"`, which retrieves any rules in the `attack-responses` category or in the `backdoor` category.

Keysight, EX-1004
Page 420

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 SEU mechanism to add new 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 422.	No	A grouping	keywords
Rule Content	Finds rules according to the content of the rule. See Understanding Rule Content Filters on page 425.	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 427.	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 Specifying Rule Classifications in Event Searches on page 314 and Defining the Event Classification on page 772.	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

Keysight, EX-1004
Page 421

Rule Filter Groups (Continued)

Filter Group	Description	Multiple Argument Support?	Heading is...	Items in List are...
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.
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 484 and Setting Rule States on page 435 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.
SEU	Finds rules added through a specific SEU.	No	A keyword	arguments

Understanding Rule Configuration Filters

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 rule state recommended by RNA, you can filter on rule state by selecting **Does not match recommendation**.

Keysight, EX-1004
Page 422

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: P&R Admin/Admin

1. 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 that recommended by RNA, 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: P&R Admin/Admin

1. Click **Recommendation**.
2. Select the RNA rule state recommendation to filter by.

   Requires: DC + IPS + RNA The Rules page updates to display rules according to recommended rule state.

To use the Threshold filter:

Access: P&R Admin/Admin

1. 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**.

Keysight, EX-1004
Page 423

- 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: P&R Admin/Admin

1. 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: P&R Admin/Admin

1. 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**.

Keysight, EX-1004
Page 424

- 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: P&R Admin/Admin

1. 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 425

Rule Content Filters (Continued)

To use this filter, click...	Then...	Result
Reference	Type the reference string to filter by, and click **OK**.	Finds rules that contain the supplied string in the reference field.
Action	Select the action to filter by: • To find alert rules, select **Alert**, 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: • 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, 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, 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.

Keysight, EX-1004
Page 426

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 `metadata:"service http"` to locate rules with metadata relating to the HTTP service.	Find rules with metadata containing the matching key-value pair.

Understanding Rule Categories

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 the rules in a category. For example, if you do not have any Microsoft IIS web servers on your network, you might filter by the **web-iis** category and then disable all the rules showing to disable the entire **Web IIS** category.

IMPORTANT! The Sourcefire VRT may use the Security Enhancement Update (SEU) mechanism to add new rule categories.

Keysight, EX-1004
Page 427

The Rule Categories table lists each rule category and describes the types of events the rules in that category detect.

Rule Categories

Rule Category	Contents
Attack Responses	Rules that detect common attack responses.
Backdoor	Rules that detect backdoor-type exploits and activity.
Bad Traffic	Rules that detect anomalous traffic.
Chat	Rules that detect the use of unauthorized instant messaging programs.
Content Replace	Rules that use the `replace` keyword to neutralize attacks but not drop them in inline deployments, and to stop instant messaging and peer-to peer file transfers without adversely affecting the other allowed functionality of these applications.
Distributed Denial of Service (DDoS)	Rules that detect Distributed Denial of Service (DDoS) attacks.
Decoder	Rules that detect anomalous network traffic during packet decoding.
Deleted	Deleted rules that have not yet been removed from the product. These rules exist for historical purposes.
DNS	Rules that detect DNS-based exploits.
Denial of Service (DoS)	Rules that detect Denial of Service (DoS) attacks.
Exploit	Rules that detect known exploits.
Finger	Rules that detect finger-based exploits.
FTP	Rules that detect FTP-based exploits.
ICMP	Rules that detect ICMP-based exploits.
ICMP Info	Rules that detect exploits that use ICMP information requests.
IMAP	Rules that detect IMAP-based exploits.
Info	Rules that detect possible reconnaissance activity.

Keysight, EX-1004
Page 428

Rule Categories (Continued)

Rule Category	Contents
Local	User-defined rules that you have created. The Local category appears only if user-defined rules exist. See Writing New Rules on page 861 for more information about creating custom rules.
Misc	Miscellaneous rules that do not fall under other categories.
Multimedia	Rules that detect the use of authorized multimedia applications.
MySQL	Rules that detect exploits that target MySQL database servers.
Netbios	Rules that detect NETBIOS/SMB-based exploits.
NNTP	Rules that detect NNTP-based exploits.
Oracle	Rules that detect exploits that target Oracle database servers and ports.
Other IDS	Rules that detect connections from other intrusion systems.
(P2P) Peer to Peer	Rules that detect the use of unauthorized peer-to-peer downloading applications.
Policy	Rules that detect miscellaneous suspicious activity.
POP2	Rules that detect POP2-based exploits.
POP3	Rules that detect POP3-based exploits.
Preprocessor	Rules that detect anomalous network traffic or exploits during preprocessing.
RPC	Rules that detect RPC-based exploits.
Rservices	Rules that detect UNIX or Linux-based remote services (rlogin, rsh) login exploits.
SCADA	Rules that detect unusual SCADA traffic on a network.
Scan	Rules that generate events when scanning activity is detected.
Sensitive Data	Rules that detect sensitive data types.
Shellcode	Rules that detect shell code.
SMTP	Rules that detect exploits that target mail servers.

Keysight, EX-1004
Page 429

Rule Categories (Continued)

Rule Category	Contents
SNMP	Rules that detect SNMP-based exploits.
Specific Threats	Rules that generate events on high-profile network threats that have widely recognized names such "sasser" and "netsky".
Spyware- PUT	Rules that detect attempts to infect hosts with spyware, key loggers, trojans, and other potential unwanted technology (PUT). These rules also detect hosts that are already infected. **IMPORTANT!** You should read the rule documentation accompanying these rules because many are resource-intensive and can produce performance problems.
SQL	Rules that detect attacks on Microsoft SQL Servers.
Telnet	Rules that detect exploits that target or originate from telnet servers.
TFTP	Rules that detect TFTP-based exploits.
Virus	Rules that detect known worms and viruses.
VoIP	Rules targeted at detection of threats in Voice Over IP traffic.
Web ActiveX	Rules that detect attacks that exploit ActiveX controls.
Web CGI	Rules that detect attacks that exploit CGI scripts.
Web Client	Rules that detect suspicious traffic sent to or from web servers.
Web Coldfusion	Rules that detect attacks that exploit Cold Fusion scripts.
Web Frontpage	Rules that detect attacks that exploit FrontPage scripts.
Web IIS	Rules that detect attacks against IIS web servers.
Web Misc	Rules that detect miscellaneous exploits targeting web servers.
Web PHP	Rules that detect attacks that exploit PHP scripts.
X11	Rules that detect X11-based exploits.

Keysight, EX-1004
Page 430

Editing a Rule Filter Directly

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 422, Understanding Rule Content Filters on page 425, and Understanding Rule Categories on page 427. 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 405 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 772. For information on rule SIDs and GIDs, see Reading Preprocessor

Keysight, EX-1004
Page 431

Generator IDs on page 488. 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 419.

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

Keysight, EX-1004
Page 432

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

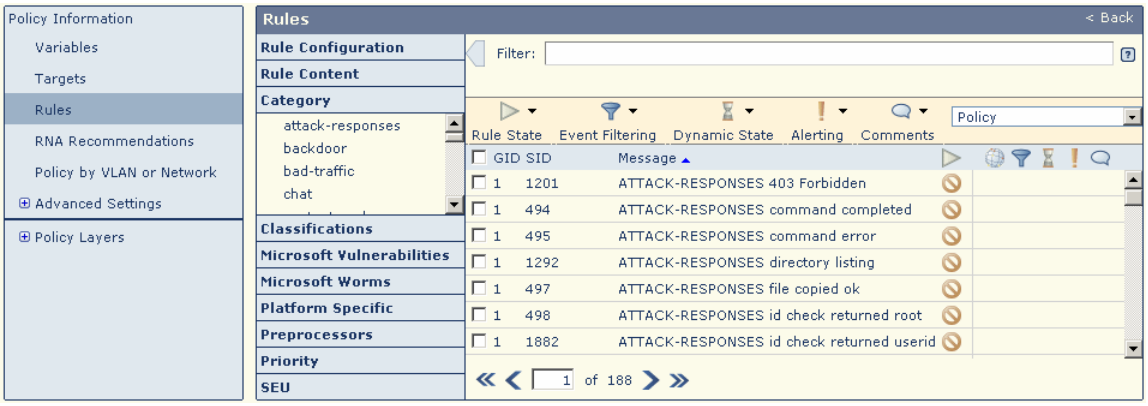
2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 419.
   - Understanding Rule Configuration Filters on page 422
   - Understanding Rule Content Filters on page 425

Keysight, EX-1004
Page 433

- Understanding Rule Categories on page 427
- Editing a Rule Filter Directly on page 431 for more information.

The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

The page refreshes to display all matching rules.

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 435 for information on enabling and disabling rules on the Rules page.
   - See Filtering Intrusion Event Notification Per Policy on page 439 for information on adding thresholding and suppression to rules.
   - See Adding Dynamic Rule States on page 452 for information on setting dynamic rule states that trigger when rate anomalies occur in matching traffic.
   - See Adding Alerts on page 458 for information on adding SNMP or OPSEC alerts to specific rules.
   - See Adding Rule Comments on page 466 for more information on adding rule comments to rules.
7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 434

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Setting Rule States

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Sourcefire Vulnerability Research Team (VRT) sets the default state of each intrusion, packet decoder, 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. Rules in an intrusion policy 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, 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. If you assign inline or inline with fail open interface sets to your IPS detection engines, 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 an inline interface 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 IPS 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, IPS drops the malicious packet and generates an intrusion event. The packet never reaches its target.

Keysight, EX-1004
Page 435

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 IPS 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 IPS to detect a specific intrusion attempt, then drop the packet containing the attack and generate an intrusion event when it finds matching traffic in an inline deployment, or to generate an intrusion event when it finds matching traffic in a passive deployment.

  Note that your intrusion policy must be set to drop rules in an inline deployment for IPS to drop packets; see Setting Drop Behavior in an Inline Deployment on page 358 for more information.
- Set the rule state to **Disable** if you do not want IPS to evaluate matching traffic.

To use drop rules, you must:

1. Assign an inline or inline with fail open interface set to an inline IPS detection engine.
2. Enable the **Drop when Inline** option in the policy you apply.
3. Set the rule state to **Drop and Generate Events** for any rules that should drop all packets that match the rule.
4. Apply the policy to your inline detection engine.

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 417.

See Understanding and Writing Intrusion Rules on page 761 for information about rule anatomy, rule keywords and their options, and rule writing syntax.

The VRT sometimes uses a Security Enhancement Update (SEU) to change the default state of one or more rules in a default policy. If you allow SEUs to update your base policy, you also allow the SEU 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 SEU will not override your change.

To change the rule state for one or more rules:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 436

2. Click **Edit** next to the policy you want to edit.

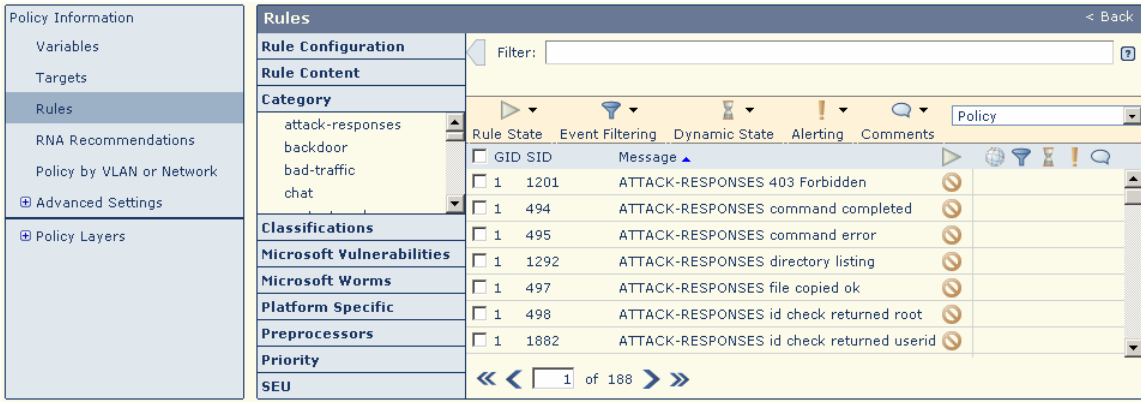
   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information 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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

     The page refreshes to display all matching rules.

Keysight, EX-1004
Page 437

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**.

     Note that in passive deployments, including when an inline interface is in tap mode, setting the rule state to **Drop and Generate Events** causes rules to generate events without dropping packets.
   - 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 sensor 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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 438

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Filtering Intrusion Event Notification Per Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 439 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 447 explains how to suppress notification of specified events per source or destination IP address per policy.

Configuring Event Thresholding

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 439

For more information see the following sections:

- Understanding Event Thresholding on page 440
- Adding and Modifying Intrusion Event Thresholds on page 442
- Viewing and Deleting Intrusion Event Thresholds on page 445
- Setting a Threshold for a Rule on page 412

Understanding Event Thresholding

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, IPS stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the count argument) trigger the rule during the specified time period. Note that the counter for the time restarts once you hit the threshold count of events and IPS 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. IPS 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, IPS logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to two, and **Seconds** to 10, the following event counts result: • If the rule is triggered once in 10 seconds, IPS does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, IPS generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, IPS generates one event (the threshold is met when the rule triggered the second time and following events are ignored)

Keysight, EX-1004
Page 440

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 IPS tracks event instances.

Thresholding IP Options

Option	Description
Source IP	Calculates event instance count per source IP address.
Destination IP	Calculates the event instance count per destination IP address.

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, IPS 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 452, Filtering Events on page 849, and Configuring Suppression Per Intrusion Policy on page 447 for more information.

Keysight, EX-1004
Page 441

See the following sections for more information:

- Adding and Modifying Intrusion Event Thresholds on page 442
- Setting a Threshold for a Rule on page 412
- Viewing and Deleting Intrusion Event Thresholds on page 445

TIP! You can also add thresholds from within the packet view of an intrusion event. See Viewing Event Information on page 293 for more information.

Adding and Modifying Intrusion Event Thresholds

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can set a threshold for one or more specific rules. You can also separately or simultaneously modify existing threshold settings. You can set a a single threshold for each. Adding a threshold overwrites any existing threshold for the rule.

For more information on viewing and deleting threshold configurations, see Viewing and Deleting Intrusion Event Thresholds on page 445.

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 705.

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! If you apply a policy containing a global or individual threshold to a detection engine running on a sensor with multiple CPUs, you may receive a higher number of events than expected.

To add or modify event thresholds:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 442

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon () appears when you add the first keyword or type the first character and disappears if you delete all characters.

     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.

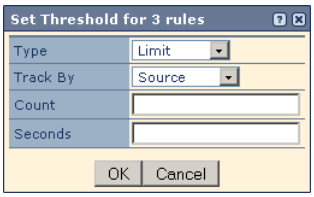
     Note that you can also modify the global threshold that applies by default to all rules and preprocessor-generated events. See Using Global Rule Thresholding on page 705 for more information.

   Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

Keysight, EX-1004
Page 443

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**.

    IPS 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, IPS includes an indication over the icon of the number of event filters.
12. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 444

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Viewing and Deleting Intrusion Event Thresholds

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

To view or delete a threshold:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

Keysight, EX-1004
Page 445

- 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. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

  The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

  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.

     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 705 for more information.

   Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.
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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 446

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Configuring Suppression Per Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 452, Filtering Events on page 849, and Configuring Event Thresholding on page 439 for more information.

Keysight, EX-1004
Page 447

See the following sections for more information:

- Suppressing Intrusion Events on page 448
- Viewing and Deleting Suppression Conditions on page 450

TIP! You can also add suppressions from within the packet view of an intrusion event. See Viewing Event Information on page 293 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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:

Keysight, EX-1004
Page 448

- 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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

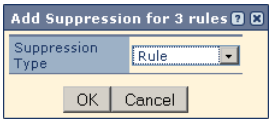
  The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

  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.

   Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.
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, CIDR block, or variable you want to specify as the source or destination IP address.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

Keysight, EX-1004
Page 449

9. Click **OK**.

   IPS 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, IPS includes an indication over the icon of the number of event filters.

10. You have the following options:

    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Viewing and Deleting Suppression Conditions

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 450

To view or delete a defined suppression condition:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon () appears when you add the first keyword or type the first character and disappears if you delete all characters.

     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.

   Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

Keysight, EX-1004
Page 451

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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Adding Dynamic Rule States

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 service requests. You can use rate-based prevention to change the action of a rule in response to excessive rule matches for specific rules.

Keysight, EX-1004
Page 452

For more information, see the following sections:

- Understanding Dynamic Rule States on page 453
- Setting a Dynamic Rule State on page 454

Understanding Dynamic Rule States

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 3D Sensors 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-base attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate service 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 traffic detected by the detection engine.

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 once 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 IPS does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action can 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.

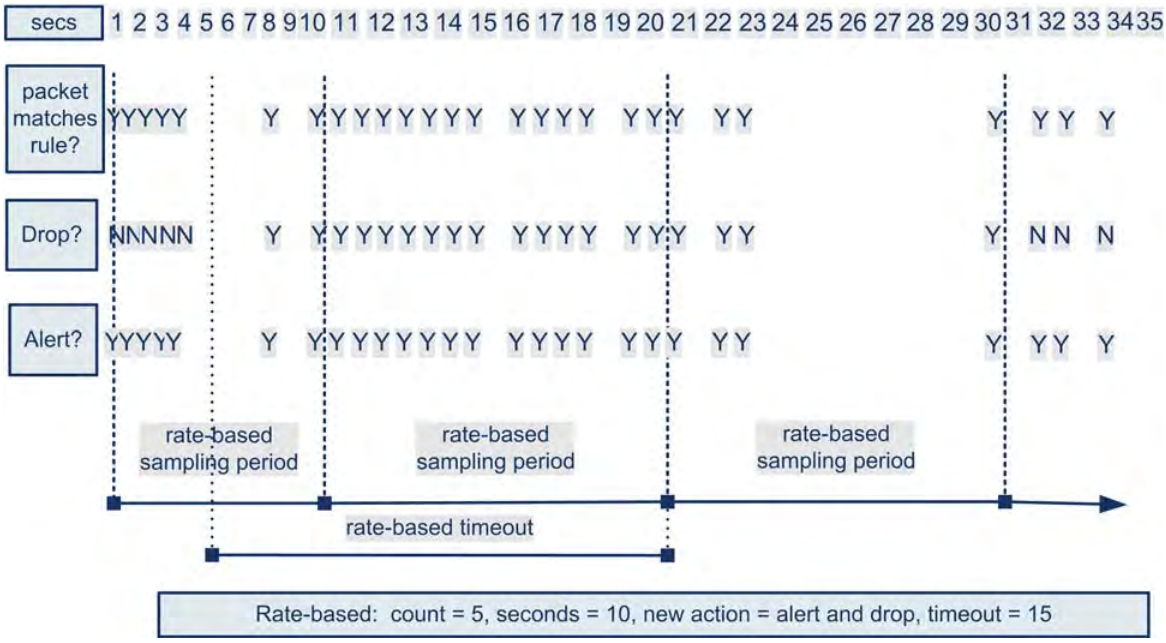
IMPORTANT! Rate-based actions cannot enable disabled rules or drop traffic that matches disabled rules.

Keysight, EX-1004
Page 453

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 454

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

Keysight, EX-1004
Page 455

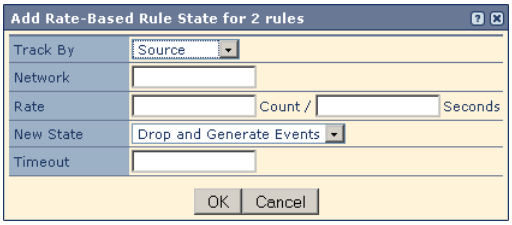
4. Locate the rule or rules where you want to add a dynamic rule state. You have the following options:
   - To sort the current display, click on a column heading or icon. To reverse the sort, click again.
   - Construct a filter by clicking on keywords or arguments in the filter panel on the left. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. Multiple keywords are combined using AND logic. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

     The page refreshes to display all matching rules.

     Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.
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 for the detection engine.

Keysight, EX-1004
Page 456

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, CIDR block, variable, or a comma-separated list of these. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

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**.

    IPS 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, IPS includes an indication over the icon of 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. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 457

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Adding Alerts

Platform: DC/MDC/Standalone Sensor

Component: IPS

If you configure SNMP or OPSEC alerting for your Sourcefire 3D System, you can add an alert to specific rules in your intrusion policy. For more information, see the following sections:

- Adding SNMP Alerts on page 458
- Adding OPSEC Alerts on page 460

Adding SNMP Alerts

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 458

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon () appears when you add the first keyword or type the first character and disappears if you delete all characters.

     The page refreshes to display all matching rules.

     Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

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**.

   IPS 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, IPS includes an indication over the icon of 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**.

Keysight, EX-1004
Page 459

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Adding OPSEC Alerts

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can add individual SAM responses to intrusion rules and configure how your Check Point™ OPSEC™ Suspicious Activity Monitor (SAM) responds when specific standard text rules, shared object rules, or preprocessor rules within an intrusion policy trigger.

You must configure the SAM Response independently for each rule. This configuration is saved separately from the rule within a policy, so that any changes to the rule do not change the configured response. If you are using your Defense Center, the firewall responses are pushed to your managed 3D Sensors when you apply the affected intrusion policy. For information on creating the OPSEC SAM Application and connecting the 3D Sensor and the SAM server, see Using Check Point OPSEC Responses on page 735.

Keysight, EX-1004
Page 460

You can disable individual OPSEC SAM responses. For information on clearing all SAM responses for a detection engine, see Clearing All OPSEC SAM Responses for a Detection Engine on page 745

IMPORTANT! OPSEC settings in intrusion policies applied to 3Dx800 sensors have no effect.

See Understanding OPSEC Filter Options on page 464 for more information.

If you configure an OPSEC firewall response 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.

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 an OPSEC alert:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 OPSEC alerts. You have the following options:

   - To sort the current display, click on a column heading or icon. To reverse the sort, click again.

Keysight, EX-1004
Page 461

- 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. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

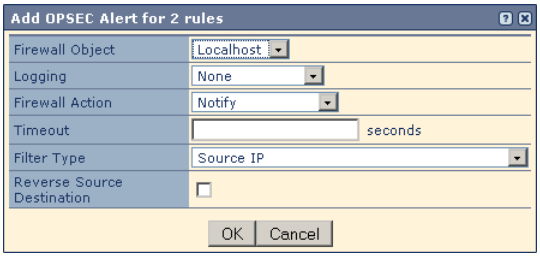
  The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

  The page refreshes to display all matching rules.

  Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

5. Select the rule or rules where you want to set OPSEC 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 OPSEC Alert**.

   The Add OPSEC Alert page appears.



7. From the **Firewalls Object** list, select the firewall object that receives this response.

   You can use the defaults of All, Gateways, or Localhost, or can create additional firewall objects using the SAM Control Panel. For details, refer to Adding and Deleting Firewall Objects on page 745.

Keysight, EX-1004
Page 462

8. From the **Logging** list, select one of the following log types.

Log Type	Description
None	Does not log SAM responses to this rule.
Long	Performs detailed logging, but does not store the event that caused the SAM response.
Long with Alert	Keeps more detailed logs along with the event that caused the SAM response.

9. From the **Firewall Action** list, select the action you want the firewall to take. You can select one of the following actions.

Action	Description
Notify	All connection attempts are logged as specified by the logging attribute. Traffic is not impeded.
Inhibit	All connection attempts are rejected. Existing connections will continue to function.
Inhibit and Close	All connection attempts are rejected. Existing connections are closed.
Drop	All connection attempts are dropped. Existing connections will continue to function.
Drop and Close	All connection attempts are dropped. Existing connections are closed.

10. In the **Timeout** field, specify the number of seconds that the specified action continues after it is initiated by the firewall.

IMPORTANT! The maximum value is 2,147,483 seconds. Specifying a value of 0 causes the action to continue until it is manually disabled by a firewall administrator.

11. From the **Filter Type** list, select the filter type you want to apply from the drop-down list.

Sourcefire supports 15 different filter types. For details, refer to Understanding OPSEC Filter Options on page 464.

Keysight, EX-1004
Page 463

12. Depending on the filter type you selected, the **Network Mask** field may appear. If it appears, enter the appropriate network mask value (for example, 255.255.255.0).
13. To install the filter with reversed source and destination information from the packet that triggered the rule, check **Reverse Source Destination**.
14. Click **OK**.

    IPS 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, IPS includes an indication over the icon of the number of alert types.

TIP! To remove an OPSEC alert from a rule, click the check box next to the rule and select **Alerting > Remove OPSEC Alerts**.

- You can continue editing your policy.

  See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding OPSEC Filter Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

When configuring firewall responses, you must specify a filter. Filters define parameters of traffic against which the specified action is taken. The supported filters apply information from packets that trigger intrusion rules to traffic at the

Keysight, EX-1004
Page 464

firewall, and take the specified action against those packets. The following list describes the available filters.

- **Source IP** performs the specified action against network traffic that has the same source IP address as the packet that triggered the intrusion rule.
- **Destination IP** performs the specified action against network traffic that has the same destination IP address as the packet that triggered the intrusion rule.
- **Source or Destination IP** performs the specified action against network traffic that has either the same source or destination IP address as the source IP address of the packet that triggered the intrusion rule.
- **Source IP, Destination IP, and Service** performs the specified action against network traffic that has the same source IP address, destination IP address, IP protocol, and destination port as the packet that triggered the intrusion rule.
- **Destination IP and Service** performs the specified action against network traffic that has the same destination IP address, IP protocol, and destination port as the packet that triggered the intrusion rule.
- **Source IP and Protocol** performs the specified action against network traffic that has the same source IP address and IP protocol as the packet that triggered the intrusion rule.
- **Destination IP and Protocol** performs the specified action against network traffic that has the same destination IP address and IP protocol as the packet that triggered the intrusion rule.
- **Source Network** performs the specified action against network traffic identified as coming from the same source network as the packet that triggered the intrusion rule, calculated using the Source Network Mask attribute specified in the Check Point OPSEC SAM response configuration.
- **Destination Network** performs the specified action against traffic whose source address matches the destination network, calculated using the Network Mask attribute, of the packet that triggered the intrusion rule.
- **Source or Destination Network** performs the specified action against any packet whose source or destination network, calculated using the Network Mask attribute, matches the address of the packet that triggered the intrusion rule.
- **Source Network and Protocol** performs the specified action against any network traffic whose source network, calculated using the Source Network Mask attribute, and IP protocol match the source network and IP protocol of the packet that triggered the intrusion rule.
- **Destination Network and Protocol** performs the specified action against any network traffic whose destination network, calculated using the Destination Network Mask attribute, and IP protocol match the destination network and IP protocol of the packet that triggered the intrusion rule.

Keysight, EX-1004
Page 465

- **Destination Network and Service** performs the specified action against any network traffic with the same destination network, calculated by comparing the Destination Network Mask attribute, IP protocol, and destination port to those of the packet that triggered the intrusion rule.
- **Source Network, Destination IP and Service** performs the specified action against any network traffic with the same source network, calculated by comparing the Source Network Mask, destination IP address, IP protocol, and destination port to those of the packet that triggered the intrusion rule.
- **Source IP, Destination Network and Service** performs the specified action against any network traffic with the same source IP address, destination network, calculated by comparing the Destination Network Mask, IP protocol, and destination port to those of the packet that triggered the intrusion rule.

Adding Rule Comments

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 Edit Rule page. For more information on editing rules, see Modifying Existing Rules on page 864.

To add a comment to a rule:

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

Keysight, EX-1004
Page 466

- 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. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

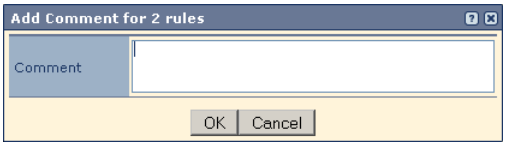
  The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

  The page refreshes to display all matching rules.

  Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

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**.

   IPS adds the comment and displays a comment icon ( ) next to the rule in the Comments column. If you add multiple comments to a rule, IPS includes an indication over the icon of 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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 467

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Managing RNA Rule State Recommendations

Platform: DC/MDC/Standalone Sensor

Component: IPS + RNA

On a Sourcefire 3D System with both the IPS component and the RNA component, you can use the RNA Recommended Rules feature to associate hosts and services on your network with rules written to protect those hosts and services. See Introduction to Sourcefire RNA on page 878 for information on deploying the RNA component.

When you configure the RNA Recommended Rules feature, the system identifies hosts and services on your network, searches your base policy for rules that protect against vulnerabilities associated with the identified hosts and services, and identifies the current state of rules in your base policy. The system then

Keysight, EX-1004
Page 468

recommends rule states and, optionally, sets the rules to the recommended states using the following criteria:

- The system recommends that you set the rule state to Generate Events when the rule is set to either Generate Events or Disable in the base policy and the rule protects against vulnerabilities associated with hosts and services on your network.
- The system recommends that you set the rule state to Drop and Generate Events when the rule is set to Drop and Generate Events in the base policy and the rule protects against vulnerabilities associated with hosts and services on your network.
- The system recommends that you set the rule state to Disable when the rule protects against vulnerabilities not associated with hosts and services on your network.

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 RNA-recommended rule states is that the rules in your intrusion policy match the settings recommended by Sourcefire for the specific hosts and services on your network. See Using Default Intrusion Policies on page 362 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 subsequently. Advanced recommendations options allow you to tailor your configuration.

Note that RNA does not recommend rule state changes for preprocessor rules, including sensitive data preprocessor 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 Recommended Rule State Generation on page 1725.

See the following sections for more information:

- Understanding Basic Rule State Recommendations
- Understanding Advanced Rule State Recommendations
- Using RNA Recommendations

Understanding Basic Rule State Recommendations

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.

Keysight, EX-1004
Page 469

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 435 for information on manually changing the state of selected rules. See Managing Rules in an Intrusion Policy on page 403 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 generating recommendations prevents the system from modifying the states of those rules, and manually setting the states of specified rules after generating recommendations overrides recommended state modifications of those rules. Optionally, you can choose to include in the intrusion policy report a list of rules whose rule state differs from the recommended state. See Viewing an Intrusion Policy Report on page 350 for more information.

Note that choosing to use recommended rule states adds a read-only RNA 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 494 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 RNA recommended rules, the system recommends rule state changes for hosts identified by the $HOME_NET variable for the detection engines where you apply your intrusion policy. 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 470 for more information.

Understanding Advanced Rule State Recommendations

Advanced settings allow you to redefine which hosts on your network RNA monitors for vulnerabilities, to influence which rules RNA recommends based on rule overhead, and to specify whether to generate recommendations to disable rules.

If you want IPS to dynamically adapt active rule processing for specific packets based on RNA host information, you can also enable adaptive profiles. For more information, see Adaptive Profiles and RNA Recommended Rules on page 702.

See the following sections for more information:

- Understanding the Network to Examine on page 471
- Understanding Rule Overhead on page 471

Keysight, EX-1004
Page 470

Understanding the Network to Examine

You configure the RNA Recommended Rules feature by identifying a network in the network map to examine for hosts and services. The system then recommends the rules you can activate to protect a network comprised of those hosts and services. For information on the network map, see Using the Network Map on page 923.

You can specify the network to examine for RNA recommendations using the following methods, separately or combined:

- by using the value defined for the $HOME_NET variable for detection engines targeted in your intrusion policy

  Note that the value of any is used if you do not explicitly set a policy-specific value and your policy does not target detection engines with specified values for $HOME_NET. See Managing Targets on page 368 and Understanding Existing Variables on page 372 for more information.
- by specifying hosts to examine for RNA 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 use detection engine $HOME_NET values or specify the hosts to examine. When you use both a policy-specific $HOME_NET value and detection engine $HOME_NET values, the values for both are combined.

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

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 410 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. 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 IPS factors rule overhead into recommendations to generate events or to drop and generate events. IPS does not factor rule overhead into

Keysight, EX-1004
Page 471

recommendations to disable rules. Note also that local rules do not have any overhead, unless they are mapped to a third-party RNA vulnerability. See Importing Local Rule Files on page 1857 and Managing Third-Party Product Mappings on page 1408 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 services on your network have not changed, both sets of recommendations with overhead set to medium will be the same for that rule set.

Using RNA Recommendations

Platform: DC/MDC/Standalone Sensor

Component: IPS + RNA

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 469 and Understanding Advanced Rule State Recommendations on page 470 for more information. After generating recommendations, you can elect to use the recommended rule states; you can also view recommended states on the Rules page and use any of the features available on the Rules page.

To use RNA rule state recommendations:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 472

3. You have two options:
   - If you have not already generated recommendations, click **No recommendations have been generated. Click here to set up RNA recommendations.** on the Policy Information page.
   - If you have already generated recommendations, select **Click to change recommendations**.

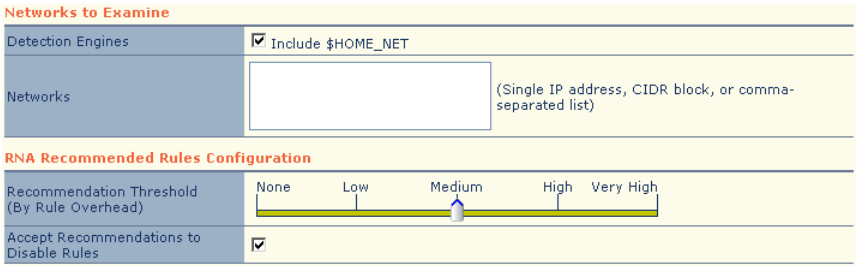
   The RNA Recommended Rules Configuration page appears.



4. 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 RNA recommended state, select **Include all differences between recommendations and rule states in policy reports**. See Viewing an Intrusion Policy Report on page 350 for more information.
   - To generate RNA recommended rule states using the default settings, go to step 9.
   - To modify the advanced recommendations options, go to step 5.

5. Click the **+** sign to expand the Advanced Settings section.

   The advanced RNA recommendations options appear.



Keysight, EX-1004
Page 473

6. You have three choices:
   - To specify the network to examine for RNA recommendations using the value of the `$HOME_NET` variable for detection engines where you will apply the policy, ensure that the **Detection Engine** check box is selected. See Managing Targets on page 368 for more information.

     Note that this option is selected by default. When not selected, a default value of any is used; when used in a policy with a policy-specific $HOME_NET value, detection engine and policy-specific values are combined. See Modifying Variables on page 374 for information on configuring the value of $HOME_NET.
   - To manually specify the network to examine for RNA recommendations, clear the **Detection Engine** check box and identify the hosts to examine for RNA recommendations in the **Networks** field.

     You can enter a single IP address or CIDR block, or a comma-separated list comprised of either or both, including negation. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
   - To specify the network to examine using both the `$HOME_NET` variable and additional hosts you specify, ensure that the **Detection Engine** check box is selected and, in the **Networks** field, identify additional hosts to examine for RNA 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 471 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 471 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.

Keysight, EX-1004
Page 474

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 RNA 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 982.

   Note that updating the policy to use or not use recommendations can take several minutes depending on the size of your network and rule set.

   IMPORTANT! RNA always recommends that you enable a local rule associated with a third party vulnerability mapped to a host. RNA does not make state recommendations for unmapped local rules. For more information, see Managing Third-Party Product Mappings on page 1408.

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. See Understanding Basic Rule State Recommendations on page 469 for more information.

Keysight, EX-1004
Page 475

11. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 476

Chapter 11
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 361 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 IPS 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 478 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 484 explains how preprocessors normalize traffic for use by the rules engine.

Keysight, EX-1004
Page 477

- Automatically Enabling Advanced Settings on page 490 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 492 explains troubleshooting options that you should set only when asked to do so by Sourcefire Support.
- Using Layers in an Intrusion Policy on page 494 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

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.

You must enable an advanced setting 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

Keysight, EX-1004
Page 478

following listing provides 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 513 for more information.

For information on...	See...
DCE/RPC Configuration	Configuring the DCE/RPC Preprocessor on page 526
DNS Configuration	Configuring the DNS Preprocessor on page 535
FTP and Telnet Configuration	Decoding FTP and Telnet Traffic on page 538
HTTP Configuration	Decoding HTTP Traffic on page 560
Sun RPC Configuration	Configuring the Sun RPC Preprocessor on page 579
SMTP Configuration	Configuring SMTP Decoding on page 585
SSH Configuration	Selecting SSH Preprocessor Options on page 591
SSL Configuration	Configuring the SSL Preprocessor on page 599

Keysight, EX-1004
Page 479

- 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.

For information on...	See...
Checksum Verification	Verifying Checksums on page 602
Inline Normalization	Normalizing Inline Traffic on page 604
IP Defragmentation	Configuring IP Defragmentation on page 618
Packet Decoding	Configuring Packet Decoding on page 625
TCP Stream Configuration	Configuring TCP Stream Preprocessing on page 640
UDP Stream Configuration	Configuring UDP Stream Preprocessing on page 645

Keysight, EX-1004
Page 480

- 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.

For information on...	See...
Back Orifice Detection	Detecting Back Orifice on page 648
Portscan Detection	Configuring Portscan Detection on page 655
Rate-Based Attack Prevention	Configuring Rate-Based Attack Prevention on page 673
Sensitive Data Detection	Configuring Sensitive Data Detection on page 683

- Detection Enhancement

  With adaptive profiles, IPS can adapt to network traffic by associating traffic with host and service information from the RNA network map and then processing the traffic accordingly.

For information on...	See...
Adaptive Profiles	Configuring Adaptive Profiles on page 702

- Intrusion Rule Thresholds

  Global rule thresholding can prevent your system from being overwhelmed with a large number of events by allowing you to use thresholds to limit the number of times the system logs and displays intrusion events.

For information on...	See...
Global Rule Thresholding	Configuring Global Thresholds on page 708

Keysight, EX-1004
Page 481

- Performance Settings

  The IPS component provides server settings for improving system performance.

For information on...	See...
Event Queue Configuration	Event Queue Configuration on page 714
Latency-Based Packet Handling	Configuring Packet Latency Thresholding on page 719
Latency-Based Rule Handling	Configuring Rule Latency Thresholding on page 724
Performance Statistics Configuration	Performance Statistics Configuration on page 726
Regular Expression Limits	Constraining Regular Expressions on page 728
Rule Processing Configuration	Rule Processing Configuration on page 731

- 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. Additionally, you can specify Check Point™ OPSEC™ firewall responses for intrusion events generated when a standard text rule triggers. You can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events.

For information on...	See...
OPSEC Configuration	Using Check Point OPSEC Responses on page 735
SNMP Alerting	Configuring SNMP Responses on page 749
Syslog Alerting	Configuring Syslog Responses on page 754

Keysight, EX-1004
Page 482

To modify advanced settings:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 483

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding Preprocessors

Platform: DC/MDC/Standalone Sensor

Component: IPS

3D Sensors with IPS use preprocessors to 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 IPS component 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 command syntax, and validating checksums.

You can configure these preprocessors to ensure that the packets IPS 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 IPS, detects the state of a TCP

Keysight, EX-1004
Page 484

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 3D Sensor 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 435.

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 341, Managing Rules in an Intrusion Policy on page 403, and Setting Rule States on page 435 for more information.

In addition to preprocessors, the IPS component also provides advanced settings for detecting anomalous traffic, enhancing detection, applying a global rule threshold, tuning performance, and configuring external SNMP, syslog, and OPSEC alerting.

See the following sections for more information:

- Meeting Traffic Challenges with Preprocessors on page 485 describes both normal traffic and the inspection challenges experienced at the network layer, transport layer, and application layer.
- Understanding Preprocessor Execution Order on page 486 explains the order of execution in Sourcefire 3D System preprocessors.
- Reading Preprocessor Events on page 488 describes preprocessor events and the information they contain.

Meeting Traffic Challenges with Preprocessors

Platform: DC/MDC/Standalone Sensor

Component: IPS

The IPS component 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

Keysight, EX-1004
Page 485

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 IPS 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 against both IPS and the analyst. Attackers can also launch subterfuge attacks by transmitting TCP, UDP and ICMP packets with invalid checksums in an attempt to cause IPS 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 IPS 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

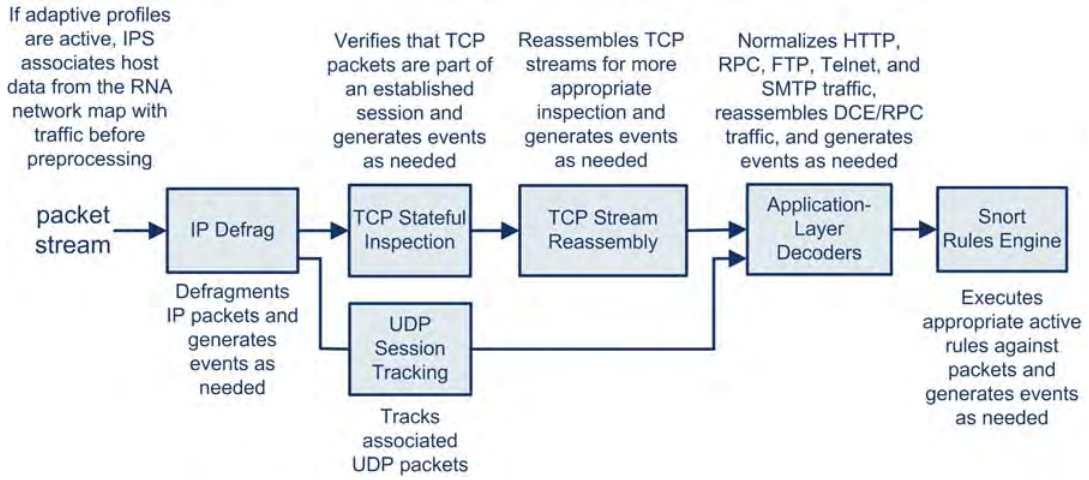
Platform: DC/MDC/Standalone Sensor

Component: IPS

Protocol decoders, preprocessors, and rules run in a specific order so that they can perform IP transfer layer protocol decoding first, then perform data

Keysight, EX-1004
Page 486

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:

- IPS can generate an intrusion event against fragmented IP datagrams that cannot be defragmented, and then stop inspecting those packets.
- IPS can generate an event against TCP packets whose state cannot be validated, and then stop inspecting those packets.
- IPS 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.
- IPS can adapt IP defragmentation and stream preprocessing behavior to fit the operating system formats on the target host using adaptive profiles, which require both IPS and RNA, or target-based policies, which require only IPS, 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.

Keysight, EX-1004
Page 487

Reading Preprocessor Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 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 488 describes the information contained in a preprocessor-generated event.
- Reading Preprocessor Generator IDs on page 488 details the information provided by the preprocessor generator ID.

Understanding the Preprocessor Event Packet Display

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 488

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 403 and the Rule Types table on page 405 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 405 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 `tag` rule option is used. For more information, see Evaluating Post-Attack Traffic on page 850.
3	Shared Object Rule	The event was generated when the packet triggered a shared object rule. See the Rule Types table on page 405 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.
106	RPC Decoder	The RPC decoder decoded the packet.
116	Packet Decoder	The event was generated by the packet decoder.
119	HTTP Inspection Preprocessor	The event was generated by the HTTP Inspection preprocessor.
120	HTTP Inspection Anomalous Server Preprocessor	The event was generated by the HTTP Inspection anomalous server preprocessor. This processor contains a single SID (1), which generates an event when HTTP traffic is received by ports not specified as web server ports in your intrusion policy.
122	Portscan Detector	The event was generated by the portscan flow detector.

Keysight, EX-1004
Page 489

Generator IDs (Continued)

ID	Component	Description
123	IP Defragmentor	The event was generated when a fragmented IP datagram could not be properly reassembled.
124	SMTP Decoder	The event was generated when the SMTP preprocessor detected an exploit against an SMTP verb. See Understanding SMTP Decoding on page 582 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 546 and Understanding Client-Level FTP Options on page 555 for more information.
126	Telnet Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within telnet traffic.
128	SSH Preprocessor	The event was generated when the SSH preprocessor detected an exploit within SSH traffic.
129	Stream Preprocessor	The event was generated during stream preprocessing by the stream preprocessor.
131	DNS Preprocessor	The event was generated by the DNS preprocessor.
133	DCE/RPC Preprocessor	The event was generated by the DCE/RPC preprocessor.
134	Suspended or Re-enabled Intrusion Rules	The event was generated when rule latency suspended or re-enabled a group of intrusion rules. Suspended rules have a SID of 1, and re-enabled rules have a SID of 2. For more information, see Understanding Rule Latency Thresholding on page 720.

Automatically Enabling Advanced Settings

Platform: DC/MDC/Standalone Sensor

Component: IPS

The IPS component can enable advanced settings when they are required by a standard text rule, a shared object rule, or a preprocessor rule. When you save an intrusion policy with a disabled advanced setting that is required by a rule or rule option, you are prompted whether you want IPS to automatically enable the required advanced setting. Before you can save the policy, you must either manually enable the required advanced setting configuration, allow IPS to automatically enable the required advanced setting, or disable any rules that require the advanced setting.

Keysight, EX-1004
Page 490

Note that IPS 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	Rules and Rule Options Causing Automatic Enable
Application Layer Preprocessors	DCE/RPC preprocessor	Keyword: • 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 Inspect preprocessor	Keyword: • 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	SSL preprocessor	Keyword: • ssl_state • ssl_version
Transport/Network Layer Preprocessors	TCP or UDP stream preprocessing	Keyword: • flow • flowbits • stream_size
Transport/Network Layer Preprocessors	TCP stream preprocessing	Keyword: stream_reassemble
Performance Settings	Regular Expression Limit	Keyword: pcre
Specific Threat Detection	Sensitive Data Detection	Generator ID: • 138 • 139

Keysight, EX-1004
Page 491

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 DNS preprocessor
- the FTP/Telnet preprocessor
- the HTTP Inspect preprocessor
- the SMTP preprocessor
- the SSL preprocessor
- the DCE/RPC preprocessor when the RPC over HTTP proxy, RPC over HTTP server, TCP, or SMB transport protocol is selected
- portscan detection when the TCP protocol is selected
- sensitive data detection

You are prompted whether to enable UDP stream preprocessing when it is disabled and you enable the DCE/RPC preprocessor with the UDP transport protocol selected.

Understanding Troubleshooting Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 492

The Troubleshooting Options table describes these troubleshooting options.

Troubleshooting Options

Advanced setting	Option	Description
TCP Stream Configuration (Global)	Session Termination Logging Threshold	This global TCP option logs a message when an individual connection exceeds the specified threshold. A value of 0 turns off the message. The upper limit of 1GB is also restricted by the amount of memory on the sensor allocated for stream processing. See Using TCP Stream Preprocessing on page 627 for more information.
TCP Stream Configuration (Policy)	Maximum Queued Bytes	This TCP target-based policy option specifies the amount of data that can be queued on one side of a TCP connection. A value of 0 specifies an unlimited number of bytes. See Using TCP Stream Preprocessing on page 627 for more information.
TCP Stream Configuration (Policy)	Maximum Queued Segments	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. A value of 0 specifies an unlimited number of data segment bytes. See Using TCP Stream Preprocessing on page 627 for more information.
Latency-Based Packet Threshold	Log packet when threshold is exceeded	Enables or disables logging of syslog messages that you can use in conjunction with Sourcefire Support to troubleshoot the system. See Understanding Packet Latency Thresholding on page 716 and Viewing the System Log on page 1981 for more information.
Latency-Based Rule Threshold	Log when threshold is exceeded	Enables or disables logging of syslog messages that you can use in conjunction with Sourcefire Support to troubleshoot the system. See Understanding Rule Latency Thresholding on page 720 and Viewing the System Log on page 1981 for more information.

Keysight, EX-1004
Page 493

Chapter 12
Using Layers in an Intrusion Policy

Larger organizations with many 3D Sensors with IPS may have many IPS 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 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 494 explains the layers that comprise a basic policy and how you can use them.
- Configuring User Layers on page 507 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

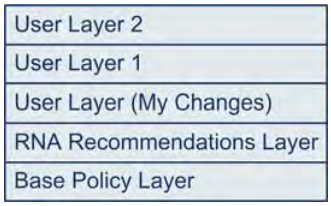
A policy where you do not add layers includes the read-only base policy layer and a single user-configurable layer that is initially named My Changes. If you generate and use rule state recommendations based on RNA data, the system automatically inserts a read-only RNA Recommendations layer immediately above

Keysight, EX-1004
Page 494

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 IPS 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 RNA 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.



When the highest layer in your policy is a read-only layer, or a shared layer as described in Sharing Layers on page 496, 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 403 for more information.
- enable, disable, or modify an advanced setting. See Modifying Advanced Settings on page 478 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 you apply a policy, IPS *flattens* the layers; that is, IPS 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, IPS applies the setting that is configured at the highest layer.

Note that regardless of whether you allow SEUs to modify your policy, changes in an SEU never override changes you make in a layer. This is because changes in an SEU 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

Keysight, EX-1004
Page 495

that an SEU makes to your default policy. See Importing SEUs and Rule Files on page 1848 for more information.

TIP! You can create an intrusion policy based solely on the default settings in the base policy and, optionally, RNA rule state recommendations by not adding layers, deleting or not modifying the My Changes layer, and, optionally, generating and using RNA recommendations.

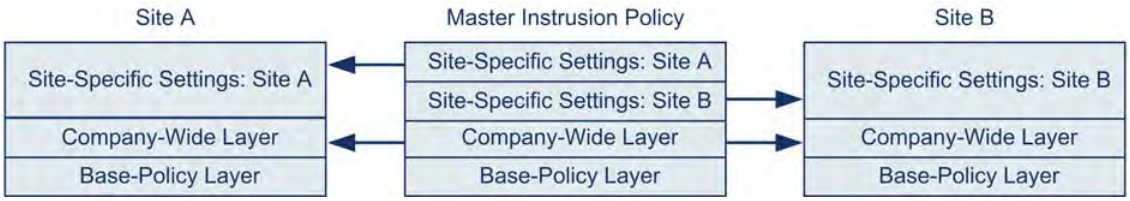
See the following sections for more information on using policy layers:

- Sharing Layers on page 496 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 497 explains how you can work with rules in an intrusion policy layer.
- Removing Multi-Layer Rule Settings on page 499 explains how you can remove settings for event filters, dynamic states, and alerting from multiple layers using the intrusion policy Rules page.
- Using the RNA Recommendations Layer on page 503 explains how you can view and delete rule attributes in layers.
- Using Layers with Advanced Settings on page 504 explains how you can work with advanced settings in an intrusion policy layer.

Sharing Layers

You can share any user-configurable layer with other intrusion policies. When you share a layer and then edit a configuration within that layer, IPS 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 you would not apply to detection engines but, instead, would use as the source for site-specific policies that you do apply to detection engines.



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

Keysight, EX-1004
Page 496

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 applying to a detection engine unless the settings for Site A that override the settings for Site B happened to be suitable for a different site. However, 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 508 for instructions on configuring shared layers.

Using Rules in Layers

Platform: DC/MDC/Standalone Sensor

Component: IPS

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:



Keysight, EX-1004
Page 497

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	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 435 for more information. 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. 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.
one	Threshold SNMP Alert OPSEC Alert	override a setting of the same type for the rule in a lower layer. Note that setting a threshold overwrites any existing threshold for the rule in the layer. See Configuring Event Thresholding on page 439 and Adding Alerts on page 458 for more information.
one or more	Suppression Rate-Based Rule State	cumulatively combine settings of the same type for each selected rule down to the first layer where a rule state is set for the rule. Settings below the layer where a rule state is set are ignored. See Configuring Suppression Per Intrusion Policy on page 447 and Adding Dynamic Rule States on page 452 for more information.
one or more	Comment	add a comment to a rule. Comments are rule-specific, not policy- or layer-specific. You can add one or more comments to a rule in any layer. See Adding a Rule Comment for a Rule on page 417 for more information.

For example, if you set a rule state to Drop and Generate Events in one layer and to Disabled in a higher layer, the intrusion policy Rules page shows that the rule is disabled.

In another example, if you set a source-based suppression for a rule to 192.168.1.1 in one layer, and you also set a destination-based suppression for the rule to 192.168.1.2 in another layer, the Rules page shows that the cumulative effect is to suppress events for the source address 192.168.1.1 and the

Keysight, EX-1004
Page 498

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** 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. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 498.

   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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 499

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.
- 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 435 for more information.

To remove settings in multiple layers using the Rules page:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** 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. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

Keysight, EX-1004
Page 500

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. Note that if you click an argument for a keyword already in the filter it replaces the existing argument. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy on page 419 and Setting a Rule Filter in an Intrusion Policy on page 432.

     The filter clearing icon ( ) appears when you add the first keyword or type the first character and disappears if you delete all characters.

     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.

   Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062). Preprocessor events have the generator ID for the preprocessor. For more information on preprocessor GIDs and SIDs, see Reading Preprocessor Generator IDs on page 488.

Keysight, EX-1004
Page 501

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.
   - To remove all OPSEC alert settings for a rule, select **Alerting > Remove OPSEC 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 435 for more information.

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 502

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Using the RNA Recommendations Layer

Platform: DC/MDC/Standalone Sensor

Component: IPS + RNA

When you have generated RNA 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 RNA Recommendations system layer immediately above the base layer in your intrusion policy. Subsequently choosing not to use the recommended rule states removes the RNA Recommendations system layer. Note that you can repeatedly remove and restore the RNA Recommendations layer by choosing to use or not use recommendations, but you cannot delete the layer manually.

Adding the RNA Recommendations layer adds an RNA Recommendations link under Policy Layers in the navigation panel. That link leads you to a read-only view of the RNA Recommendations layer page. From the RNA 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 403 for more information on working with rules on the Rules page.

Adding the RNA Recommendations layer also adds a Rules sublink beneath the RNA Recommendations link in the navigation panel. The Rules sublink provides

Keysight, EX-1004
Page 503

access to a read-only display of the Rules page in the RNA 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 RNA Recommendations column in this or other Rules page views, there is no recommendation for this rule.

  Note that when a rule in the RNA 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 471 for more information.

See Managing RNA Rule State Recommendations on page 468 for more information.

Using Layers with Advanced Settings

Platform: DC/MDC/Standalone Sensor

Component: IPS

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. So, 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 478 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 496 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

Keysight, EX-1004
Page 504

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.

IPS uses the configuration for an advanced setting in the highest layer where the configuration is enabled. Unless you explicitly modify the configuration, IPS 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, IPS 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 507 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 polices**. See Using Layers with Advanced Settings on page 504 and Configuring User Layers on page 507 for more information. Note that this action is not available on the Advanced Settings page.

Keysight, EX-1004
Page 505

Layer Summary Page Actions (Continued)

To...	You can...
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 478 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** 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. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 505.

Keysight, EX-1004
Page 506

5. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Configuring User Layers

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 507

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	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. 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. 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.
add a layer to your policy	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. 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.
enable or disable sharing a layer in your policy with other policies	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. 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.
move a layer above or below another layer	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. The screen refreshes and the layer appears in the new location.
manage rules or modify advanced setting configurations in a layer	click the edit icon ( ) for the layer. 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 403 and Modifying Advanced Settings on page 478 for more information. 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.

Keysight, EX-1004
Page 508

Policy Layer Configuration Actions (Continued)

To...	You can...
merge a layer into the next layer beneath it	click the merge icon ( ) for the layer you want to merge, then click **OK** when prompted or click **Cancel** to abandon the merge. The page refreshes and the layer is merged with the layer beneath it. A merged layer retains all settings that were unique to either layer, and accepts the settings from the higher layer if both layers included settings for the same rule or advanced setting. The merged layer retains the name of the lower layer. 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. 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.
copy a layer	click the copy icon ( ) of the layer you want to copy. 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.
delete a layer	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. The page refreshes and the layer is deleted. 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.
display the Policy Information page	click **Policy Summary**. See Managing Intrusion Policies on page 335 for an explanation of the actions you can take from the Policy Information page.

Keysight, EX-1004
Page 509

Policy Layer Configuration Actions (Continued)

To...	You can...
display the Layer summary page for a layer	click the layer name in the summary for the layer. The Layer summary page for the layer appears. 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 496, Using Layers with Advanced Settings on page 504, and Using Rules in Layers on page 497 for more information. 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.
display the Layer summary page for the base policy	click the base policy name in the base policy summary. The Layer summary page for the base policy appears. From this page you can select a different base policy for your intrusion policy and specify whether changes in an imported SEU 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 361, Allowing SEUs to Modify the Base Policy on page 363, Using Layers with Advanced Settings on page 504, and Using Rules in Layers on page 497 for more information.
display a layer-level advanced setting configuration page	click the advanced setting name in the summary for the layer. Note that configuration pages are read-only in the base policy and in shared layers. See Sharing Layers on page 496 and Understanding the Base Policy on page 361, and for more information.
display rules in a layer by rule state type	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. 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.

Keysight, EX-1004
Page 510

To configure layers in your intrusion policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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
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

Keysight, EX-1004
Page 511

4. Optionally, you can take any of the actions in the Policy Layer Configuration Actions on page 508.
5. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 512

Chapter 13
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.

When enabled, preprocessors normalize traffic according to the configured options, but generate no events in most cases before sending the normalized traffic to the rules engine unless you also enable the accompanying preprocessor rules on the Rule page. See Setting Rule States on page 435 for more information.

See the following sections for more information:

- Decoding DCE/RPC Traffic on page 514 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 531 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 538 describes the FTP/Telnet decoder and explains how to configure it to normalize and decode FTP and Telnet traffic.
- Decoding HTTP Traffic on page 560 describes the HTTP decoder and explains how to configure it to normalize HTTP traffic.
- Using the Sun RPC Preprocessor on page 578 describes the RPC decoder and explains how to configure it to normalize RPC traffic.

Keysight, EX-1004
Page 513

- Decoding SMTP Traffic on page 582 describes the SMTP decoder and explains how to configure it to normalize SMTP traffic.
- Detecting Exploits Using the SSH Preprocessor on page 590 explains how to identify and process exploits in SSH-encrypted traffic.
- Using the SSL Preprocessor on page 596 explains how you can use the SSL preprocessor to identify encrypted traffic and eliminate false positives by stopping IPS inspection of that traffic.

Decoding DCE/RPC Traffic

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 the IIS RPC over HTTP service, 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, and descriptions of SMB options and functionality refer to both SMB and Samba.

Most DCE/RPC exploits occur in DCE/RPC client requests targeted for DCE/RPC servers, which at some time or another could be practically any host on your network that is running Windows or Samba. The DCE/RPC preprocessor detects DCE/RPC requests 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 client requests. It also analyzes SMB data streams and detects anomalous SMB behavior and evasion techniques. It determines if SMB traffic contains a DCE/RPC request, continues processing if it does, and ceases processing if it does not.

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 since, ultimately, IP transports all DCE/RPC traffic. See Using TCP Stream Preprocessing on page 627 and Defragmenting IP Packets on page 613.

Keysight, EX-1004
Page 514

Finally, the DCE/RPC preprocessor normalizes DCE/RPC traffic for processing by the rules engine. See DCE/RPC Keywords on page 832 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.

When enabled, the preprocessor normalizes traffic according to the configured options, but generates no events in most cases before sending the normalized traffic to the rules engine unless you also enable the accompanying rules with generator ID (GID) 133 on the Rule page. See Setting Rule States on page 435 for more information.

Note that when a shared object rule or standard text rule that requires this preprocessor is enabled in an intrusion policy, you must enable the preprocessor or choose to allow IPS to enable it automatically before you can save the policy. For more information, see Automatically Enabling Advanced Settings on page 490.

Normally the DCE/RPC preprocessor only processes traffic traveling over standard DCE/RPC ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify DCE/RPC traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the preprocessor engine checks each packet for service identifiers to see if the packet is DCE/RPC traffic. If the identifier is found, the packet is processed even if it came over a non-standard DCE/RPC port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

See the following sections for more information:

- Selecting Global DCE/RPC Options on page 515
- Understanding Target-Based DCE/RPC Server Policies on page 517
- Understanding DCE/RPC Transports on page 518
- Selecting DCE/RPC Target-Based Policy Options on page 522
- Configuring the DCE/RPC Preprocessor on page 526

Selecting Global DCE/RPC Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 515

them unless you have a thorough understanding of the preprocessor and the interaction between the preprocessor and enabled DCE/RPC rules.

Maximum Fragment Size

When **Enable Defragmentation** is selected, specifies the maximum DCE/RPC fragment length allowed between 1514 and 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 or a value of 0 disables this option.

You should not adjust this option unless you have a thorough understanding of the DCE/RPC rules that are enabled in the intrusion policy. In particular, make sure that the maximum fragment size is greater than or equal to the depth to which the rules need to detect. For more information, see Constraining Content Matches on page 778 and Using Byte_Jump and Byte_Test on page 794.

This option does not generate events.

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.

You should not adjust this option unless you have a thorough understanding of the DCE/RPC rules that are enabled in the intrusion policy. In particular, make sure that the reassembly threshold is greater than or equal to the depth to which the rules need to detect. For more information, see Constraining Content Matches on page 778 and Using Byte_Jump and Byte_Test on page 794.

This option does not generate events.

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.

This option does not generate events.

Keysight, EX-1004
Page 516

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.

This option generates events when rule 133:1 is enabled.

Understanding Target-Based DCE/RPC Server Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can create one or more target-based server policies to configure the DCE/RPC preprocessor to look at DCE/RPC requests the same way they are looked at by receiving hosts acting as DCE/RPC servers on your network. 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 requests 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 requests to other hosts running different Windows or Samba versions. The default target-based policy applies to any host not included in another target-based policy.

You select one of the following Windows or Samba versions for each target-based policy:

- Windows 2000
- Windows 2003
- Windows XP
- Windows Vista
- Samba (3.0.23 and later)

Keysight, EX-1004
Page 517

- Samba 3.0.21 and 3.0.22)
- Samba (3.0.20 and earlier)

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 518 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.

Understanding DCE/RPC Transports

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 of 1024-65535 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 522 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 637 and Selecting Stream Reassembly Options on page 638 for more information.

Keysight, EX-1004
Page 518

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 in either of the following cases where it is not necessary to enable any transport in the default target-based policy:

- if 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.
- if you have enabled adaptive profiles in your intrusion policy

  Normally the DCE/RPC preprocessor only processes traffic traveling over standard DCE/RPC ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify DCE/RPC traffic travelling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the preprocessor engine checks each packet for service identifiers to see if the packet is DCE/RPC traffic. If the identifier is found, the packet is processed even if it came over a non-standard DCE/RPC port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

See the following sections for more information:

- Understanding Connectionless and Connection-Oriented DCE/RPC Traffic on page 519
- Understanding the RPC over HTTP Transport on page 521

Understanding Connectionless and Connection-Oriented DCE/RPC Traffic

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

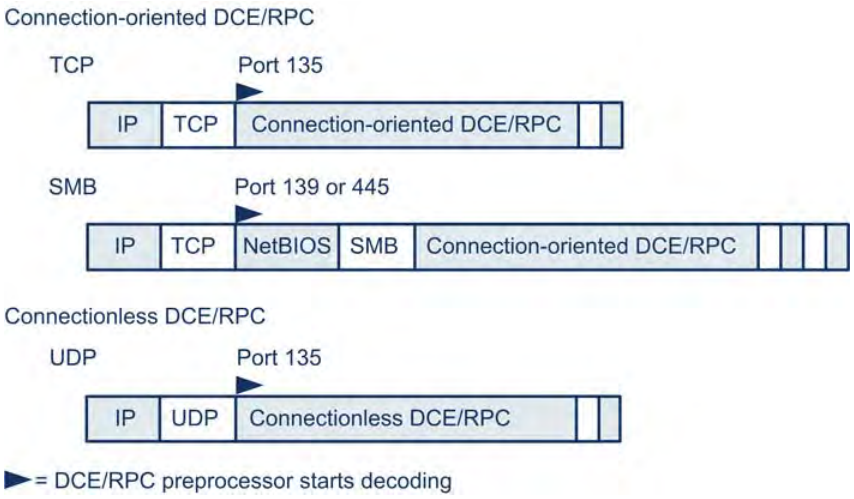
  The DCE/RPC preprocessor detects connectionless DCE/RPC in the UDP transport.

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

Keysight, EX-1004
Page 519

by connectionless DCE/RPC header values; in contrast, the transport protocol ensures correct fragment order for connection-oriented DCE/RPC. The DCE/RPC preprocessor uses these and other protocol-specific characteristics to monitor both protocols for anomalies and other evasion techniques, and to decode and defragment traffic before passing it to the rules engine.

The following diagram illustrates the point at which the DCE/RPC preprocessor begins processing DCE/RPC traffic for the different transports.



Note the following in the figure:

- The well-known TCP or UDP port 135 identifies DCE/RPC traffic in the TCP and UDP transports.
- The figure does not include RPC over HTTP.

  For RPC over HTTP, connection-oriented DCE/RPC is transported directly over TCP as shown in the figure after an initial setup sequence over HTTP. See Understanding the RPC over HTTP Transport on page 521 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 613 for more information.

Keysight, EX-1004
Page 520

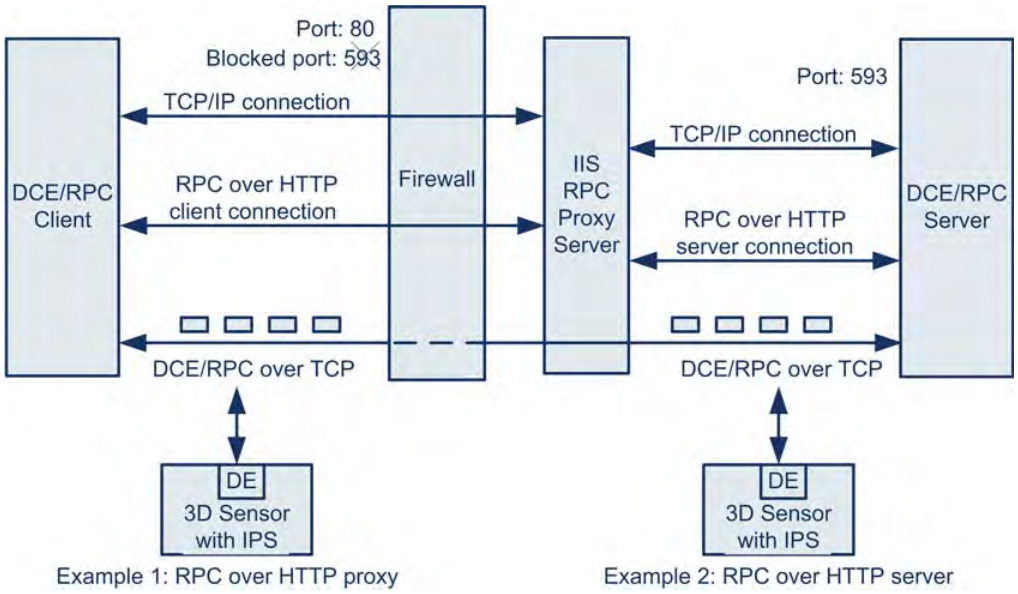
- TCP transports all connection-oriented DCE/RPC.

  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 627 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 645 for more information.

Understanding the RPC over HTTP Transport

Microsoft RPC over HTTP allows you to tunnel DCE/RPC traffic through a firewall as shown in the following diagram. The DCE/RPC preprocessor detects version 1 of Microsoft RPC over HTTP.



The Microsoft IIS proxy server and the DCE/RPC server can be on the same host or on different hosts. Separate proxy and server options provide for both cases. Note the following in the figure:

- The DCE/RPC server monitors port 593 for DCE/RPC client traffic, but the firewall blocks port 593.

  Firewalls typically block port 593 by default.
- RPC over HTTP transports DCE/RPC over HTTP using well-known HTTP port 80, which firewalls are likely to permit.
- Example 1 shows that you would select the **RPC over HTTP proxy** option when the detection engine (DE) monitors traffic between the DCE/RPC client and the MicroSoft IIS RPC proxy server.

Keysight, EX-1004
Page 521

- 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 detection engine 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Networks

The host or hosts where you want to apply the DCE/RPC target-based server policy.

You can specify a single IP address or CIDR 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 CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

IMPORTANT! The `default` setting in the default policy specifies all IP addresses (0.0.0.0/0) 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 for the default policy, and you cannot specify 0.0.0.0/0 or leave this setting blank in another policy.

This option does not generate events.

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 517 for detailed information on these policies.

This option does not generate events.

Keysight, EX-1004
Page 522

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$".

This option generates events when you enable rule 133:26.

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.

This option generates events when you enable rule 133:20.

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.

This option does not generate events.

Keysight, EX-1004
Page 523

RPC over HTTP Proxy Ports

Enables detection of DCE/RPC traffic tunneled by RPC over HTTP over each specified port when your Sourcefire 3D Sensor is positioned between the DCE/RPC client and the MicroSoft IIS RPC proxy server. See Understanding the RPC over HTTP Transport on page 521.

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.

This option does not generate events.

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 detection engine monitors traffic between the two servers. See Understanding the RPC over HTTP Transport on page 521.

Typically, when you enable this option you should also enable **RPC over HTTP Server Auto-Detect Ports** with a port range of 1025-65535 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.

This option does not generate events.

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 of 1025-65535 for that option.

This option does not generate events.

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 of 1025-65535 for that option.

This option does not generate events.

Keysight, EX-1004
Page 524

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.

This option does not generate events.

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 Sourcefire 3D Sensor is positioned between the DCE/RPC client and the MicroSoft IIS RPC proxy server. See Understanding the RPC over HTTP Transport on page 521.

When enabled, you would typically specify a port range of 1025-65535 to cover the entire range of ephemeral ports.

This option does not generate events.

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 detection engine monitors traffic between the two servers. See Understanding the RPC over HTTP Transport on page 521.

Typically, when you enable **RPC over HTTP Server Ports** you also enable this option, even if you are not aware of any proxy web servers on your network, and specify a port range of 1025-65535 for this option to cover the entire range of ephemeral ports.

This option does not generate events.

TCP Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in TCP on the specified ports.

Typically, when **TCP Ports** are enabled, you also enable this option and specify a port range of 1025-65535 to cover the entire range of ephemeral ports.

This option does not generate events.

UDP Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in UDP on each specified port.

Typically, when **UDP Ports** are enabled, you also enable this option and specify a port range of 1025-65535 to cover the entire range of ephemeral ports.

This option does not generate events.

SMB Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in SMB.

Keysight, EX-1004
Page 525

Typically, use the default settings.

This option does not generate events.

Configuring the DCE/RPC Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 515 and Selecting DCE/RPC Target-Based Policy Options on page 522 for rules associated with specific detection options; see also Setting Rule States on page 435.

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
Connection-Oriented DCE/RPC	133:27 through 133:39
Detect Connectionless DCE/RPC	133:40 through 133:43

If you are managing IPS using a Defense Center with RNA, note that you can process DCE/RPC traffic using non-standard ports. To check each packet for DCE/RPC service identifiers and process packets where they are found, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

To configure the DCE/RPC preprocessor:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 526

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

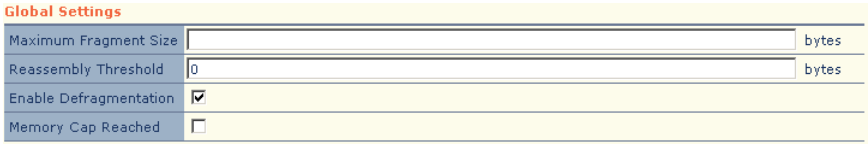
3. Click **Advanced Settings** in the navigation panel on the left.

   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**.

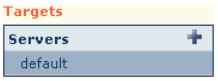
   The DCE/RPC Configuration page appears.



   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 494 for more information.

5. You can modify any of the options described in Selecting Global DCE/RPC Options on page 515.

   Note that, 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.



Keysight, EX-1004
Page 527

6. You have two options:

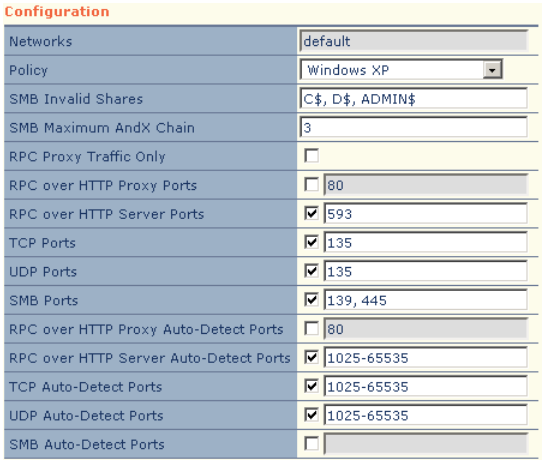
   - Add a new target-based policy. Click the **+** sign next to **Servers** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses in the **Server Address** field and click **OK**.

     You can specify a single IP address or CIDR block, or a comma-separated list comprised of either or both. You can configure up to 255 policies, including the default policy. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     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.



Keysight, EX-1004
Page 528

7. You can modify any of the following target-based policy options. See Selecting DCE/RPC Target-Based Policy Options on page 522 for more information:

   - To specify the host or hosts where you want to apply the DCE/RPC target-based server policy, enter a single IP address or CIDR 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. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
   - 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.
   - 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.

Keysight, EX-1004
Page 529

- 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 518, Understanding the RPC over HTTP Transport on page 521, and Selecting DCE/RPC Target-Based Policy Options on page 522 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 of 1025-65535 for auto-detection ports that you enable to cover the entire range of ephemeral ports. See Understanding DCE/RPC Transports on page 518, Understanding the RPC over HTTP Transport on page 521, and Selecting DCE/RPC Target-Based Policy Options on page 522 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.

8. Optionally, click **Configure Rules for DCE/RPC Configuration** at the top of the page to display rules associated with individual options.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Selecting Global DCE/RPC Options on page 515, Selecting DCE/RPC Target-Based Policy Options on page 522, and the Traffic-Associated DCE/RPC Rules table on page 526.

   Click **Back** to return to the DCE/RPC Configuration page.

Keysight, EX-1004
Page 530

9. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Detecting Exploits in DNS Name Server Responses

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Normally the DNS preprocessor only processes traffic traveling over standard DNS ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify DNS traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the preprocessor

Keysight, EX-1004
Page 531

engine checks each packet for service identifiers to see if the packet is DNS traffic. If the identifier is found, the packet is processed even if it came over a non-standard DNS port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

See the following sections for more information:

- Understanding DNS Preprocessor Resource Record Inspection on page 532
- Detecting Overflow Attempts in RData Text Fields on page 534
- Detecting Obsolete DNS Resource Record Types on page 534
- Detecting Experimental DNS Resource Record Types on page 535
- Configuring the DNS Preprocessor on page 535

Understanding DNS Preprocessor Resource Record Inspection

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

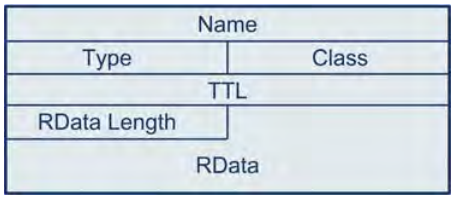
This section...	Includes...	For example...
Answer	Optionally, one or more resource records that provide a specific answer to a query	The IP address corresponding to a domain name

Keysight, EX-1004
Page 532

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:



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. You must enable TCP stream preprocessing 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 627 and Using UDP Stream Preprocessing on page 645.

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.

Keysight, EX-1004
Page 533

Detecting Overflow Attempts in RData Text Fields

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Detecting Obsolete DNS Resource Record Types

Platform: DC/MDC/Standalone Sensor

Component: IPS

RFC 1035 identifies several resource record types as obsolete. Because these are obsolete record types, some systems do not account for them and can 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

For details on configuring DNS obsolete record detection, see Configuring the DNS Preprocessor on page 535.

Keysight, EX-1004
Page 534

Detecting Experimental DNS Resource Record Types

Platform: DC/MDC/Standalone Sensor

Component: IPS

RFC 1035 identifies several resource record types as experimental. Because these are experimental record types, some systems do not account for them and can 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

For details on configuring DNS experimental record detection, see Configuring the DNS Preprocessor on page 535.

Configuring the DNS Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

Use the following procedure to configure the DNS preprocessor.

If you are managing IPS using a Defense Center with RNA, note that you can process DNS traffic using non-standard ports. If you want to check each packet for DNS service identifiers, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

To configure the DNS preprocessor:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 535

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

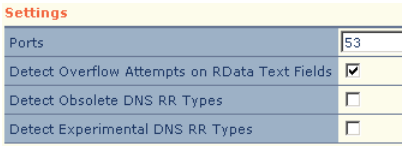
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



   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 494 for more information.

5. Optionally, modify the source port or ports the DNS preprocessor should monitor for DNS server responses. Separate multiple ports with commas.

6. Optionally, select the **Detect Overflow Attempts on RData Text fields** check box to enable detection of buffer overflow attempts in RData text fields.

   See Detecting Overflow Attempts in RData Text Fields on page 534 for more information.

   Note that you must set the accompanying rule (GID: 131, SID: 3) to Generate Events or, optionally, to Drop and Generate Events in an inline deployment, if you want the preprocessor to log intrusion events for this option.

7. Optionally, select the **Detect Obsolete DNS RR Types** check box to enable detection of obsolete resource record types.

   See Detecting Obsolete DNS Resource Record Types on page 534 for more information.

   Note that you must set the accompanying rule (GID: 131, SID: 1) to Generate Events or, optionally, to Drop and Generate Events in an inline deployment, if you want the preprocessor to log intrusion events for this option.

Keysight, EX-1004
Page 536

8. Optionally, select the **Detect Experimental DNS RR Types** check box to detect experimental resource record types.

   See Detecting Experimental DNS Resource Record Types on page 535 for more information.

   Note that you must set the accompanying rule (GID: 131, SID: 2) to Generate Events or, optionally, to Drop and Generate Events in an inline deployment, if you want the preprocessor to log intrusion events for this option.

9. Optionally, click **Configure Rules for DNS Configuration** at the top of the page.

   A filtered view appears of DNS 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 435 for more information.

   Click **Back** to return to the DNS Configuration page.

10. You have the following options:

    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

    - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 537

Decoding FTP and Telnet Traffic

Platform: DC/MDC/Standalone Sensor

Component: IPS

The FTP/Telnet decoder analyzes FTP and telnet data streams, normalizing FTP and telnet commands prior to processing by the rules engine. You must ensure that TCP stream preprocessing is enabled to use the FTP/Telnet decoder.

When enabled, the decoder normalizes traffic according to the configured options, but generates no events in most cases before sending the normalized traffic to the rules engine unless you also enable the accompanying rules with generator ID (GID) 125 (FTP) and 126 (Telnet) on the Rule page. 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 435 for more information.

For more information, see the following topics:

- Understanding Global FTP and Telnet Options on page 538
- Configuring Global FTP/Telnet Options on page 539
- Understanding Telnet Options on page 542
- Configuring Telnet Options on page 543
- Understanding Server-Level FTP Options on page 546
- Configuring Server-Level FTP Options on page 550
- Understanding Client-Level FTP Options on page 555
- Configuring Client-Level FTP Options on page 557

Understanding Global FTP and Telnet Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Normally the FTP/Telnet decoder only processes traffic traveling over standard FTP or Telnet ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify FTP or Telnet traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the decoder engine checks each packet for service identifiers to see if the packet is FTP or Telnet traffic. If the identifier is found, the packet is processed even if it came over a non-standard FTP or Telnet port. For more information on how

Keysight, EX-1004
Page 538

adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

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.

This option does not generate events.

Detect Encrypted Traffic

Detects encrypted telnet and FTP sessions.

This option generates events when you enable rules 125:7 and 126:2.

Continue to Inspect Encrypted Data

Instructs the preprocessor to continue checking a data stream after it is encrypted, looking for eventual decrypted data.

This option does not generate events.

Configuring Global FTP/Telnet Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 538.

If you are managing IPS using a Defense Center with RNA, note that you can decode FTP or Telnet traffic using non-standard ports. To check each packet for FTP and Telnet service identifiers and process packets where they are found, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

To configure global options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 539

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 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 494 for more information.

TIP! For more information on configuring the other options on this page, see Configuring Telnet Options on page 543, Configuring Server-Level FTP Options on page 550, and Configuring Client-Level FTP Options on page 557.



Keysight, EX-1004
Page 540

5. Optionally, you can modify any of the following in the Global Settings page area:

   - Select or clear **Stateful Inspection**:

     To examine reassembled TCP streams containing FTP packets, select **Stateful Inspection**.

     To inspect only unreassembled packets, clear **Stateful Inspection**.

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 627 and Reassembling TCP Streams on page 637.

   - Select or clear the **Detect Encrypted Traffic**:

     To detect encrypted traffic, select **Detect Encrypted Traffic**.

     To ignore encrypted traffic, clear **Detect Encrypted Traffic**.

     Note that you must enable the accompanying FTP rule (GID:125, SID:7) and telnet rule (GID:126, SID:2) before the preprocessor, when enabled, will generate events for the corresponding traffic. Click **Configure Rules for FTP and Telnet Configuration** at the top of the page to display 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 435 for more information.
   - 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.

Keysight, EX-1004
Page 541

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

   For more information on telnet options, see Understanding Telnet Options on page 542. For more information on additional FTP options, see Understanding Server-Level FTP Options on page 546 and Understanding Client-Level FTP Options on page 555.

Understanding Telnet Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 542

Ports

Indicates the ports whose telnet traffic you want to normalize. In the interface, list multiple ports separated by commas.

IMPORTANT! Any ports you add to the telnet **Ports** list should also be added to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see Reassembling TCP Streams on page 637.

This option does not generate events.

Normalize

Normalizes telnet traffic to the specified ports.

This option does not generate events.

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. Since FTP uses the Telnet protocol on the control connection, it is also susceptible to this behavior.

This option generates events when you enable rule 126:3.

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.

This option generates events when you enable rule 126:1.

Configuring Telnet Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 542.

Keysight, EX-1004
Page 543

To configure telnet options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

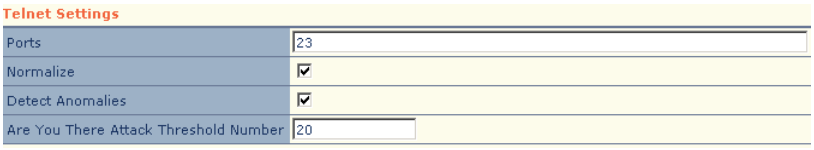
4. You have two choices, depending on whether **FTP and Telnet Configuration** under Application Layer Preprocessors is enabled:

   - If the configuration is enabled, click **Edit**.
   - If the configuration is disabled, click **Enabled**, then click **Edit**.

   The FTP and Telnet 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 494 for more information.

TIP! For more information on configuring the other options on this page, see Configuring Global FTP/Telnet Options on page 539, Configuring Server-Level FTP Options on page 550, and Configuring Client-Level FTP Options on page 557.



Keysight, EX-1004
Page 544

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 637.

   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.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Understanding Telnet Options on page 542.

   Click **Back** to return to the FTP and Telnet Configuration page.

Keysight, EX-1004
Page 545

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

   For more information on global FTP and telnet options, see Understanding Global FTP and Telnet Options on page 538. For more information on additional FTP options, see Understanding Server-Level FTP Options on page 546 and Understanding Client-Level FTP Options on page 555.

Understanding Server-Level FTP Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 546

Networks

Use this option to specify one or more IP addresses of FTP servers.

You can specify a single IP address or CIDR block, or a comma-separated list of either or both up to 1024 characters, and you can specify up to 255 profiles including the default profile. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

IMPORTANT! The `default` setting in the default profile specifies all IP addresses (0.0.0.0/0) 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 for the default profile, and you cannot specify 0.0.0.0/0 or leave this setting blank in another profile.

This option does not generate events.

Ports

Use this option to specify the ports on the FTP server where the sensor should monitor traffic. In the interface, list multiple ports separated by commas.

IMPORTANT! Any ports you add to the server-level FTP **ports** list should also be added to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see Reassembling TCP Streams on page 637.

This option does not generate events.

Log FTP Command Validation Configuration

Use this option to turn on printing of the configuration information for each ftp command listed for this server.

This option does not generate events.

Additional FTP Commands

Use this line to specify the additional commands that the decoder should detect. Separate additional commands by spaces.

This option does not generate events.

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.

This option generates events when you enable rule 125:3.

Keysight, EX-1004
Page 547

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.

This option does not generate events.

Check Commands for String Format Attacks

Use this option to check the specified commands for string format attacks.

This option generates events when you enable rule 125:5.

Command Validity

Use this option to enter a valid format for a specific command. See Creating FTP Command Parameter Validation Statements on page 548 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.

This option generates events when you enable rules 125:2 and 125:4.

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.

This option does not generate events.

Detect Telnet Escape Codes within FTP Commands

Use this option to detect when telnet commands are used over the FTP command channel.

This option generates events when you enable rule 125:1.

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.

This option does not generate events.

Creating FTP Command Parameter Validation Statements

When setting up a validation statement for an FTP command, you can specify a group of alternative parameters. You can also create a binary OR relationship between two parameters by separating them with a pipe character (|) in the

Keysight, EX-1004
Page 548

validation statement. Surrounding parameters by square brackets ([]) indicates that those parameters are optional.

You can create FTP command parameter validation statements to validate the syntax of a parameter received as part of an FTP communication.

Any of the parameters listed in the following table can be used in FTP command parameter validation statements:

If you use...	The following validation occurs...
`int`	The represented parameter must be an integer.
`number`	The represented parameter must be an integer between 1 and 255.
`char _chars`	The represented parameter must be a single character and a member of the characters specified in the *`_chars`* argument. For example, defining the command validity for `MODE` with the validation statement `char SBC` checks that the parameter for the `MODE` command comprises the character `S` (representing Stream mode), the character `B` (representing Block mode), or the character `C` (representing Compressed mode).
`date _datefmt`	If *`_datefmt`* contains `#`, the represented parameter must be a number. If *`_datefmt`* contains `C`, the represented parameter must be a character. If *`_datefmt`* contains literal strings, the represented parameter must match the literal string.
`string`	The represented parameter must be a string.
`host_port`	The represented parameter must be a valid host port specifier as defined by RFC 959, the File Transfer Protocol specification by the Network Working Group.

You can combine the syntax in the table above as needed to create parameter validation statements that correctly validate each FTP command where you need to validate traffic.

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`.

Keysight, EX-1004
Page 549

To group multiple parameters:

Access: P&R Admin/Admin

- Separate the parameters by spaces.

  For the `TYPE` command, place a space after the opening bracket and a space before the closing bracket.

  For example, to evaluate the `MODE` command parameter to see if it is a character that represents both the Stream mode and the Block mode, type this command validation statement `char S,B`.

To accept one or the other of two parameters:

Access: P&R Admin/Admin

- Separate the parameters by a pipe character (|).

  For the `TYPE` command, place a space before and after the pipe character.

  For example, to evaluate the `MODE` command parameter to see if it is a character that represents either the Stream mode or the Block mode, type this command validation statement `char S|B`.

To designate a parameter as optional:

Access: P&R Admin/Admin

- Enclose the parameter in square brackets ([]).

  For the `TYPE` command, place a space before and after the pipe character.

See Understanding Server-Level FTP Options on page 546 for more information.

Configuring Server-Level FTP Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 546.

To configure server-level FTP options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 550

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 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 494 for more information.

TIP! For more information on configuring the other options on this page, see Configuring Global FTP/Telnet Options on page 539, Configuring Telnet Options on page 543, and Configuring Client-Level FTP Options on page 557.



Keysight, EX-1004
Page 551

5. You have two options:

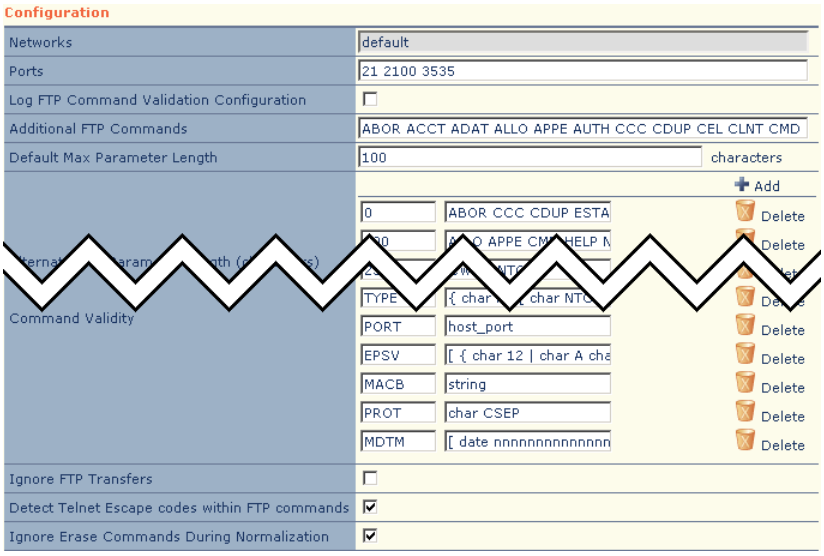
   - Add a new server profile. Click the **+** sign next to **FTP Server** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**.

     You can specify a single IP address or CIDR block, or a comma-separated list comprised 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 CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     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.



Keysight, EX-1004
Page 552

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 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 637.

   - To turn on printing of the configuration information for each FTP command listed for this server, check **Log FTP Command Validation Configuration**.
   - 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.

     If the parameter length for a command exceeds this threshold and, the accompanying preprocessor rule with generator ID (GID) 126 and Snort ID (SID) 3 is enabled, the preprocessor generates an event.
   - 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.

Keysight, EX-1004
Page 553

- 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 546.
- 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 538.

- 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.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Understanding Server-Level FTP Options on page 546.

   Click **Back** to return to the FTP and Telnet Configuration page.

8. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 554

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

For more information on telnet options, see Understanding Telnet Options on page 542. For more information on additional FTP options, see Understanding Global FTP and Telnet Options on page 538 and Understanding Client-Level FTP Options on page 555.

Understanding Client-Level FTP Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Networks

Use this option to specify one or more IP addresses of FTP clients.

You can specify a single IP address or CIDR block, or a comma-separated list of either or both up to 1024 characters, and you can specify up to 255 profiles including the default profile. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

IMPORTANT! The `default` setting in the default profile specifies all IP addresses (0.0.0.0/0) 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 for the default profile, and you cannot specify 0.0.0.0/0 or leave this setting blank in another profile.

This option does not generate events.

Max Response Length

Use this option to specify the maximum length of a response string from the FTP client.

This option generates events when you enable rule 125:6.

Keysight, EX-1004
Page 555

Detect FTP Bounce Attempts

Use this option to detect FTP bounce attacks.

This option generates events when you enable rule 125:8.

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.

This option does not generate events.

Detect Telnet Escape Codes within FTP Commands

Use this option to detect when telnet commands are used over the FTP command channel.

This option generates events when you enable rule 125:1.

Ignore Erase Commands During Normalization

When **Detect Telnet Escape Codes within FTP Commands** is enabled, 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.

This option does not generate events.

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.

This option does not generate events.

Keysight, EX-1004
Page 556

Configuring Client-Level FTP Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 555.

To configure client-level FTP options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

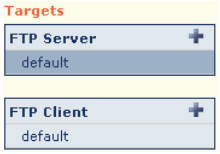
3. Click **Advanced Settings** in the navigation panel on the left.

   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 page appears.



Keysight, EX-1004
Page 557

5. You have two options:

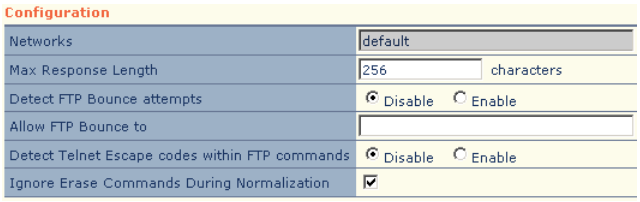
   - Add a new client profile. Click the **+** sign next to **FTP Client** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Client Address** field and click **OK**.

     You can specify a single IP address or CIDR block, or a comma-separated list comprised 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 CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     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.



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, enable **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.

Keysight, EX-1004
Page 558

- 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 in the Sourcefire 3D System, see IP Address Conventions on page 59.

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, enable **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.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Understanding Client-Level FTP Options on page 555.

   Click **Back** to return to the FTP and Telnet Configuration page.

Keysight, EX-1004
Page 559

8. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

   For more information on telnet options, see Understanding Telnet Options on page 542. For more information on additional FTP options, see Understanding Server-Level FTP Options on page 546 and Understanding Global FTP and Telnet Options on page 538.

Decoding HTTP Traffic

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 560

- 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.

The preprocessor engine performs HTTP normalization *statelessly*. That is, it normalizes HTTP strings on a packet-by-packet basis, and can only process HTTP strings that have been reassembled by the TCP stream preprocessor. You must ensure that TCP stream preprocessing is enabled to use the HTTP preprocessor.

When enabled, the preprocessor normalizes traffic according to the configured options, but generates no events in most cases before sending the normalized traffic to the rules engine unless you also enable the accompanying rules with generator ID (GID) 119 on the Rule page. A link on the configuration page takes you to a filtered view of HTTP Inspect 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 435 for more information.

Note that when a rule that requires this preprocessor is enabled in an intrusion policy, you must enable the preprocessor or choose to allow IPS to enable it automatically before you can save the policy. For more information, see Automatically Enabling Advanced Settings on page 490.

Normally the HTTP preprocessor only processes traffic traveling over standard HTTP ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify HTTP traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the preprocessor engine checks each packet for service identifiers to see if the packet is HTTP traffic. If the identifier is found, the packet is processed even if it came over a non-standard HTTP port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

See the following sections for more information:

- Selecting Global HTTP Normalization Options on page 562
- Configuring Global HTTP Configuration Options on page 563
- Selecting Server-Level HTTP Normalization Options on page 564
- Selecting Server-Level HTTP Normalization Encoding Options on page 570
- Configuring Server-Level HTTP Configuration Options on page 574
- Enabling Additional HTTP Inspect Preprocessor Rules on page 577

Keysight, EX-1004
Page 561

Selecting Global HTTP Normalization Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

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.

This option generates events when you enable rule 120:1.

Detect HTTP Proxy Servers

Detects HTTP traffic using proxy servers not defined by the **Allow HTTP Proxy Use** option.

This option generates events when you enable rule 119:17.

Maximum Compressed Data Depth

Sets the maximum size of compressed data to decompress when **Inspect Compressed Data** is enabled. You can specify 1 to 65535 bytes. See Selecting Server-Level HTTP Normalization Options on page 564 for more information.

This option does not generate events.

Maximum Decompressed Data Depth

Sets the maximum size of the normalized decompressed data when **Inspect Compressed Data** is enabled. You can specify 1 to 65535 bytes. See Selecting Server-Level HTTP Normalization Options on page 564 for more information.

This option does not generate events.

Keysight, EX-1004
Page 562

Configuring Global HTTP Configuration Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure detection of HTTP traffic to non-standard ports and on HTTP traffic using proxy servers.

If you are managing IPS using a Defense Center with RNA, note that you can process HTTP traffic using non-standard ports. To check each packet for HTTP service identifiers and process packets where they are found, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

To configure global HTTP configuration options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

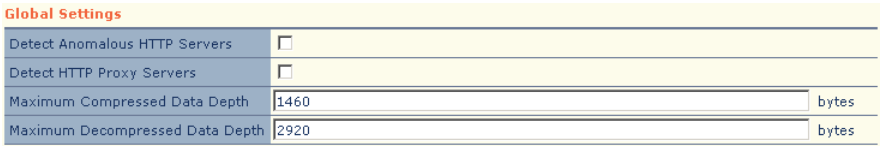
3. Click **Advanced Settings** in the navigation panel on the left.

   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 562.

Keysight, EX-1004
Page 563

6. Optionally, click **Configure Rules for HTTP Configuration** at the top of the page to display rules associated with individual options.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Selecting Global HTTP Normalization Options on page 562.

   Click **Back** to return to the HTTP Configuration page.

7. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Selecting Server-Level HTTP Normalization Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 564

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.

Networks

Use this option to specify one or more servers.

You can specify a single IP address or CIDR block, or a comma-separated list of either or both. 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 CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

IMPORTANT! The `default` setting in the default profile specifies all IP addresses (0.0.0.0/0) 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 for the default profile, and you cannot specify 0.0.0.0/0 or leave this setting blank in another profile.

This option does not generate events.

Ports

The ports whose HTTP traffic the preprocessor engine normalizes. Separate multiple port numbers with commas.

IMPORTANT! Any ports you add to the server-level HTTP port list should also be added to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see Reassembling TCP Streams on page 637.

This option does not generate events.

Oversize Dir Length

Detects URL directories longer than the specified value.

This option generates events when you enable rule 119:15.

Keysight, EX-1004
Page 565

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 783 for more information.

You can specify a value -1 through 1460. Sourcefire recommends that you set client flow depth to its maximum value. Specify any of the following:

- 1-1460 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.

This option does not generate events.

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 783 for more information.

Unlike client flow depth, server flow depth specifies the number of bytes per TCP session, not per HTTP packet, for rules to inspect.

Keysight, EX-1004
Page 566

You can specify a value -1 through 65535. Sourcefire recommends that you set the server flow depth to its maximum value. You can specify any of the following:

- 1-65535.

  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 can 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 can 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.**

This option does not generate events.

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 of 1 to 65535 to enable it.

This option generates events when you enable rule 119:19.

Maximum Number of Headers

Detects when the number of headers exceeds this setting in an HTTP request. Specify a value of 1 to 1024 to enable it.

This option generates events when you enable rule 119:20.

Keysight, EX-1004
Page 567

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 783 for more information.

Specify a value of -1 to 65535. 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 of 0 to 65535 for the **HTTP Client Body** option to function in an intrusion rule.

This option does not generate events.

HTTP Methods

Specifies HTTP request methods for the preprocessor to inspect in addition to GET and POST. Use a comma to separate multiple values.

This option does not generate events.

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.

This option does not generate events.

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.

This option does not generate events.

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.

This option does not generate events.

Keysight, EX-1004
Page 568

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.

This option does not generate events.

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.

This option does not generate events.

Inspect URI Only

Inspects only the URI portion of the normalized HTTP request packet.

This option does not generate events.

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 IPS 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 783, Generating Events on the HTTP Encoding Type and Location on page 855, and Pointing to a Specific Payload Type on page 857.

This option generates events when you enable rules 120:2 and 120:3.

Normalize UTF Encodings to UTF-8

When **Inspect HTTP Responses** is enabled, detects UTF-16LE, UTF-16-BE, UTF-32LE, and UTF32BE encodings in HTTP responses and normalizes them to UTF-8.

This option generates events when you enable rule 120:4.

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. IPS inspects chunked and non-chunked HTTP response data. IPS inspects decompressed data packet by packet across multiple packets as needed; that is, IPS 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

Keysight, EX-1004
Page 569

decompressed data ends when **Server Flow Depth** is reached unless you also select **Inspect Unlimited Packets**. You can use the `file_data` rule keyword to inspect decompressed data; see Pointing to a Specific Payload Type on page 857 for more information.

This option does not generate events.

Inspect Unlimited Packets

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. See Selecting Global HTTP Normalization Options on page 562.

This option does not generate events.

Extract Original Client IP Address

Extracts the original client IP address from the X-Forwarded-For (XFF) or True-Client-IP HTTP headers. You can display the extracted original client IP address in the intrusion events table view. See Understanding Intrusion Events on page 276 for more information.

Profile

Specifies the types of encoding that are normalized for HTTP traffic. IPS 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 570 for more information.

This option does not generate events.

Selecting Server-Level HTTP Normalization Encoding Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

ASCII Encoding

Decodes encoded ASCII characters and specifies whether the rules engine generates an event on ASCII-encoded URIs.

This option generates events when you enable rule 119:1.

Keysight, EX-1004
Page 570

UTF-8 Encoding

Decodes standard UTF-8 Unicode sequences in the URI.

This option generates events when you enable rule 119:6.

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.

This option generates events when you enable rule 119:3.

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.

This option generates events when you enable rule 119:4.

Base 36 Encoding

Decodes base 36 encoding, which mainly appears in Asian versions of IIS. This option does not work with %u or utf_8 options enabled.

This option generates events when you enable rule 119:5.

Microsoft IIS Encoding

Decodes using Unicode codepoint mapping.

TIP! Sourcefire recommends enabling this option, because it is seen mainly in attacks and evasion attempts.

This option generates events when you enable rule 119:7.

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 only found in attack scenarios.

Keysight, EX-1004
Page 571

This option generates events when you enable rule 119:2.

Multi-Slash Obfuscation

Normalizes multiple slashes in a row into a single slash.

This option generates events when you enable rule 119:8.

IIS Backslash Obfuscation

Normalizes backslashes to forward slashes.

This option generates events when you enable rule 119:9.

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.

This option generates events when you enable rules 119:10 and 119:11.

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 whitespace if a space character (0x20) precedes it.

This option generates events when you enable rule 119:12.

Invalid RFC Delimiter

Normalizes line breaks (\n) in URI data.

This option generates events when you enable rule 119:13.

Webroot Directory Traversal

Detects directory traversals that traverse past the initial directory in the URL.

This option generates events when you enable rule 119:18.

Keysight, EX-1004
Page 572

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 whitespace if a space character (0x20) precedes it.

This option does not generate events.

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.

This option generates events when you enable rule 119:14.

Max Chunk Encoding Size

Detects abnormally large chunk sizes in URI data.

This option generates events when you enable rules 119:16 and 119:22.

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.

This option does not generate events.

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.

This option does not generate events.

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.

This option does not generate events.

Keysight, EX-1004
Page 573

Configuring Server-Level HTTP Configuration Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

Use the following procedure to configure server-level HTTP configuration options.

If you are managing IPS using a Defense Center with RNA, note that you can process HTTP traffic using non-standard ports. To check each packet for HTTP service identifiers and process packets where they are found, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

To configure server-level HTTP configuration options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

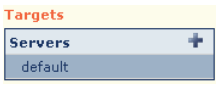
3. Click **Advanced Settings** in the navigation panel on the left.

   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 494 for more information.

Keysight, EX-1004
Page 574

5. You have two options:

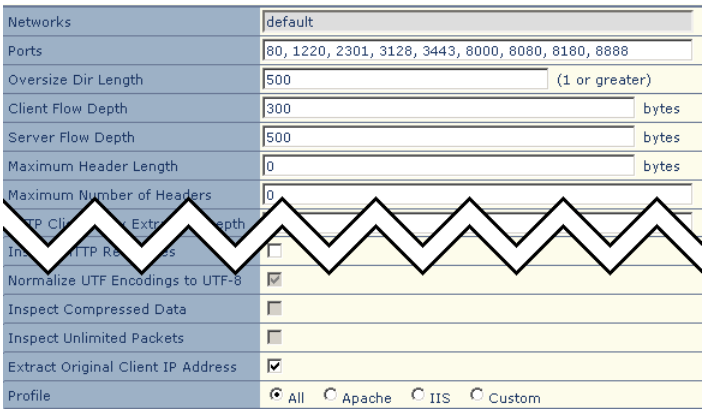
   - Add a new server profile. Click the **+** sign next to **Servers** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**.

     You can specify a single IP address or CIDR block, or a comma-separated list comprised 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 CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     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.

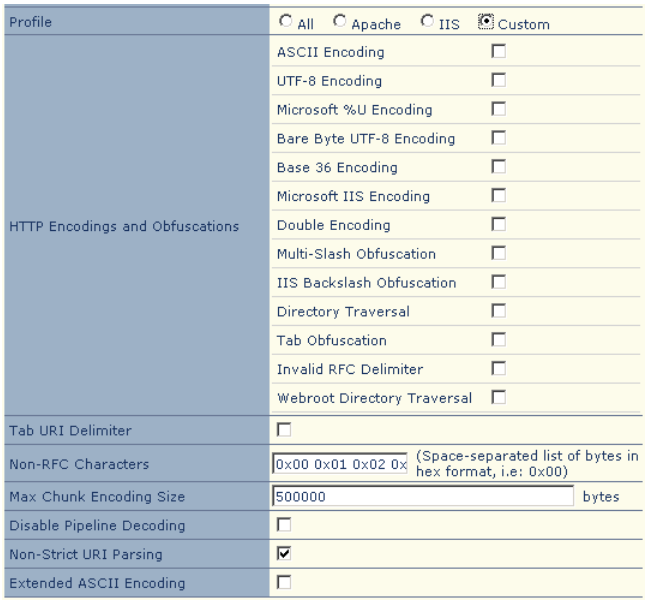
   The highlighted address updates on the left side of the page.

   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.

Keysight, EX-1004
Page 575

8. Select a server profile as follows:
   - Select **Custom** to create your own server profile (see Selecting Server-Level HTTP Normalization Options on page 564 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.
9. If you selected **Custom**, the custom options appear.



10. Configure the HTTP decoding options you want in your profile.

    See Selecting Server-Level HTTP Normalization Options on page 564 for details on available normalization options.
11. Optionally, click **Configure Rules for HTTP Configuration** at the top of the page to display rules associated with individual options.

    You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

    To identify the rules associated with different configuration options, see Selecting Server-Level HTTP Normalization Options on page 564 and Selecting Server-Level HTTP Normalization Encoding Options on page 570.

    Click **Back** to return to the HTTP Configuration page.

Keysight, EX-1004
Page 576

12. You have the following options:

    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

    - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Enabling Additional HTTP Inspect Preprocessor Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 577

preprocessor rules that are not associated with specific configuration options. See Setting Rule States on page 435 for more information.

Additional HTTP Inspect Preprocessor Rules

Preprocessor Rule GID:SID	Description
119:21	Detects an HTTP request header with more than one `content-length` field.
120:5	Detects UTF-7 encoding in HTTP response traffic; UTF-7 should only appear where 7-bit parity is required, such as in SMTP traffic.

Using the Sun RPC Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS provides RPC (Remote Procedure Call) normalization, which takes fragmented RPC records and normalizes them to a single record so the rules engine can inspect the complete record. For example, an attacker may attempt to discover the port where RPC `admind` runs. Some UNIX hosts use RPC `admind` to perform remote distributed system tasks. If the host performs weak authentication, a malicious user could take control of remote administration. The standard text rule (generator ID: 1) with the Snort ID (SID) 575 detects this attack by searching for content in specific locations to identify inappropriate `portmap GETPORT` requests.

Normally the RPC preprocessor only processes traffic traveling over standard RPC ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify RPC traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the preprocessor engine checks each packet for service identifiers to see if the packet is RPC traffic. If the identifier is found, the packet is processed even if it came over a non-standard RPC port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

Keysight, EX-1004
Page 578

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 ports you add to the RPC **ports** list should also be added to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see Reassembling TCP Streams on page 637.

Detect fragmented RPC records

Detects RPC fragmented records.

This option generates events when you enable rules 106:1 and 106:5.

Detect multiple records in one packet

Detects more than one RPC request per packet (or reassembled packet).

This option generates events when you enable rule 106:2.

Detect fragmented record sums which exceed one fragment

Detects reassembled fragment record lengths that exceed the current packet length.

This option generates events when you enable rule 106:3.

Detect single fragment records which exceed the size of one packet

Detects partial records

This option generates events when you enable rule 106:4.

Configuring the Sun RPC Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use the following procedure to configure the Sun RPC preprocessor.

To configure the Sun RPC preprocessor:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 579

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

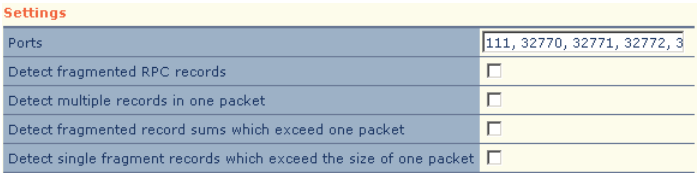
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



   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 494 for more information.

5. In the **Ports** field, type the port numbers where you want to decode RPC traffic. Separate multiple ports with commas.

   If you are managing IPS using a Defense Center with RNA, note that you can normalize RPC traffic using non-standard ports. To check each packet for RPC service identifiers, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

   For more information, see Using the Sun RPC Preprocessor on page 578.

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**

   See Using the Sun RPC Preprocessor on page 578 for more information.

Keysight, EX-1004
Page 580

7. Optionally, click **Configure Rules for Sun RPC Configuration** at the top of the page to display rules associated with individual options.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Using the Sun RPC Preprocessor on page 578.

   Click **Back** to return to the Sun RPC Configuration page.

8. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 581

Decoding SMTP Traffic

Platform: DC/MDC/Standalone Sensor

Component: IPS

The SMTP decoder instructs the rules engine to normalize SMTP commands. You must ensure that TCP stream preprocessing is enabled to use the SMTP decoder.

When enabled, the decoder normalizes traffic according to the configured options, but generates no events in most cases before sending the normalized traffic to the rules engine unless you also enable the accompanying rules with generator ID (GID) 124 on the Rule page. A link on the configuration page takes you to a filtered view of SMTP 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 435 for more information.

For more information, see the following sections:

- Understanding SMTP Decoding on page 582
- Configuring SMTP Decoding on page 585

Understanding SMTP Decoding

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can enable or disable normalization, and you can configure options to control the types of anomalous traffic the SMTP decoder detects. You must ensure that TCP stream preprocessing is enabled to use the SMTP decoder.

Normally the SMTP decoder only processes traffic traveling over standard SMTP ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify SMTP traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the decoder engine checks each packet for service identifiers to see if the packet is SMTP traffic. If the identifier is found, the packet is processed even if it came over a non-standard SMTP port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

WARNING! If you added a `preprocessor xlink2state` directive to the `user.conf` file on your appliance to enable the xlink2state preprocessor for a release prior to Version 4.5.x, delete that directive before attempting to configure SMTP preprocessing.

Keysight, EX-1004
Page 582

Ports

Specifies the ports whose SMTP traffic you want to normalize. Separate multiple ports with commas.

IMPORTANT! Any ports you add to the SMTP **ports** list should also be added to the TCP client reassembly port list. For more information on configuring TCP reassembly ports, see Reassembling TCP Streams on page 637.

This option does not generate events.

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.

This option does not generate events.

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**.

This option does not generate events.

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.

This option does not generate events.

Ignore Data

Does not process mail data; only processes MIME mail header data.

This option does not generate events.

Ignore TLS Data

Does not process data encrypted under the Transport Layer Security protocol.

This option does not generate events.

Keysight, EX-1004
Page 583

No Alerts

Disables intrusion events when accompanying preprocessor rules are enabled.

This option does not generate events.

Detect Unknown Commands

Detects unknown commands in SMTP traffic.

This option generates events when you enable rules 124:5 and124:6.

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.

This option generates events when you enable rule 124:1.

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.

This option generates events when you enable rules 124:2 and 124:7.

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.

This option generates events when you enable rule 124:3.

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.

This option generates events when you enable rule 124:3.

Invalid Commands

Detects if these commands are sent from the client side.

This option generates events when you enable rule 124:5 and 124:6.

Keysight, EX-1004
Page 584

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.

This option generates events when you enable rule 124:4.

Detect xlink2state

Detects packets that are part of X-Link2State Microsoft Exchange buffer data overflow attacks. For inline 3D Sensors with IPS, can also drop those packets.

This option generates events when you enable rule 124:8.

Enable MIME Decoding

Enables Base64 decoding of MIME email attachments. The preprocessor pipelines multiple Base64-encoded MIME attachments in one packet. You must also select the SMTP **Stateful Inspection** option to decode cross-packet data or data crossing multiple TCP segments. You can use the `file_data` rule keyword with its **MIME** argument to inspect the decoded data. See Pointing to a Specific Payload Type on page 857 for more information.

This option does not generate events.

Maximum MIME Decoding Depth

Specifies the maximum number of bytes of Base64-encoded data to decode per SMTP session. You can specify 5 to 20480 bytes.

This option does not generate events.

Configuring SMTP Decoding

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use the SMTP Configuration page of an intrusion policy to configure SMTP normalization.

Keysight, EX-1004
Page 585

If you are managing IPS using a Defense Center with RNA, note that you can process SMTP traffic using non-standard ports. To check each packet for SMTP service identifiers and process packets where they are found, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

To configure SMTP decoding options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

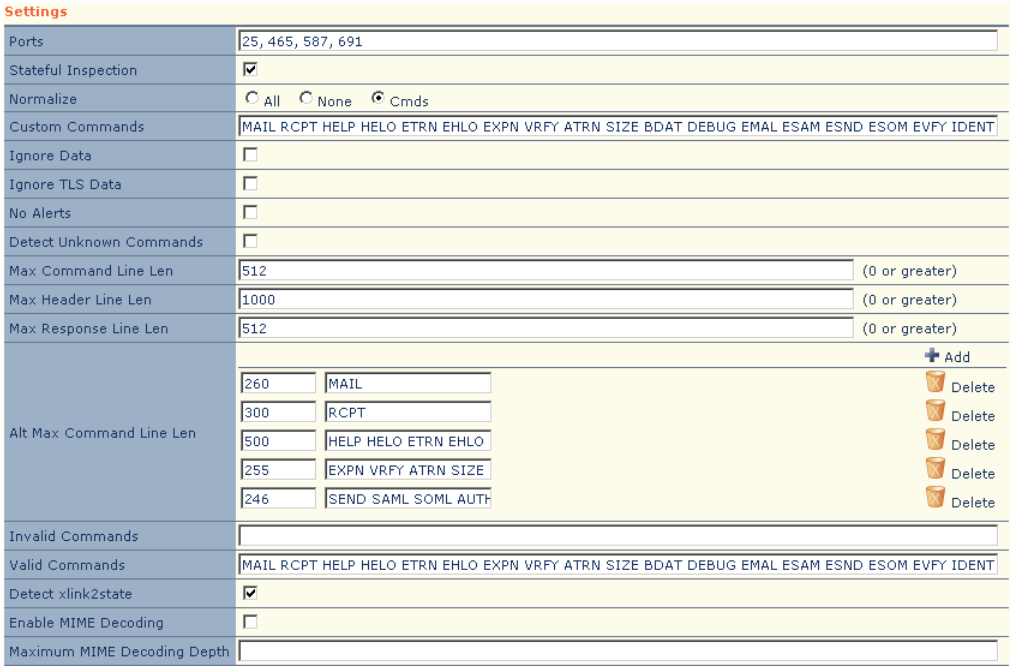
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 586

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.



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 494 for more information.

IMPORTANT! For more information on SMTP decoding options, see Understanding SMTP Decoding on page 582.

5. Specify the **Ports** where SMTP traffic should be decoded, separated by commas.
6. Select or clear **Stateful Inspection**:
   - To examine reassembled TCP streams containing SMTP packets, select **Stateful Inspection**.
   - To inspect only unreassembled SMTP packets, clear **Stateful Inspection**.

Keysight, EX-1004
Page 587

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, QUEU, STARTTLS, TICK, TIME, TURN, TURNME, VERB, VRFY, X-EXPS, X-LINK2STATE, XADR, XAUTH, XCIR, XEXCH50, XGEN, XLICENSE, XQUE, XSTA, XTRN, or XUSR.

Keysight, EX-1004
Page 588

14. To detect packets that are part of X-Link2State Microsoft Exchange buffer data overflow attacks, select **Detect xlink2state**.
15. To enable decoding of MIME email attachments for inspection using the `file_data` rule keyword with its `MIME` argument, select **Enable MIME Decoding**.

    You must also select the SMTP **Stateful Inspection** option to decode cross-packet data or data crossing multiple TCP segments.

    Optionally, you can also specify a maximum of 5 to 20480 bytes of Base64-encoded data to decode per SMTP session in **Maximum MIME Decoding Depth**.
16. Optionally, click **Configure Rules for SMTP Configuration** at the top of the page to display rules associated with individual options.

    You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

    To identify the rules associated with different configuration options, see Understanding SMTP Decoding on page 582.

    Click **Back** to return to the SMTP Configuration page.
17. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
    - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 589

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Detecting Exploits Using the SSH Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 application attempts to connect to a secure SSH service 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.

Keysight, EX-1004
Page 590

When enabled, the preprocessor detects exploits and anomalies according to the configured options, but generates no events in most cases unless you also enable the accompanying rules with generator ID (GID) 128 on the Rule page. 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 435 for more information.

See the following sections for more information:

- Selecting SSH Preprocessor Options on page 591
- Configuring the SSH Preprocessor on page 594

Selecting SSH Preprocessor Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

This section describes the options you can use to configure the SSH preprocessor.

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

Keysight, EX-1004
Page 591

Finally, the preprocessor will automatically detect any version string other than version 1 or version 2.

You must ensure that TCP stream preprocessing is enabled to use the SSH preprocessor. See Using TCP Stream Preprocessing on page 627 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 452.

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.

This option does not generate events.

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.

This option does not generate events.

Number of Encrypted Packets to Inspect

Specifies the number of encrypted packets to examine per session.

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.

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.

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.

This option does not generate events.

Keysight, EX-1004
Page 592

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.

If 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.

Increase the value for this option if the preprocessor generates false positives on the Challenge-Response Buffer Overflow or CRC-32 exploit.

This option does not generate events.

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.

This option does not generate events.

Detect Challenge-Response Buffer Overflow Attack

Enables or disables detecting the Challenge-Response Buffer Overflow exploit.

This option generates events when you enable rule 128:1.

Detect SSH1 CRC-32 Attack

Enables or disables detecting the CRC-32 exploit.

This option generates events when you enable rule 128:2.

Detect Server Overflow

Enables or disables detecting the SecureCRT SSH Client Buffer Overflow exploit.

This option generates events when you enable rule 128:3.

Detect Protocol Mismatch

Enables or disables detecting protocol mismatches.

This option generates events when you enable rule 128:4.

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).

Keysight, EX-1004
Page 593

This option generates events when you enable rule 128:5.

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.

This option generates events when you enable rule 128:6.

Detect Bad Version String

Note that, when enabled, the preprocessor detects without configuration any version string other than version 1 or 2.

This option generates events when you enable rule 128:7.

Configuring the SSH Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

This section explains how to configure the SSH preprocessor.

To configure the SSH preprocessor:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

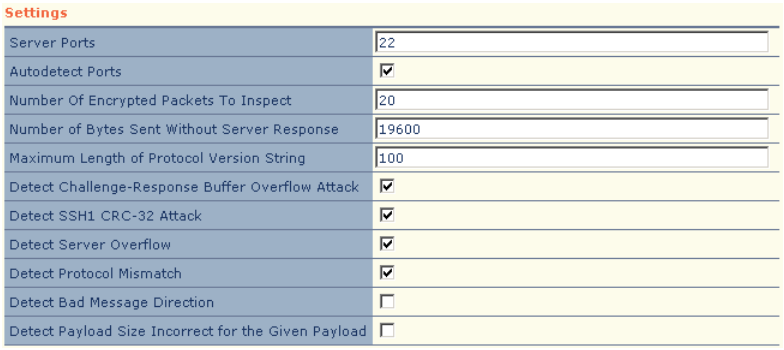
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 594

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 494 for more information.

5. You can modify any of the options on the SSH Configuration preprocessor page. See Selecting SSH Preprocessor Options on page 591 for more information.

6. Optionally, click **Configure Rules for SSH Configuration** at the top of the page to display rules associated with individual options.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Selecting SSH Preprocessor Options on page 591.

   Click **Back** to return to the SSH Configuration page.

Keysight, EX-1004
Page 595

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Using the SSL Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

Although IPS 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, IPS 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 IPS 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 that when an intrusion rule that requires this preprocessor is enabled in an intrusion policy, you must enable the preprocessor or choose to allow IPS to

Keysight, EX-1004
Page 596

enable it automatically before you can save the policy. For more information, see Automatically Enabling Advanced Settings on page 490.

For more information, see the following sections:

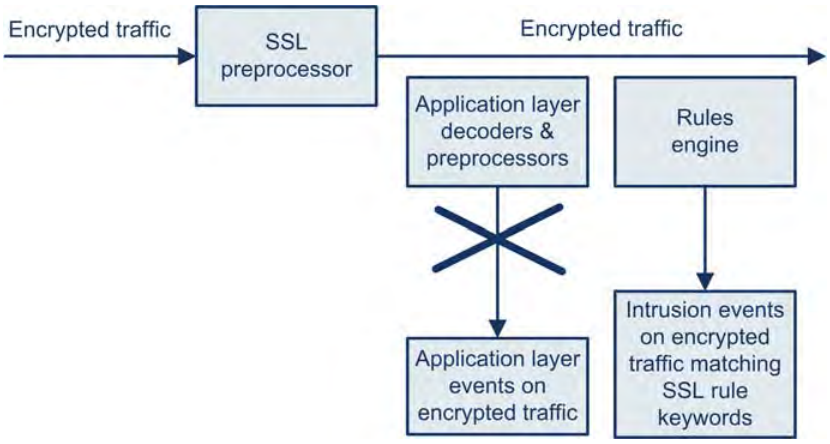
- Understanding SSL Preprocessing on page 597
- Enabling SSL Preprocessor Rules on page 598
- Configuring the SSL Preprocessor on page 599

Understanding SSL Preprocessing

Platform: DC/MDC/Standalone Sensor

Component: IPS

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. Once the preprocessor detects that a session state is encrypted, IPS marks the traffic in that session as encrypted. You can configure IPS to stop processing on all packets in an encrypted session once 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

Keysight, EX-1004
Page 597

for SSL preprocessing. For qualifying traffic, IPS uses the following scenarios to determine whether the traffic is encrypted:

- IPS 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
- IPS 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
- IPS 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
- IPS 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, IPS ignores future packets in a session once it marks the session as encrypted.

Normally the SSL preprocessor only processes traffic traveling over standard SSL ports. However, if you are managing IPS using a Defense Center with RNA, you can enable the adaptive profiles feature to use service information from your RNA network map to identify SSL traffic traveling over non-standard ports. When you apply an intrusion policy with adaptive profiles enabled, the preprocessor engine checks each packet for service identifiers to see if the packet is SSL traffic. If the identifier is found, the packet is processed even if it came over a non-standard SSL port. For more information on how adaptive profiles affect preprocessing of traffic, see Using Adaptive Profiles on page 697.

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 825. Note that you must enable the SSL preprocessor to allow processing of rules that contain SSL keywords.

Enabling SSL Preprocessor Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

When enabled, the SSL preprocessor inspects the contents of the handshake and key exchange messages exchanged at the beginning of an SSL session, but generates no events in most cases unless you also enable the accompanying rules with generator ID (GID) 137 on the intrusion policy Rules page. See Setting Rule States on page 435 for more information.

Keysight, EX-1004
Page 598

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 599 for more information.

Configuring the SSL Preprocessor

Platform: DC/MDC/Standalone Sensor

Component: IPS

By default, IPS attempts to inspect encrypted traffic. When you enable the SSL preprocessor, it detects when a session becomes encrypted. Once 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 IPS policy, you should also enable the SSL preprocessor in that policy.

In addition, you can enable the **Stop inspecting encrypted traffic** option to cause IPS 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. IPS 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 IPS 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 IPS does not continue inspecting traffic in the encrypted session.

You can specify the ports where the preprocessor monitors traffic for encrypted sessions.

Keysight, EX-1004
Page 599

If you are managing IPS using a Defense Center with RNA, note that you can process SSL traffic using non-standard ports. To check each packet for SSL service identifiers, enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

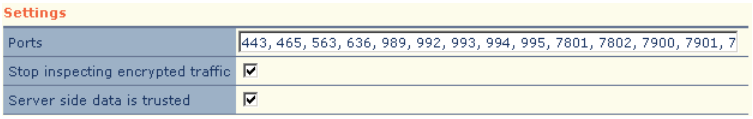
3. Click **Advanced Settings** in the navigation panel on the left.

   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 494 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.

Keysight, EX-1004
Page 600

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.

   A filtered view appears 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 435 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 to the preprocessor to log intrusion events.
9. Click **Back** to return to the SSH Configuration page.
10. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
    - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 601

Chapter 14
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 602
- Normalizing Inline Traffic on page 604
- Defragmenting IP Packets on page 613
- Understanding Packet Decoding on page 621
- Using TCP Stream Preprocessing on page 627
- Using UDP Stream Preprocessing on page 645

Verifying Checksums

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS 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

Keysight, EX-1004
Page 602

algorithm to verify the integrity of a protocol in the packet. The packet is considered to be unchanged if the detection engine computes the same value that is written in the packet by the end host.

Disabling checksum verification can leave your network susceptible to insertion attacks. Note that IPS 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

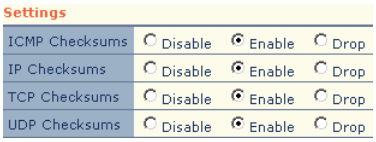
3. Click **Advanced Settings** in the navigation panel on the left.

   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 494 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**

Keysight, EX-1004
Page 603

- **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 you apply. See Setting Drop Behavior in an Inline Deployment on page 358 for more information. Note also that setting these options to **Drop** in a passive deployment is the same as setting them to **Enable**.

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Normalizing Inline Traffic

Platform: DC/MDC/Standalone Sensor

Component: IPS

The inline normalization preprocessor normalizes traffic to minimize the chances of attackers evading detection in inline deployments. When you apply an intrusion

Keysight, EX-1004
Page 604

policy with the inline normalization preprocessor enabled, IPS tests the following two conditions to ensure that you are using an inline deployment:

- **Drop when Inline** is enabled
- The policy is applied to an inline or inline with failopen interface 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 normalizations 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 IPS processes are the same as the packets received by the hosts on your network.

See the following sections for more information:

- Understanding Protocol Normalization on page 605
- Configuring Inline Normalization on page 607

Understanding Protocol Normalization

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

This section lists the base normalizations and any optional normalizations for each protocol type. See Configuring Inline Normalization on page 607 for information on configuring normalization of traffic for different protocols.

IPv4 Normalization

When you enable **Normalize IPv4**, IPS 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)

Keysight, EX-1004
Page 605

In addition, IPS performs the following optional normalizations when you enable IPV4 normalization and select the corresponding option:

- clears the single-bit Don't Fragment subfield of the IPv4 Flags header field; see the **Normalize Don't Fragment Bit** option in Configuring Inline Normalization on page 607 for more information
- clears the single-bit Reserved subfield of the IPv4 Flags header field; see the **Normalize Reserved Bit** option in Configuring Inline Normalization on page 607 for more information
- sets the Time to Live (TTL) field as needed to a specified minimum value; see the **Reset TTL** and **Minimize TTL** options in Configuring Inline Normalization on page 607 for more information

IPv6 Normalization

When you enable **Normalize IPv6**, IPS 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, IPS 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 607 for more information.

ICMPv4 and ICMPv6 Normalization

When you enable **Normalize ICMPv4**, **Normalize ICMPv6**, or both, IPS clears the 8-bit Code field in Echo (Request) and Echo Reply messages in the corresponding ICMP traffic.

TCP Normalizations

When you enable **Normalize TCP**, IPS performs the following base normalizations:

- removes data in synchronization (SYN) packets
- 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
- 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

Keysight, EX-1004
Page 606

IPS also performs the following normalizations unless you specify that IPS should allow any TCP option in TCP traffic:

- 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 607 for more information.

In addition, IPS performs the following optional normalizations when you enable TCP normalization and select the corresponding option:

- sets the two-byte Urgent Pointer header field to the payload length if the pointer is greater than the payload length; see the **Normalize Urgent Pointer** option in Configuring Inline Normalization on page 607 for more information
- normalizes the TCP Data field to ensure consistency in retransmitted data and drops any segments that cannot be properly reassembled; see the **Normalize TCP Payload** option in Configuring Inline Normalization on page 607 for more information
- optionally, clears ECN flags on a per-packet basis regardless of negotiation, or on a per-stream basis if usage was not negotiated; see the **Explicit Congestion Notification** option in Configuring Inline Normalization on page 607 for more information

Configuring Inline Normalization

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 607

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 358
- you must apply your policy to an inline or inline with failopen interface set; you can identify the interface type on the Apply Policy page where you apply your policy; see Applying an Intrusion Policy on page 347

You must also ensure that the TCP stream preprocessor is enabled when you enable TCP normalization; see Modifying Advanced Settings on page 478.

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 IPS 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 IPS 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**

IPS assumes a value of 1 when the field is empty.

Note that you can enable the following rules in the packet decoder rule category to generate events for this option:

- You can enable rule 116:428 to generate an event when IPS detects an IPv4 packet with a TTL less than the specified minimum.
- You can enable rule 116:270 to generate an event when IPS 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 625 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

IPS 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. IPS assumes a value of 0 when the field is empty.

Keysight, EX-1004
Page 608

Normalize IPv4

Enables normalization of IPv4 traffic. See IPv4 Normalization on page 605 for information on specific IPv4 normalizations. IPS 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 IPv6

Sets all Option Type fields in the Hop-by-Hop Options and Destination Options extension headers to 00 (Skip and continue processing). IPS 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 Normalizations on page 606 for information on specific TCP normalizations. When this option is enabled, you can also enable **Normalize Urgent Pointer**, **Normalize TCP 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 478.

Keysight, EX-1004
Page 609

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.

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 633 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.

IPS does not normalize options that you explicitly allow. IPS normalizes options that you do not explicitly allow by setting the options to No Operation (TCP Option 1).

IPS always allows the Maximum Segment Size (MSS), Window Scale, and Time Stamp TCP options because these options are commonly used for optimal TCP performance. IPS normalizes these commonly-used options as described in TCP Normalizations on page 606 regardless of the configuration of **Allow These TCP Options**. IPS 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.

Keysight, EX-1004
Page 610

You can also specify any, which allows all TCP options and effectively disables normalization of all TCP options. See TCP Normalizations on page 606 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, IPS 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

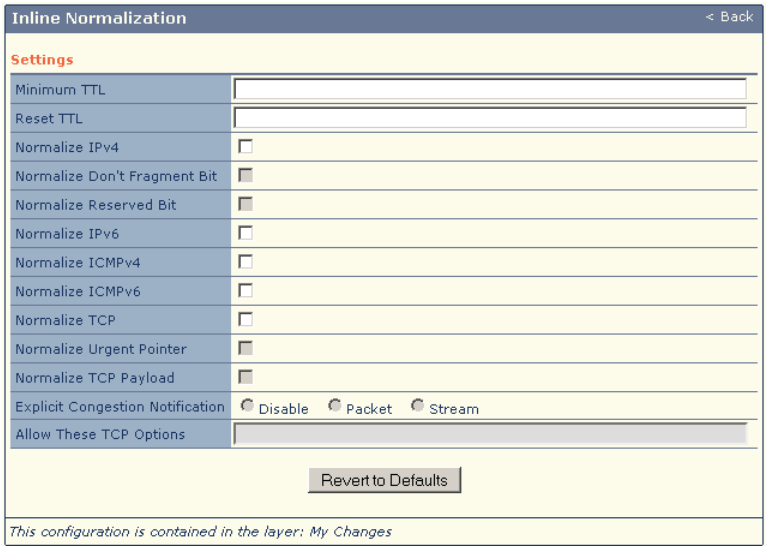
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 611

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 494 for more information.

5. You can set any of the options described in Configuring Inline Normalization on page 607.

Keysight, EX-1004
Page 612

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Defragmenting IP Packets

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

When enabled, the preprocessor defragments traffic according to the configured options, but generates no events in most cases before sending the decoded

Keysight, EX-1004
Page 613

traffic to the rules engine unless you also enable the accompanying rules with generator ID (GID) 123 on the intrusion policy Rules page. 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 435 for more information.

See the following sections for more information:

- Understanding IP Fragmentation Exploits on page 614
- Target-Based Defragmentation Policies on page 615
- Selecting Defragmentation Options on page 616
- Configuring IP Defragmentation on page 618

Understanding IP Fragmentation Exploits

Platform: DC/MDC/Standalone Sensor

Component: IPS

Enabling IP defragmentation helps you detect attacks against hosts on your network, like the teardrop attack, and resource consumption attacks against IPS 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 IPS self-preservation above exhaustive inspection. IPS 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 IPS 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

Keysight, EX-1004
Page 614

host on your network. See Target-Based Defragmentation Policies on page 615 for more information.

TIP! If you are managing IPS using a Defense Center with RNA and you want to dynamically select target-based policies using host operating system information for the target host in a packet, you can enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

Target-Based Defragmentation Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 IPS and exploit that host by sending malicious content hidden in overlapping packet fragments. This packet, when reassembled and inspected by IPS, 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

Keysight, EX-1004
Page 615

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
Linux	• Linux • OpenBSD
Last	• Cisco IOS
Solaris	• SunOS
Windows	• Windows

Selecting Defragmentation Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

You can configure the global **Preallocated Fragments** option:

Keysight, EX-1004
Page 616

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 3D Sensor, the memory limit for the 3D Sensor 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 CIDR 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 CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

Note that the `default` setting in the default policy specifies all IP addresses (0.0.0.0/0) 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 for the default policy, and you cannot specify 0.0.0.0/0 or leave this setting blank in another policy.

This option does not generate events.

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 615 for detailed information on these policies.

This option does not generate events.

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.

This option does not generate events.

Keysight, EX-1004
Page 617

Minimum TTL

Specifies the minimum acceptable TTL value a packet may have. This option detects TTL-based insertion attacks against IPS.

This option generates events when you enable rule 123:1.

Detect Anomalies

Identifies fragmentation problems such as overlapping fragments.

This option generates events when you enable the following rules:

- 123:1 through 123:4,
- 123:5 (BSD policy)
- 123:6 through 123:10

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.

This option generates events when you enable rule 123:12.

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.

This option generates events when you enable rule 123:13.

Configuring IP Defragmentation

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use the following procedure to configure the IP defragmentation preprocessor.

To configure IP defragmentation:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 618

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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**.

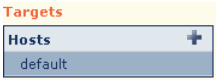
   The IP Defragmentation page appears.



   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 494 for more information.

5. Optionally, you can modify the setting for **Preallocated Fragments** in the Global Settings page area.

   See Selecting Defragmentation Options on page 616 for more information.



6. Optionally, click **Configure Rules for IP Defragmentation** at the top of the page to display rules associated with individual options.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Selecting Defragmentation Options on page 616.

   Click **Back** to return to the IP Defragmentation page.

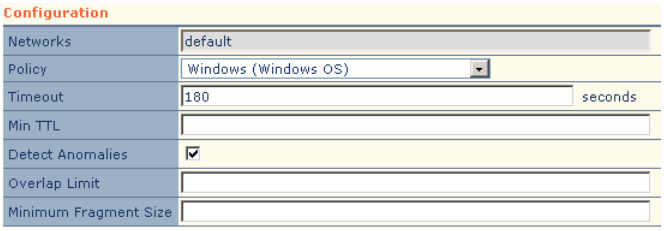
Keysight, EX-1004
Page 619

7. You have two options:
   - Add a new target-based policy. Click the **+** sign 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 CIDR block, or a comma-separated list comprised of either or both. You can create a total of 255 target-based policies including the default policy. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     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.



8. Optionally, you can modify any of the options in the Configuration page area. See Selecting Defragmentation Options on page 616 for more information.

Keysight, EX-1004
Page 620

9. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding Packet Decoding

Platform: DC/MDC/Standalone Sensor

Component: IPS

Before sending captured packets to a preprocessor, the detection engine first sends the packets to the packet decoder. The packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and the rules engine. 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 257.

When enabled, the packet decoder decodes traffic according to the configured options, but generates no events in most cases before sending the decoded traffic to the preprocessors unless you also enable the accompanying rules with generator ID (GID) 116 on the intrusion policy Rules page. A link on the

Keysight, EX-1004
Page 621

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 435 for more information.

Detect Teredo on Non-Standard Ports

Inspects Teredo tunneling of IPv6 traffic that is identified on a UDP port other than port 3544.

IPS 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. IPS 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 IPS 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** 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 417 and Setting Rule States on page 435 for more information.

This option does not generate events.

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.

This option generates events when you enable rules 116:4 and 116:5.

Keysight, EX-1004
Page 622

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
19	MD5 Signature Option
20	Space Communications Protocol Standards (SCPS)
21	Selective Negative Acknowledgements (SCPS)
22	Record Boundaries (SCPS)
23	Corruption (SPCS)
24	SNAP
26	TCP Compression Filter

Because these are experimental options, some systems do not account for them and can 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.

This option generates events when you enable rule 116:58.

Keysight, EX-1004
Page 623

Detect Obsolete TCP Options

Detects TCP headers with obsolete TCP options. Because these are obsolete options, some systems do not account for them and can be open to exploits. The following table describes these options.

TCP Option	Description
6	Echo
7	Echo Reply
16	Skeeter
17	Bubba
25	Unassigned

This option generates events when you enable rule 116:57.

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 can be open to exploits.

This option generates events when you enable rule 116:56.

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.

This option generates events when you enable rules 116:54, 116:55, and 116:59.

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.

This option generates events when you enable any packet decoder rule other than rules specifically associated with other packet decoder options.

Keysight, EX-1004
Page 624

Note the following rules associated with the inline normalization preprocessor **Minimum TTL** option:

- You can enable rule 116:428 to generate an event when IPS detects an IPv4 packet with a TTL less than the specified minimum.
- You can enable rule 116:270 to generate an event when IPS 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 607 for more information.

Configuring Packet Decoding

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure packet decoding on the Packet Decoding configuration page.

To configure packet decoding:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

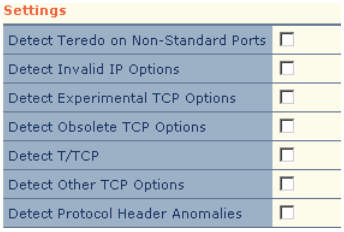
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 625

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 494 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 621 for more information.

6. Optionally, click **Configure Rules for Packet Decoding** at the top of the page to display rules associated with individual options.

   You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

   To identify the rules associated with different configuration options, see Understanding Packet Decoding on page 621.

   Click **Back** to return to the Packet Decoding page.

Keysight, EX-1004
Page 626

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Using TCP Stream Preprocessing

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

When enabled, the stream preprocessor reassembles traffic according to the configured options, but generates no events in most cases before sending the reassembled traffic to the rules engine unless you also enable the accompanying rules with generator ID (GID) 129 on the intrusion policy Rules page. 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

Keysight, EX-1004
Page 627

configure other rule actions. See Setting Rule States on page 435 for more information.

Note that when a rule that requires this preprocessor is enabled in an intrusion policy, you must enable the preprocessor or choose to allow IPS to enable it automatically before you can save the policy. For more information, see Automatically Enabling Advanced Settings on page 490.

If you enable any of the following, TCP stream preprocessing must be enabled:

- the DNS preprocessor
- the FTP/Telnet preprocessor
- the HTTP Inspect preprocessor
- the SMTP preprocessor
- the SSL preprocessor
- the DCE/RPC preprocessor when the RPC over HTTP proxy, RPC over HTTP server, TCP, or SMB transport protocol is selected
- portscan detection when the TCP protocol is selected
- TCP intrusion rules that use the `flow` or `flowbits` keyword

See the following sections for more information:

- Understanding State-Related TCP Exploits on page 628.
- Initiating Active Responses with Drop Rules on page 629.
- Selecting TCP Global Options on page 630.
- Understanding Target-Based TCP Policies on page 631.
- Selecting TCP Policy Options on page 633.
- Reassembling TCP Streams on page 637.
- Configuring TCP Stream Preprocessing on page 640.

Understanding State-Related TCP Exploits

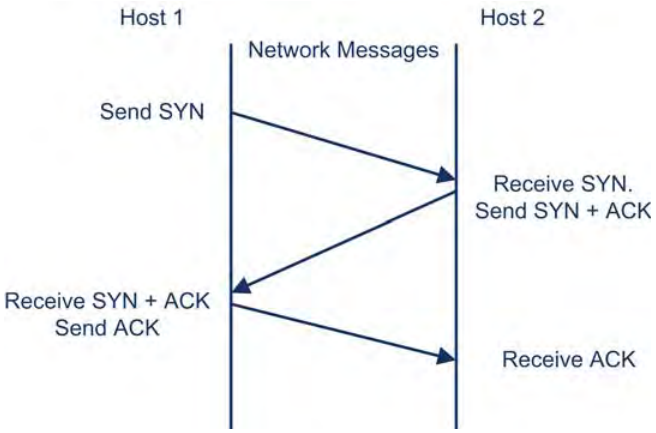
Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 628

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 IPS data.

Attacks like stick and snot use IPS'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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure the **Maximum Active Responses** option to initiate one or more *active responses* to close a TCP connection or UDP session when an offending packet triggers a TCP or UDP drop rule.

In an inline deployment, IPS responds to TCP drop rules by dropping the triggering packet and inserting a TCP Reset (RST) packet in both the client and server traffic. In a passive deployment, IPS cannot drop the packet or insert resets but sends a TCP reset to both the client and server ends of a TCP connection. In inline and passive deployments, IPS closes a UDP session by sending a UDP unreachable packet to each end of the session. Active responses are most

Keysight, EX-1004
Page 629

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, IPS can also initiate additional active responses if it sees additional traffic from either end of the connection or session. IPS 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 478 for more information.

See Selecting TCP Global Options on page 630 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 IPS 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 844 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 847 for more information.

Selecting TCP Global Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

This section describes the options that control how the TCP stream preprocessor functions.

Packet Type Performance Boost

Enables ignoring TCP traffic for all ports and services 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.

This option does not generate events.

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, IPS sends another active response unless the specified maximum has been reached. A setting of 0 disables active

Keysight, EX-1004
Page 630

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 629 and Initiating Active Responses with Rule Keywords on page 844.

This option does not generate events.

Minimum Response Seconds

Until **Maximum Active Responses** occur, specifies waiting 1 to 300 seconds before any additional traffic on a connection where IPS has initiated an active response results in a subsequent active response.

This option does not generate events.

Understanding Target-Based TCP Policies

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 IPS 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 IPS 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 CIDR 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

Keysight, EX-1004
Page 631

network that you do not identify in any other TCP policy, so there is no need to specify an IP address or CIDR block for the default TCP policy.

If you are managing IPS using a Defense Center with RNA, you can enable adaptive profiles to dynamically select target-based policies using host operating system information for the target host in a packet. For more information, see Using Adaptive Profiles on page 697.

The following list identifies the thirteen operating system policies and the host operating systems that use each.

First

- unknown OS

> **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.

Last

- Cisco IOS

BSD

- AIX
- FreeBSD
- OpenBSD

Linux

- Linux 2.4 kernel
- Linux 2.6 kernel

Old Linux

- Linux 2.2 and earlier kernel

Windows

- Windows 98
- Windows NT
- Windows 2000
- Windows XP

Keysight, EX-1004
Page 632

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The following list describes the options you can set to identify and control TCP traffic that the stream preprocessor inspects.

Network

Specifies the host or hosts to which you want to apply the TCP stream reassembly policy.

You can specify a single IP address or CIDR block. You can specify up to 255 total profiles including the default policy. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

Note that the default setting in the default policy specifies all IP addresses (0.0.0.0/0) 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 for the default policy, and you cannot specify 0.0.0.0/0 or leave this setting blank in another policy.

Keysight, EX-1004
Page 633

This option does not generate events.

Policy

Identifies the TCP policy operating system of the target host or hosts.

For more information, see Understanding Target-Based TCP Policies on page 631.

This option does not generate events.

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 3D Sensor is deployed on a segment where the network traffic is likely to reach the sensor's bandwidth limits, you should consider setting this value higher (for example, to 600 seconds) to lower the amount of processing overhead.

This option does not generate events.

Maximum TCP Window

Specifies the maximum TCP window size between 1 and 1073725440 bytes allowed by IPS 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.

This option generates events when you enable rule 129:6.

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.

This option generates events when you enable rule 129:7.

Keysight, EX-1004
Page 634

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 the be effective. See Normalizing Inline Traffic on page 604 for more information.

This option does not generate events.

Stateful Inspection Anomalies

Detects anomalous behavior in the TCP stack. When accompanying preprocessor rules are enabled, this can generate many events if TCP/IP stacks are poorly written.

This option generates events when you enable the following rules:

- 129:1 through 129:5
- 129:6 (Mac OS only)
- 129:8 through 129:11
- 129:13 through 129:19

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, IPS generates events.

This option generates events when you enable rules129:9 and 129:10.

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.

This option generates events when you enable rule 129:12.

Keysight, EX-1004
Page 635

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.

This option does not generate events.

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.

This option does not generate events.

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.

This option does not generate events.

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.

This option does not generate events.

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.

This option does not generate events.

Keysight, EX-1004
Page 636

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.

This option does not generate events.

Asynchronous Network

Specifies whether the monitored network is an asynchronous network, that is, a network where IPS sees only half the traffic. When this option is enabled, IPS does not reassemble TCP streams to increase performance.

This option does not generate events.

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 638.

This option does not generate events.

Perform Stream Reassembly on Client Services, Server Services, Both Services

If your deployment includes the RNA component, and adaptive profiles are enabled, specifies for client services, server services, or both, services to identify the traffic for the stream preprocessor to reassemble. See Selecting Stream Reassembly Options on page 638.

This option does not generate events.

Reassembling TCP Streams

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 637

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, DNS, HTTP, SMTP, or SSL port list should also be added to the TCP client stream reassembly port list for each TCP policy. For more information, see Understanding Server-Level FTP Options on page 546, Configuring the DCE/RPC Preprocessor on page 526, Detecting Exploits in DNS Name Server Responses on page 531, Selecting Server-Level HTTP Normalization Options on page 564, Decoding SMTP Traffic on page 582, Using the SSL Preprocessor on page 596, and Understanding Target-Based TCP Policies on page 631.

See the following sections for more information:

- Understanding Stream-Based Attacks on page 638
- Selecting Stream Reassembly Options on page 638

Understanding Stream-Based Attacks

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 since you are much less likely to receive malicious traffic from your own web server.

Selecting Stream Reassembly Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

In each TCP policy, you can specify a comma-separated list of ports to identify the traffic for the stream preprocessor to reassemble. If your deployment includes the RNA component, and adaptive profiles are enabled, you can also list RNA services that identify traffic to reassemble, either as an alternative to ports or in combination with ports. See Detected Operating Systems and Services on page 2019 for more information on the services RNA can detect on your network. See Using Adaptive Profiles on page 697 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

Keysight, EX-1004
Page 638

separate lists of services for any combination of client services, server services, and both. For example, assume that you wanted to reassemble the following:

- SMTP (port 25) traffic from the client
- FTP server responses (port 21)
- telnet (port 23) traffic in both directions

You could configure the following:

- For client ports, specify `23, 25`
- For server ports, specify `21, 23`

Or, instead, you could configure the following:

- For client ports, specify `25`
- For server ports, specify `21`
- For both ports, specify `23`

Additionally, consider the following example which combines ports and services and would be valid when adaptive profiles are enabled:

- For client ports, specify `23`
- For client services, specify `smtp`
- For server ports, specify `21`
- For server services, specify `telnet`

Although you can also specify `all` as the argument to provide reassembly for all ports, Sourcefire does **not** recommend setting ports to `all` because it can increase the amount of traffic inspected by this preprocessor and slow performance unnecessarily.

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.

This option does not generate events.

Perform Stream Reassembly on Client Services

If your deployment includes the RNA component and adaptive profiles are enabled, 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. See Detected Operating Systems and Services on page 2019 and Using Adaptive Profiles on page 697 for more information.

This option does not generate events.

Keysight, EX-1004
Page 639

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.

This option does not generate events.

Perform Stream Reassembly on Server Services

If your deployment includes the RNA component and adaptive profiles are enabled, 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. See Detected Operating Systems and Services on page 2019 and Using Adaptive Profiles on page 697 for more information.

This option does not generate events.

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.

This option does not generate events.

Perform Stream Reassembly on Both Services

If your deployment includes the RNA component and adaptive profiles are enabled, 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. See Detected Operating Systems and Services on page 2019 and Using Adaptive Profiles on page 697 for more information.

This option does not generate events.

Configuring TCP Stream Preprocessing

You can configure TCP stream preprocessing, including TCP policies.

To configure the stream preprocessor to track TCP sessions:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 640

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

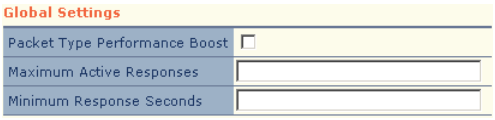
   The Advanced Settings page appears.

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**.

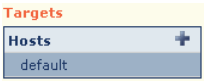
   IMPORTANT! You cannot disable TCP stream preprocessing when the DNS, FTP/Telnet, HTTP Inspection, SMTP, or SSL preprocessor is enabled, or when the DCE/RPC preprocessor is enabled with the RPC over HTTP proxy, RPC over HTTP server, TCP, or SMB transport protocol selected, or when portscan detection is 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 since 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 494 for more information.

5. Optionally, modify any of the options under Global Settings. See Selecting TCP Global Options on page 630 for more information.



Keysight, EX-1004
Page 641

6. You have two options:

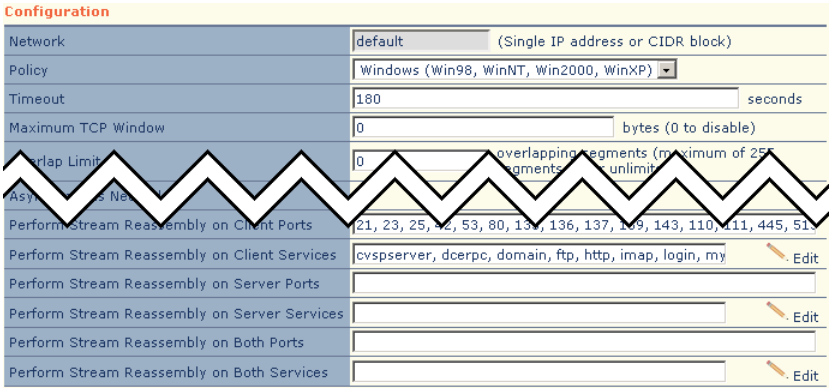
   - Add a new target-based policy. Click the **+** sign next to **Hosts** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses in the **Host Address** field and click **OK**.

     You can specify a single IP address or CIDR block. You can create a total of 255 target-based policies including the default policy. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

     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.



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.

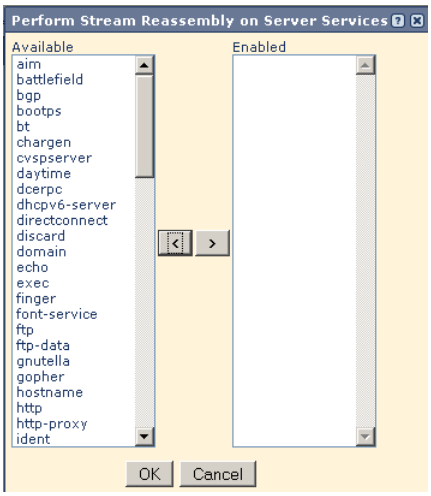
   For more information, see Selecting TCP Policy Options on page 633, and Selecting Stream Reassembly Options on page 638.

Keysight, EX-1004
Page 642

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 Client Services pop-up window.



   Note that your deployment must include the RNA component and you must enable adaptive profiles to monitor traffic for the stream preprocessor to reassemble based on RNA services. See Working with Services on page 1040 and Using Adaptive Profiles on page 697 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. You have two choices:

    - Click **OK** to add the selections.

      The TCP Stream Configuration page is displayed and the services are updated.
    - Click **Cancel** to abandon the selections.

      The TCP Stream Configuration page is displayed and services are not updated.

Keysight, EX-1004
Page 643

11. Optionally, click **Configure Rules for TCP Stream Configuration** at the top of the page to display rules associated with individual TCP policy options.

    You must set the rule state of preprocessor rules to Generate Events or, optionally, to Drop and Generate events in an inline deployment, if you want the preprocessor to log intrusion events for options associated with rules. See Setting Rule States on page 435 for more information.

    To identify the rules associated with different configuration options, see Selecting TCP Policy Options on page 633.

    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 492.

13. You have the following options:

    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

    - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.

    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 644

Using UDP Stream Preprocessing

Platform: DC/MDC/Standalone Sensor

Component: IPS

UDP stream preprocessing occurs when the rules engine processes packets against a UDP rule that includes the `flow` keyword 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 IPS 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 621.

Note 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 490.

The following preprocessors require UDP stream preprocessing to be enabled:

- DNS preprocessor
- DCE/RPC preprocessor with the UDP transport protocol selected

Configuring UDP Stream Preprocessing

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure UDP stream preprocessing.

Keysight, EX-1004
Page 645

To configure the stream preprocessor to track UDP sessions:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.

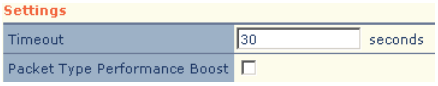
   The Advanced Settings page appears.

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 since these rules will not trigger unless UDP stream preprocessing is enabled.

   The UDP 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 494 for more information.

5. Optionally, configure a **Timeout** value to specify the number of seconds between 1 and 86400 the preprocessor keeps an inactive stream in the state table. If additional datagrams are not seen in the specified time, the preprocessor deletes the stream from the state table.

6. Optionally, select **Packet Type Performance Boost** to ignore UDP traffic for all ports and services 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.

Keysight, EX-1004
Page 646

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 647

Chapter 15
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 648 explains detection of back orifice attacks.
- Detecting Portscans on page 651 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 662 explains how to limit denial of service (DoS) and SYN flood attacks.
- Detecting Sensitive Data on page 677 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS 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.

Keysight, EX-1004
Page 648

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:SIDs 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 435 for more information.

Back Orifice GID:SIDs

Preprocessor rule GID:SID	Description
105:1	Back Orifice traffic detected
105:2	Back Orifice client traffic detected
105:3	Back Orifice server traffic detected
105:4	Back Orifice snort buffer attack detected

To view the Back Orifice Detection page:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

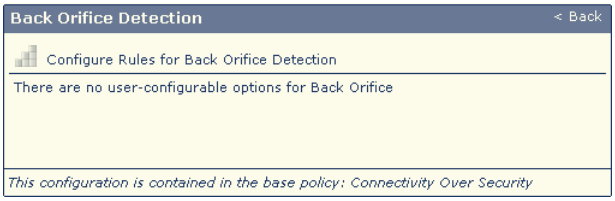
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 649

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 494 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 435 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 to the preprocessor to log intrusion events.

   Click **Back** to return to the Back Orifice Detection page.

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 650

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Detecting Portscans

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 services 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 435 and the Portscan Detection SIDs (GID:122) table on page 659 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

Keysight, EX-1004
Page 651

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 flow 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.

Keysight, EX-1004
Page 652

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	A one-to-one portscan in which an attacker uses one or a few hosts to scan multiple ports on a single target host. One-to-one portscans are characterized by: • a low number of scanning hosts • a single host that is scanned • a high number of ports scanned This option detects TCP, UDP, and IP portscans.
Port Sweep	A one-to-many portsweep in which an attacker uses one or a few hosts to scan a single port on multiple target hosts. Portsweeps are characterized by: • a low number of scanning hosts • a high number of scanned hosts • a low number of unique ports scanned This option detects TCP, UDP, ICMP, and IP portsweeps.
Decoy Portscan	A one-to-one portscan in which the attacker mixes spoofed source IP addresses with the actual scanning IP address. Decoy portscans are characterized by: • a high number of scanning hosts • a low number of ports that are scanned only once • a single (or a low number of) scanned hosts The decoy portscan option detects TCP, UDP, and IP protocol portscans.
Distributed Portscan	A many-to-one portscan in which multiple hosts query a single host for open services. Distributed portscans are characterized by: • a high number of scanning hosts • a high number of ports that are scanned only once • a single (or a low number of) scanned hosts The distributed portscan option detects TCP, UDP, and IP protocol portscans.

Keysight, EX-1004
Page 653

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	Detects only negative responses from targeted hosts. Select this sensitivity level to suppress false positives, but keep in mind that some types of portscans (slow scans, filtered scans) might be missed. This level uses the shortest time window for portscan detection.
Medium	Detects portscans based on the number of connections to a host, which means that you can detect filtered portscans. However, very active hosts such as network address translators and proxies may generate false positives. Note that you can add the IP addresses of these active hosts to the Ignore Scanned field to mitigate this type of false positive. This level uses a longer time window for portscan detection.
High	Detects portscans based on a time window, which means that you can detect time-based portscans. However, if you use this option, you should be careful to tune the detector over time by specifying IP addresses in the Ignore Scanned and Ignore Scanner fields. This level uses a much longer time window for portscan detection.

See the following sections for more information:

- Configuring Portscan Detection on page 655
- Understanding Portscan Events on page 658

Keysight, EX-1004
Page 654

Configuring Portscan Detection

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 435 and the Portscan Detection SIDs (GID:122) table on page 659 for more information.

To configure portscan detection:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

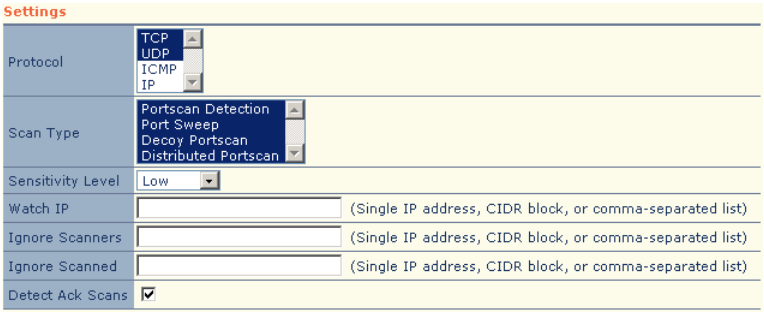
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 655

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.



   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 494 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 652 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 de-select multiple protocols. See the Portscan Types table on page 653 for more information.

7. In the **Sensitivity Level** list, specify which level you want to use: low, medium, or high.

   See the Sensitivity Levels table on page 654 for more information.

Keysight, EX-1004
Page 656

8. Optionally, in the **Watch IP** field, specify which hosts you want to watch for signs of portscan activity, or leave the field blank to watch all network traffic.

   You can enter a single IP address or CIDR block, or a comma-separated list of either or both. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

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 enter a single IP address or CIDR block, or a comma-separated list of either or both. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

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 enter a single IP address or CIDR block, or a comma-separated list of either or both. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

11. Optionally, disable **Detect Ack Scans** to discontinue monitoring of sessions picked up in mid-stream.

    IMPORTANT! Detection of mid-stream sessions helps to identify ACK scans, but can 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 or, optionally, to Drop and Generate Events in an inline deployment; click **Configure Rules for Portscan Detection** at the top of the page to display rules associated with individual TCP policy options.

    See Setting Rule States on page 435 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 659.

    Click **Back** to return to the Portscan Detection page.

13. You have the following options:

    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

Keysight, EX-1004
Page 657

- You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

  The Intrusion Policy page appears.

  See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

  You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding Portscan Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 435 for more information. The **Preprocessor Rule SID** column in the Portscan Detection

Keysight, EX-1004
Page 658

SIDs (GID:122) table lists the SID for the preprocessor rule you must enable for each portscan type.

Portscan Detection SIDs (GID:122)

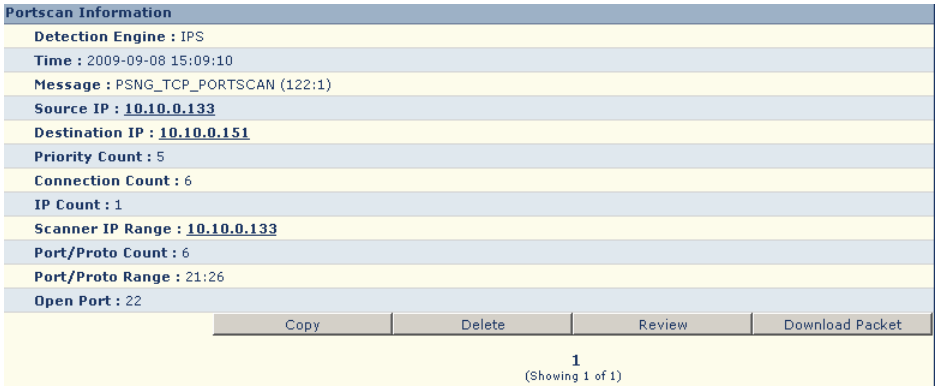
Portscan Type	Protocol:	Priority	Preprocessor Rule SID
Portscan Detection	TCP	Low	1
		Medium or High	5
	UDP	Low	17
		Medium or High	21
	ICMP:	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	9
		Medium or High	13
Port Sweep	TCP	Low	3, 27
		Medium or High	7
	UDP	Low	19
		Medium or High	23
	ICMP	Low	25
		Medium or High	26
	IP	Low	11
		Medium or High	15
Decoy Portscan	TCP	Low	2
		Medium or High	6
	UDP	Low	18
		Medium or High	22
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	10
		Medium or High	14
Distributed Portscan	TCP	Low	4
		Medium or High	8
	UDP	Low	20
		Medium or High	24
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	12
		Medium or High	16

When you enable the accompanying preprocessor rules, the portscan detector generates intrusion events that you can view just as you would any other intrusion event. However, the information presented on the packet view is different from the other types of intrusion events. This section describes the fields that appear on the packet view for a portscan event and how you can use that information to understand the types of probes that occur on your network.

Keysight, EX-1004
Page 659

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 267.

A packet view of a TCP portscan appears. For example, the following graphic illustrates the packet view for a TCP portscan event:



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.

IMPORTANT! For events generated by the portscan flow 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.

Portscan Packet View

Information	Description
Detection Engine	The detection engine that detected the event.
Time	The time when the event occurred.
Message	The event message generated by the preprocessor.

Keysight, EX-1004
Page 660

Portscan Packet View (Continued)

Information	Description
Source IP	The IP address of the scanning host. 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.
Destination IP	The IP address of the scanned host. 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.
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. Click an 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.
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.

Keysight, EX-1004
Page 661

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 662
- Rate-Based Attack Prevention and Other Filters on page 666
- Configuring Rate-Based Attack Prevention on page 673
- Understanding Dynamic Rule States on page 453
- Setting a Dynamic Rule State on page 454

Understanding Rate-Based Attack Prevention

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 3D Sensors deployed in inline mode to block rate-based attacks for a specified time and then revert to only generating events and not drop traffic.

Keysight, EX-1004
Page 662

Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate service 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 665.
- 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 665.
- 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 454.
- excessive matches for a particular rule across all traffic detected by the detection engine.

  To configure rule-based dynamic rule states, see Setting a Dynamic Rule State on page 454.

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.

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, IPS 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 once 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

Keysight, EX-1004
Page 663

set to Generate Events create events, but IPS does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action can 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.

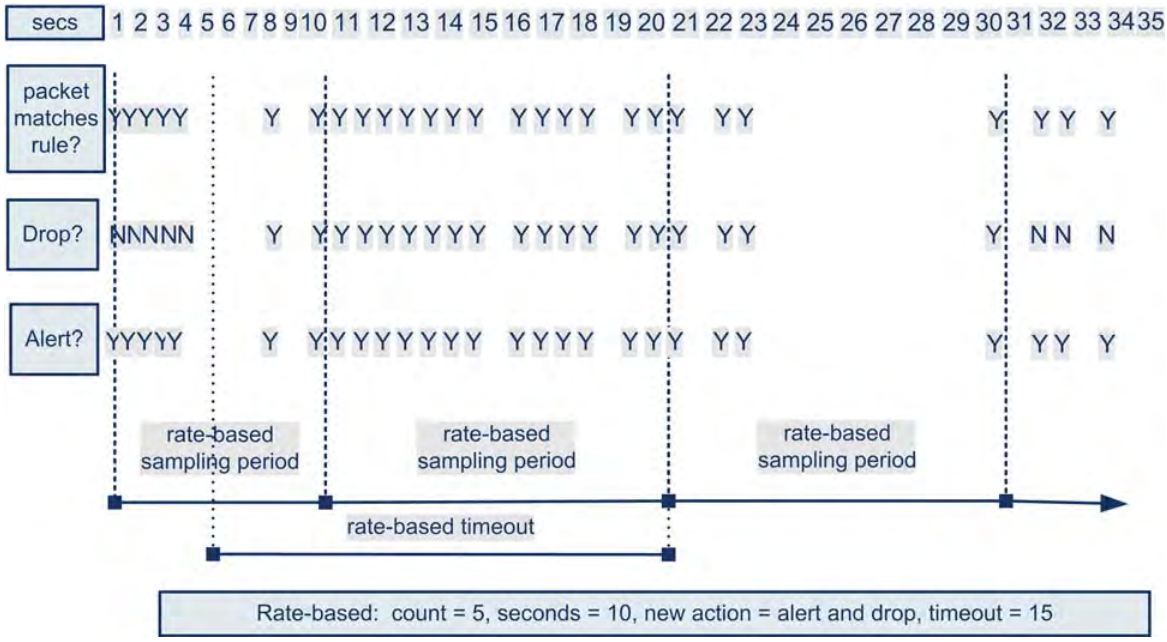
You can define multiple rate-based filters on the same rule. The first filter listed in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, IPS implements the action of the first rate-based filter. Similarly, policy-wide rate-based filters override rate-based 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

Keysight, EX-1004
Page 664

generating events only after a sampling period completes where the sampled rate is below the threshold rate.



Preventing SYN Attacks

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 sensor is deployed passively, you can generate events. If your sensor 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.

Controlling Simultaneous Connections

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 IPS detects the configured number of successful connections to or from a specified IP

Keysight, EX-1004
Page 665

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.

Rate-Based Attack Prevention and Other Filters

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `detection_filter` keyword and the thresholding and suppression features provide other ways to filter either the traffic itself or the events that IPS 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 666
- Dynamic Rule States and Thresholding or Suppression on page 668
- Policy-Wide Rate-Based Detection and Thresholding or Suppression on page 669
- Rate-Based Detection with Multiple Filtering Methods on page 671

Rate-Based Attack Prevention and Detection Filtering

Platform: DC/MDC/Standalone Sensor

Component: IPS

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, IPS tracks the number of incoming packets matching the pattern in the rule per timeout period. IPS 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.

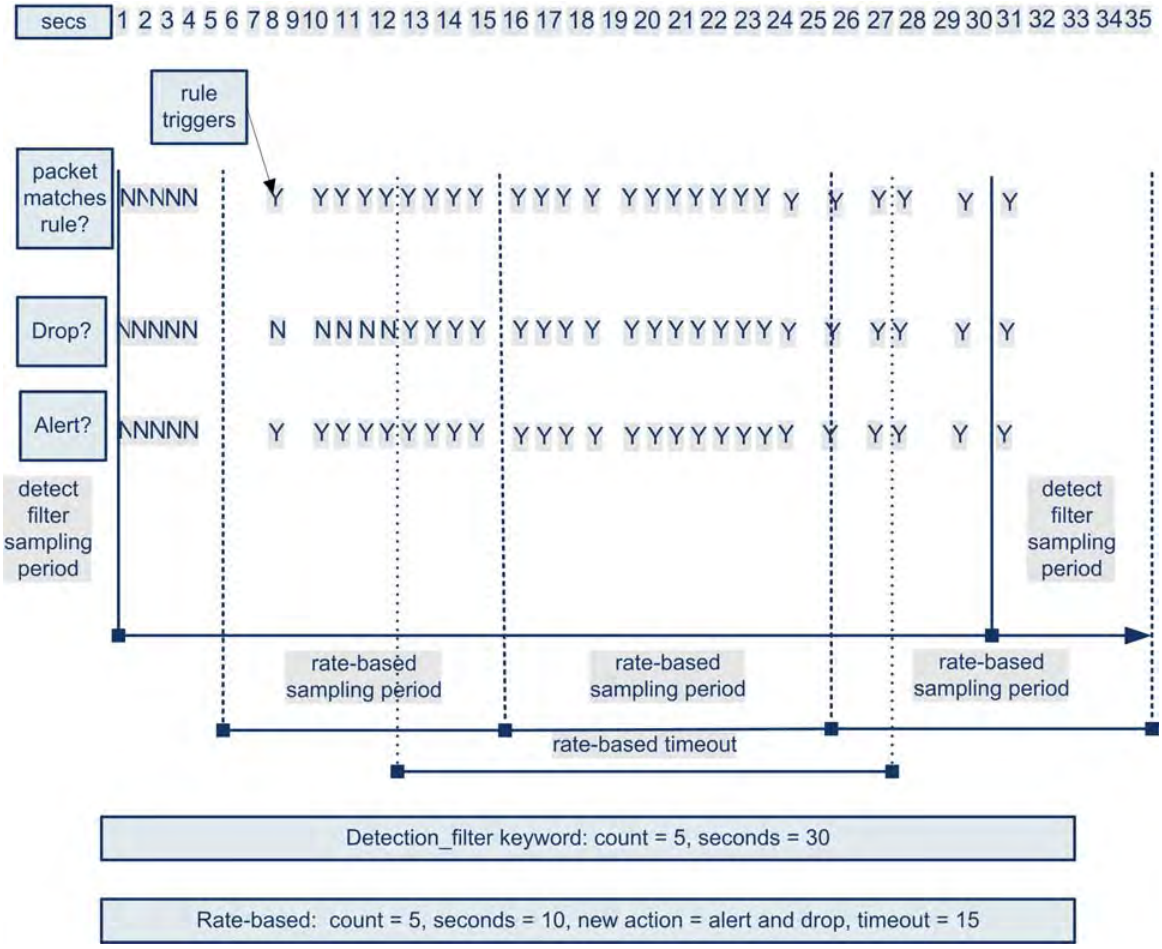
The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that also includes the `detection_filter` keyword, with a count set to 5. This rule has rate-based attack prevention configured. The rate-based settings change the rule attribute 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

Keysight, EX-1004
Page 666

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

Keysight, EX-1004
Page 667

want to manage the rate and timeout settings in the intrusion policy. For more information, see Setting Rule States on page 435.

Dynamic Rule States and Thresholding or Suppression

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 439 and Configuring Suppression Per Intrusion Policy on page 447.

If you apply suppression to a rule, IPS 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.

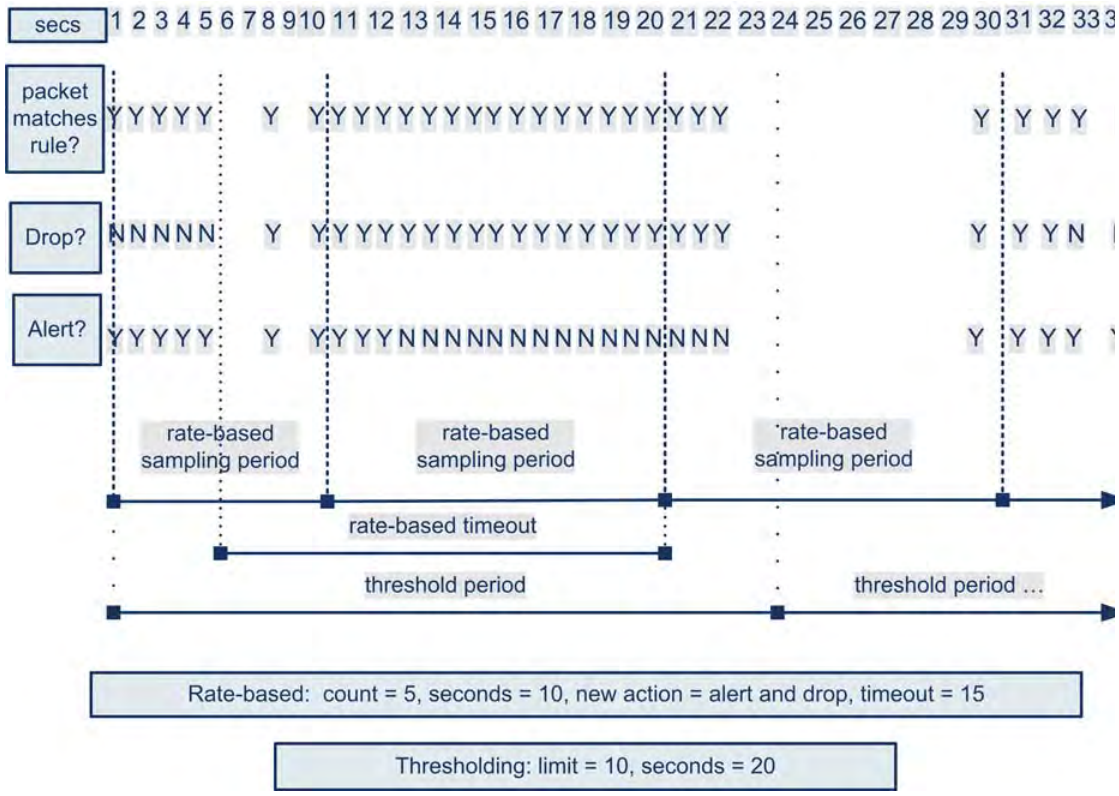
The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 15 seconds when there are five hits on the rule in 10 seconds. In addition, a limit threshold limits the number of events the rule can generate to 10 events in 20 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 IPS drops the packet. After the tenth packet, the limit threshold has been reached, so for the remaining packets IPS 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

Keysight, EX-1004
Page 668

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, IPS generates a single event to indicate the change in action. So, for example, when the limit threshold of 10 is reached and IPS stops generating events and the action changes from Generate Events to Drop and Generate Events on the 14th packet, IPS generates an eleventh event to indicate the change in action.

Policy-Wide Rate-Based Detection and Thresholding or Suppression

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 708, Configuring Event Thresholding on page 439, and Configuring Suppression Per Intrusion Policy on page 447.

Keysight, EX-1004
Page 669

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.

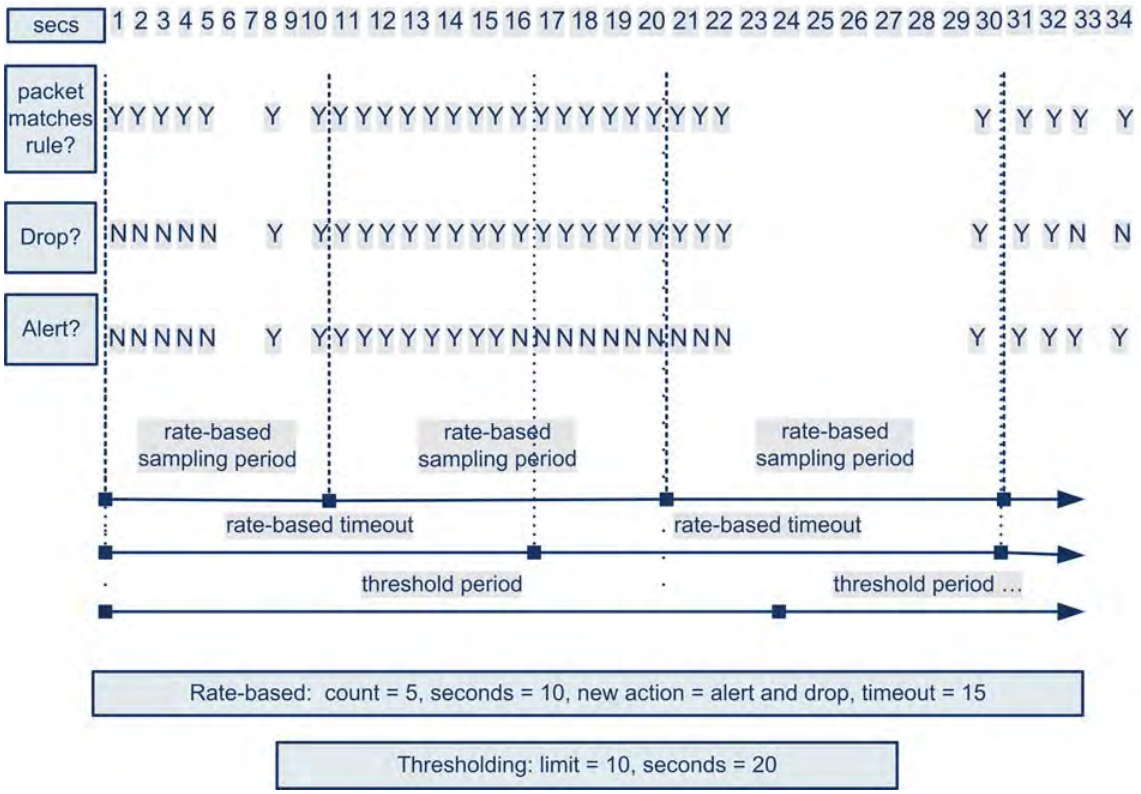
The following example shows an attacker attempting denial of service (DoS) attacks on hosts in your network. Many simultaneous connections to hosts from the same sources trigger a policy-wide Control Simultaneous Connections setting. The setting generates events and drops malicious traffic when there 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

Keysight, EX-1004
Page 670

dropping traffic continues. The rate-based action stops only after a sampling period completes where the sampled rate is below the threshold rate.



Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, IPS generates a single event to indicate the change in action. So, for example, if the limit threshold of 10 has been reached and IPS stops generating events and the action changes to Drop and Generate events on the 14th packet, IPS generates an eleventh event to indicate the change in action.

Rate-Based Detection with Multiple Filtering Methods

Platform: DC/MDC/Standalone Sensor

Component: IPS

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, therefore, describes a case where a `detection_filter` keyword, rate-based filtering, and thresholding interact.

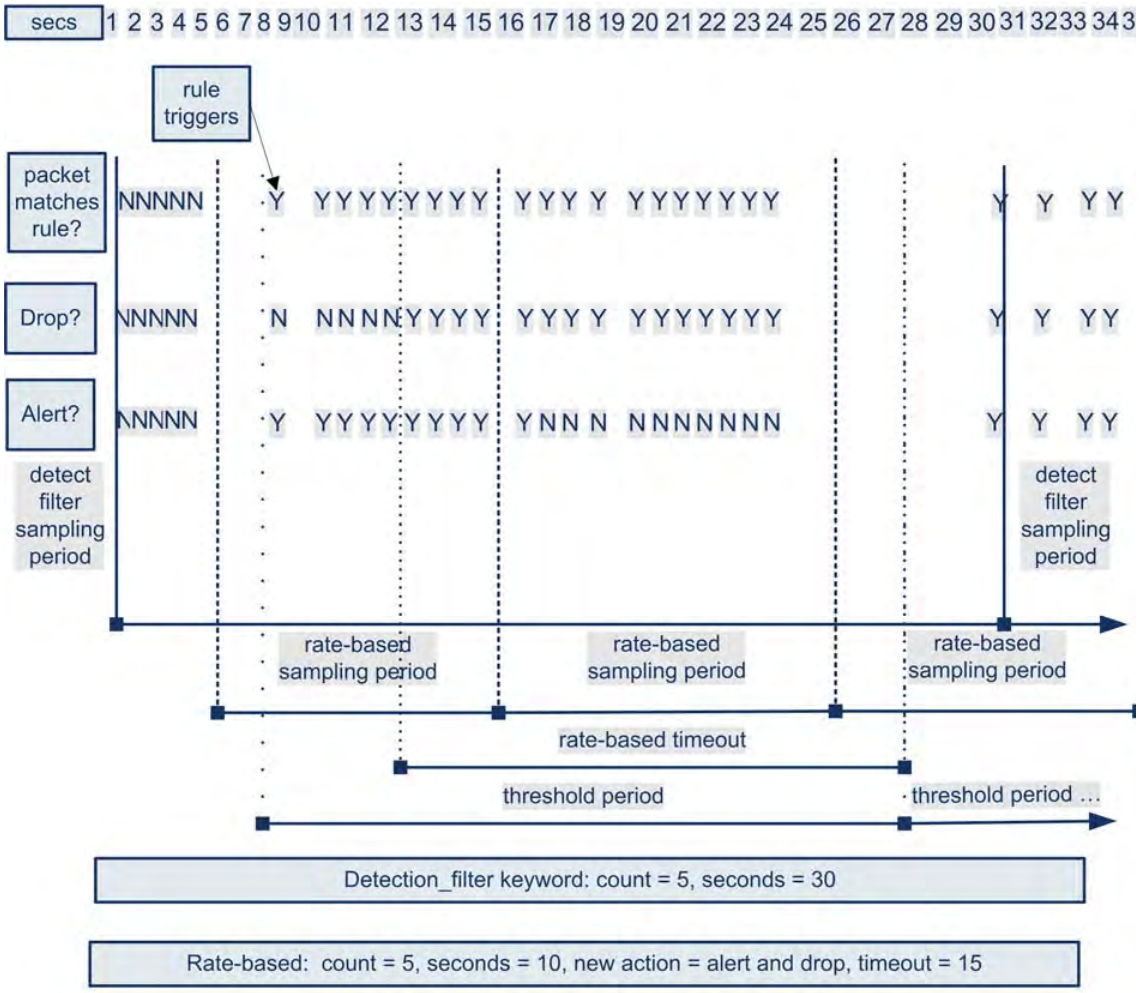
Keysight, EX-1004
Page 671

This example shows an attacker attempting a brute force login. 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, IPS 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 IPS does not generate events but does drop the packets.

Keysight, EX-1004
Page 672

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 673

To configure rate-based attack prevention:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

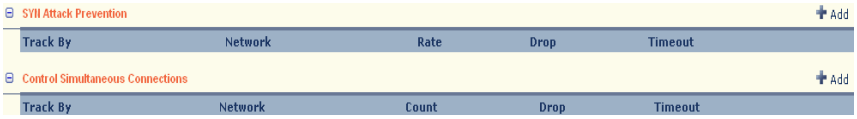
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



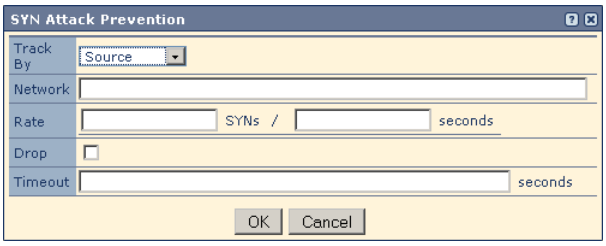
   A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in an Intrusion Policy on page 494 for more information.

Keysight, EX-1004
Page 674

5. You have two options:

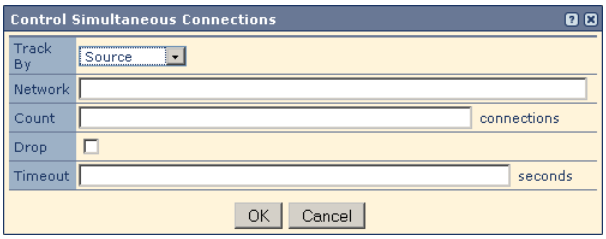
   - To prevent incomplete connections intended to flood a host, click **Add** under SYN Attack Prevention.

     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 an IP address or CIDR 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 CIDR block in the **Network** field.

   Note that IPS tracks traffic separately for each IP address included in the Network field. As an example, you might set a source CIDR block of 10.1.0.0/16 for the network setting and configure IPS 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, IPS 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 connections from 10.1.4.21 cause IPS to generate events.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

Keysight, EX-1004
Page 675

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 can entirely block connection to a host in an inline deployment.

10. You have the following options:
    - You can continue editing your policy.

      See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
    - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

      The Intrusion Policy page appears.

      See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

      You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
    - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
    - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

      See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 676

Detecting Sensitive Data

Platform: DC/MDC/Standalone Sensor

Component: IPS

Sensitive data such as Social Security numbers, credit card numbers, driver's license numbers, and so on can be leaked onto the Internet, intentionally or accidentally. IPS 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.

IPS 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, IPS would detect the phone number (555)123-4567, but not an obfuscated version where each number is separated by spaces, as in (5 5 5) 1 2 3 - 4 5 6 7, or by intervening HTML code, such as (555)-123-4567. However, IPS 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.

IPS 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 490 for more information.

Because IPS 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 627 for more information.

Keysight, EX-1004
Page 677

See the following sections for more information:

- Deploying Sensitive Data Detection on page 678
- Selecting Global Sensitive Data Detection Options on page 678
- Selecting Individual Data Type Options on page 680
- Using Predefined Data Types on page 681
- Configuring Sensitive Data Detection on page 683
- Selecting RNA Services to Monitor on page 686
- Special Case: Detecting Sensitive Data in FTP Traffic on page 689
- Using Custom Data Types on page 690

Deploying Sensitive Data Detection

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 and applying your intrusion policy:

- Select the No Rules Active default policy as your base policy; see Selecting the Base Policy on page 364 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 478 for more information.
- Apply the intrusion policy containing your sensitive data configuration to a separate detection engine reserved for sensitive data detection; see Applying an Intrusion Policy on page 347 for more information.

Selecting Global Sensitive Data Detection Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 678

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 290 for more information.
Networks	Specifies the destination host or hosts to monitor for sensitive data. You can specify a single IP address or CIDR block, or a comma-separated list of either or both. IPS interprets a blank field as any, meaning any destination IP address. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
Global Threshold	Specifies the total number of all occurrences of all data types during a single session that the preprocessor must detect in any combination before generating a global threshold event. You can specify 1 through 655355. Note that You must enable the preprocessor rule with the GID:SID combination 139:1 to detect and generate events on combined data type occurrences. See Setting Rule States on page 435 for information on enabling rules in your intrusion policy. Note also that the preprocessor generates up to one global threshold event per session. Note also that 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. 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 680 for more information.

Keysight, EX-1004
Page 679

Selecting Individual Data Type Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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
- If you are managing IPS using a Defense Center with RNA, RNA services to monitor for each data type

At a minimum, each data type must specify an event threshold and at least one port or RNA service 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 682 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 690 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.
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 RNA service for the data type.

Keysight, EX-1004
Page 680

Individual Data Type Options (Continued)

Option	Description
Threshold	Specifies the number of occurrences of the data type when IPS generates an event. You receive an error message when you save the policy if you do not set a threshold for an enabled data type. You can specify 1 through 255. Note that the preprocessor generates one event for a detected data type per session. Note also that global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the data type event threshold is reached, regardless of whether the global event threshold has been reached, and vice versa.
RNA Services	For a Defense Center with RNA, specifies up to eight RNA services 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 RNA service for the data type. You must configure RNA and enable Adaptive Profiles for this option to return a match. See Working with Services on page 1040 and Using Adaptive Profiles on page 697 for more information. See Selecting RNA Services to Monitor on page 686 for detailed instructions for selecting RNA services.
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 690 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 date 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 435 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

Keysight, EX-1004
Page 681

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 417 for more information. Predefined sensitive data rules are also listed on the Rule Editor page (**Policy & Response > IPS > 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 `(\d{3}) ?\d{3}-\d{4}`.	138:6
U.S. Social Security Numbers Without Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area number, 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 690 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.

Keysight, EX-1004
Page 682

Configuring Sensitive Data Detection

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can modify default global settings and settings for individual data types. You must also enable detection and event generation for the individual data types 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 490 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 679 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 680 for information on the options you can modify.
add or remove RNA services to monitor for a data type	click inside the **RNA Services** field, or click **Edit** next to the field. The RNA Services pop-up window appears. • To add up to eight 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. See Selecting RNA Services to Monitor on page 686 for more information. **IMPORTANT!** To detect sensitive data in FTP traffic, you must add the `ftp-data` service and enable the FTP/Telnet preprocessor. See Special Case: Detecting Sensitive Data in FTP Traffic on page 689 for more information.

Keysight, EX-1004
Page 683

Sensitive Data Configuration Actions (Continued)

To...	You can...
create a custom data type	click the **+** sign next to **Data Types** on the left side of the page. The Add Data Type pop-up window appears. Specify a unique data type name and the pattern you want to detect with this data type and click **OK**, or click **Cancel** to abandon your edits. See Using Custom Data Types on page 690 for more information.
display sensitive data preprocessor rules	click the **Configure Rules for Sensitive Data Detection** link above the Global Settings page area. A listing of all sensitive data preprocessor rules appears in a filtered display of the Rules page. Optionally, you can enable or disable any of the listed rules. Note that you must enable the sensitive data preprocessor rule for each data type that you want to use in your intrusion policy. See Setting Rule States on page 435 for more information. You can also configure sensitive data rules for any of the other actions available on the Rules page, such as rule suppression, rate-based attack prevention, and so on; see Managing Rules in an Intrusion Policy on page 403 for more information. Click **Back** to return to the Sensitive Data Detection page.
revert changes	click **Revert to Defaults**, then click **OK** at the prompt. All saved and unsaved changes to the policy-specific options in your configuration revert to the default settings in the base policy. See Understanding the Base Policy on page 361 for more information. Note the only configurable system-wide options are the custom data type name and data pattern, which do not revert. See Configuring Custom Data Types on page 693 for more information.

To configure sensitive data detection:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 684

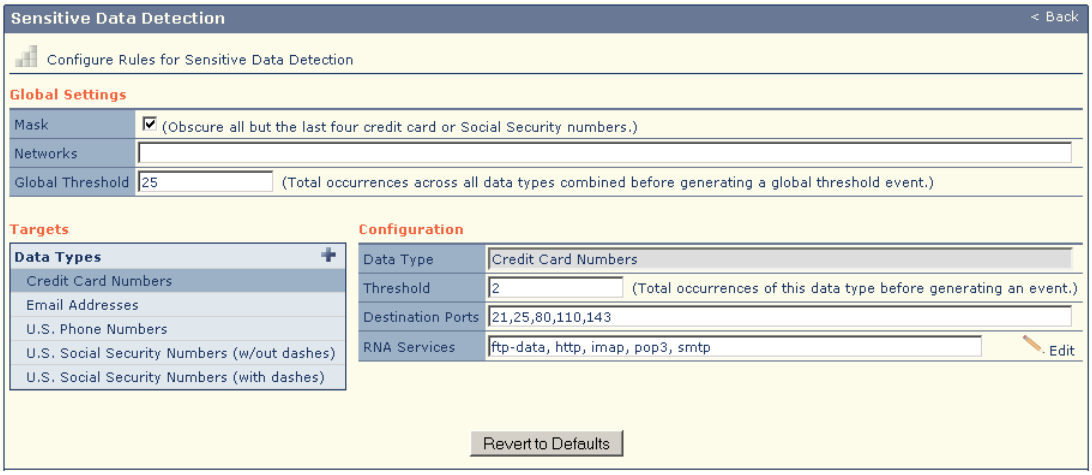
3. Click **Advanced Settings** in the navigation panel on the left.

   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.The following graphic shows the Defense Center version of the page.



   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 494 for more information.

5. You can take any of the actions described in the Sensitive Data Configuration Actions table on page 683.

6. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 685

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Selecting RNA Services to Monitor

Platform: DC

Component: IPS + RNA

If you manage IPS using a Defense Center with RNA, you can specify up to eight RNA services to monitor for each data type. See Working with Services on page 1040 for more information on the services RNA can detect on your network.

You must configure RNA and apply an RNA detection policy to your RNA detection engines for the preprocessor to monitor RNA services; see What is an RNA Detection Policy? on page 889 for more information. You must also enable Adaptive Profiles in your intrusion policy for IPS to detect data types in traffic for specific services. See Using Adaptive Profiles on page 697 for more information.

Note that IPS does not automatically enable adaptive profiles in your intrusion policy. If you enable a data type that specifies RNA services, but you do not enable adaptive profiles, IPS will not match your enabled data types against RNA services.

You must specify at least one RNA service 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 RNA services. For example, if you specify the HTTP service, you might also configure the well-known HTTP port 80. If a new host on your network implements HTTP, IPS will monitor port 80 during the interval when RNA is discovering the new HTTP service. Because you must enable adaptive profiles, IPS will also monitor HTTP service traffic on other standard and non-standard ports.

In the case where you want to detect sensitive data in FTP traffic, you must specify the `ftp-data` RNA service 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 689 and Decoding FTP and Telnet Traffic on page 538 for more information.

Keysight, EX-1004
Page 686

To modify RNA services to detect sensitive data:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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 494 for more information.

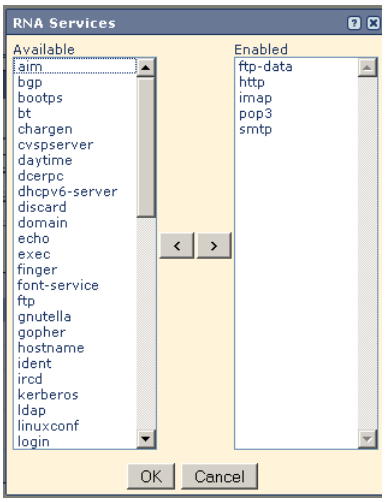
5. Click the data type name under Data Types to select the data type you want to modify.

   The Configuration area updates to display the current settings for the selected data type.

Keysight, EX-1004
Page 687

6. Click inside the **RNA Services** field, or click **Edit** next to the field.

   The RNA Services pop-up window appears.



7. You have two choices:

   - To add up to eight 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.

IMPORTANT! To detect sensitive data in FTP traffic, you must add the `ftp-data` service and enable the FTP/Telnet preprocessor. See Special Case: Detecting Sensitive Data in FTP Traffic on page 689 for more information.

8. You have two choices:

   - Click **OK** to add the service detectors.

     The Sensitive Data Detection page is displayed and the services are updated.
   - Click **Cancel** to abandon the selections.

     The Sensitive Data Detection page is displayed and services are not updated.

Keysight, EX-1004
Page 688

Special Case: Detecting Sensitive Data in FTP Traffic

Platform: DC

Component: IPS + RNA

You usually determine which traffic to monitor for sensitive data by specifying the ports to monitor or, optionally, specifying RNA services in deployments using RNA. However, specifying ports or RNA services is not sufficient for detecting sensitive data in FTP traffic. Sensitive data in FTP traffic is found in traffic for the FTP data service, 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` RNA service.

  Note that you can specify the RNA `ftp-data` service even when you do not use RNA in your deployment. In such cases, specifying the `ftp-data` service provides no benefit except to contribute to enabling detection of sensitive data in FTP traffic. See Selecting RNA Services to Monitor on page 686 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` RNA service does not invoke RNA service 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 538 for more information.
- Ensure that adaptive profiles is enabled.

  Note that you can enable adaptive profiles in a deployment that does not include RNA. In such cases, adaptive profiles provides no benefit except to contribute to enabling detection of sensitive data in FTP traffic. See Configuring Adaptive Profiles on page 702 for more information.
- Ensure that your configuration includes at least one port to monitor for sensitive data.

  Note that it is not necessary to specify an FTP port except in the unlikely case where you only want to detect sensitive data in FTP traffic. Most sensitive data configurations will include other ports such as HTTP or email service 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 683 or more information.

Keysight, EX-1004
Page 689

Using Custom Data Types

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can create and modify custom data types to detect data patterns that you specify. For example, a hospital might create a data type to protect patient numbers, or a university might create a data type to detect student numbers that have a unique numbering pattern.

Each custom date type you create also creates a single sensitive data preprocessor rule that has a generator ID (GID) of 138 and a Snort ID of 1000000 or greater, that is, a SID for a local rule. You must enable the associated sensitive data rule to enable detection, and event generation, for each custom data type you want to use in your policy. See Setting Rule States on page 435 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 417 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, and you cannot import custom sensitive data rules.

See the following sections for more information:

- Defining Data Patterns in Custom Data Types on page 690
- Configuring Custom Data Types on page 693
- Editing Custom Data Type Names and Detection Patterns on page 695

Defining Data Patterns in Custom Data Types

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 690

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.	`colou?r` matches `color` or `colour`
`{n}`	Matches the preceding character or escape sequence *n* times.	For example, `\d{2}` matches `55`, `12`, and so on; `\l{3}` matches `AbC`, `www`, and so on; `\w{3}` matches `a1B`, 25C, and so on; `x{5}` matches `xxxxx`
`\`	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class. See the Sensitive Data Pattern Character Classes table on page 692 for a description of the character classes you can use in sensitive data patterns.	`\?` matches a question mark, `\\` matches a backslash, `\d` matches numeric characters, and so on

You must use a backslash to escape the characters in the 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...
\?	?
\{	{
\}	}
\\	\

Keysight, EX-1004
Page 691

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}
```


Keysight, EX-1004
Page 692

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 680 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 435 for information on enabling rules in your intrusion policy.

To create or modify a custom data type:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 693

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:

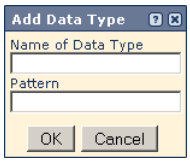
   - If the configuration is enabled, click **Edit**.
   - If the configuration is disabled, click **Enabled**, then click **Edit**.

   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 494 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 695 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 683 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 695.

   - 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.

Keysight, EX-1004
Page 694

Editing Custom Data Type Names and Detection Patterns

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 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 680 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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**.

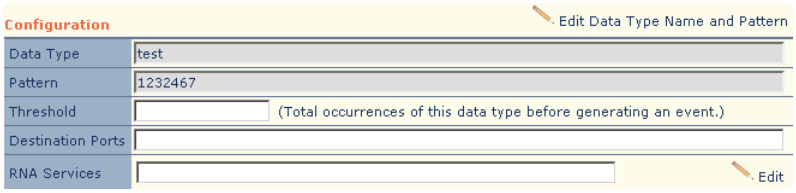
   The Sensitive Data Detection page appears.

   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 494 for more information.

Keysight, EX-1004
Page 695

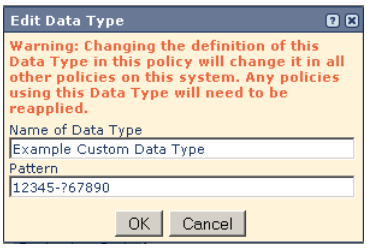
5. Click the name of the data type you want to modify.

   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 690 for information on specifying the data pattern.

   The Sensitive Data Detection page appears. If you clicked **OK**, the page displays your changes.

Keysight, EX-1004
Page 696

Chapter 16
Using Adaptive Profiles

Typically, IPS uses the static settings you configure in an intrusion policy to process and analyze traffic. With the adaptive profiles feature, however, IPS can adapt to network traffic by associating traffic with host and service information from the RNA network map and then processing the traffic accordingly.

As an example of how adaptive profiles can be useful, 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 IPS preprocessors reassemble data using a format other than that used for the operating system of the destination host, IPS may miss content that could be malicious when reassembled on the receiving host.

Adaptive profiles can also make the DCE/RPC, DNS, FTP/Telnet, HTTP, RPC, SMTP, SSH, and SSL preprocessors more efficient by checking traffic for service identification. As an example, if traffic does not contain a service identifier for HTTP and does not come over one of the standard HTTP ports, HTTP preprocessing is not attempted for that traffic.

Adaptive profiles also enable the rules engine to work more efficiently by selectively applying rules based on service metadata in a rule and the service information from the traffic.

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 698
- Configuring Adaptive Profiles on page 702

Keysight, EX-1004
Page 697

Understanding Adaptive Profiles

Platform: DC

Component: IPS + RNA

Adaptive profiles enable use of the most appropriate operating system profiles for IP defragmentation and for TCP stream preprocessing. They also allow IPS to apply rules designed for that specific type of traffic. In addition, they check for service identification information to make the DCE/RPC, DNS, FTP/Telnet, HTTP, RPC, SMTP, SSH, and SSL preprocessors operate more efficiently. For more information on the aspects of the intrusion policy affected by adaptive profiles, see the following topics:

- Using TCP Stream Preprocessing on page 627
- Defragmenting IP Packets on page 613
- Decoding HTTP Traffic on page 560
- Decoding SMTP Traffic on page 582
- Decoding FTP and Telnet Traffic on page 538
- Decoding DCE/RPC Traffic on page 514
- Detecting Exploits in DNS Name Server Responses on page 531
- Detecting Exploits Using the SSH Preprocessor on page 590
- Using the SSL Preprocessor on page 596

IPS can use data detected by RNA, obtained through an Nmap scan, or added through the host input feature to adapt processing behavior.

IMPORTANT! When you input host data from a third-party application using the command line import utility or the host input API you must first map the data to RNA product definitions so IPS can use it for adaptive profiles. For more information, see Managing Third-Party Product Mappings on page 1408.

Using Adaptive Profiles with Preprocessors

Platform: DC

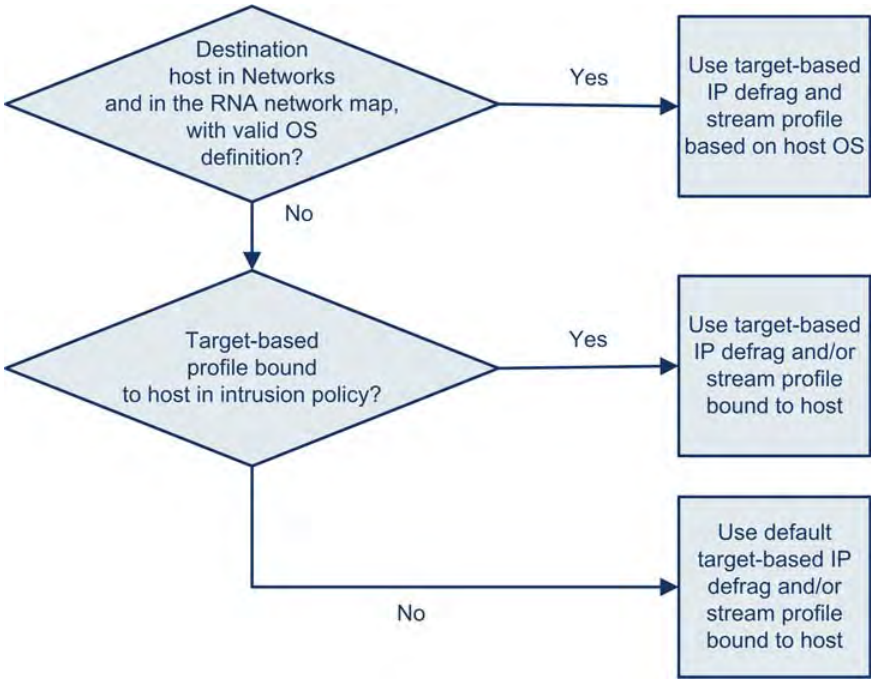
Component: IPS + RNA

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

Keysight, EX-1004
Page 698

operating system in the host profile for the target host, as illustrated in the following diagram.



For example, you create and apply 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.

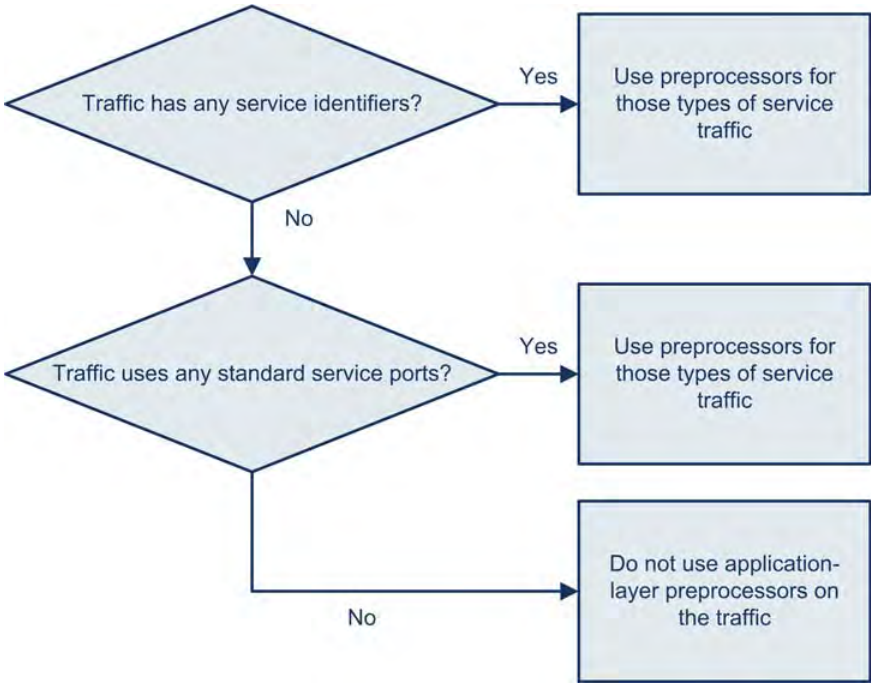
When IPS detects traffic from Host A, which is not in the 10.6.0.0/16 subnet, it uses the Linux target-based policy to reassemble IP fragments. However, when it detects traffic from Host B, which is in the 10.6.0.0/16 subnet, it retrieves Host B's operating system data from the network map, where Host B is listed as running Microsoft Windows XP Professional. IPS uses the Windows target-based profile to do the IP defragmentation for the traffic destined for Host B.

See Defragmenting IP Packets on page 613 for information on the IP Defragmentation preprocessor. See Using TCP Stream Preprocessing on page 627 for information on the stream preprocessor.

In addition, adaptive profiles use service information to identify service traffic on non-standard ports so it can be processed by the DCE/RPC, DNS, FTP/Telnet, HTTP, RPC, SMTP, SSH, and SSL preprocessors. The preprocessor engine checks each packet first to see if it has a service identifier, then checks to see if it traveled over one of the standard ports for a service, as shown in the following diagram.

Keysight, EX-1004
Page 699

Note that, as a group, the DCE/RPC, DNS, FTP/Telnet, HTTP, RPC, SMTP, SSH, and SSL preprocessors are referred to as the application-layer preprocessors.



For example, if an HTTP packet arrives over port 2020, without adaptive profiles that traffic would not be preprocessed using the HTTP preprocessor. However, with adaptive profiles enabled, if the packet has an HTTP service identifier, the HTTP preprocessor would process that packet.

Using Adaptive Profiles with Intrusion Rules

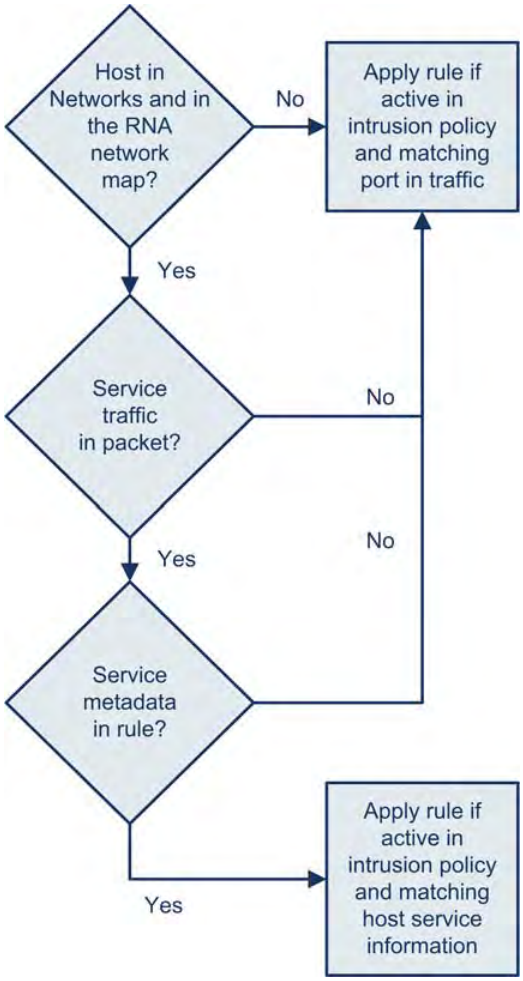
Platform: DC

Component: IPS + RNA

With adaptive profiles configured, the rules engine looks for active rules with service metadata that matches the service information for the host in a packet. If it matches, IPS uses the rule to analyze the traffic. If it does not match, IPS does not apply the rule to the traffic. If a host does not have service information, or if the rule does not have service metadata, IPS checks the port in the traffic against

Keysight, EX-1004
Page 700

the port in the rule to determine whether to apply the rule to the traffic, as shown in the following diagram.



For example, if a host is running an Apache Tomcat server that uses the HTTP service over port `8080`, a rule intended to catch HTTP exploits which checks traffic on port `80` cannot identify traffic on port `8080` as HTTP traffic. However, if adaptive profiles are active and the rule contains service metadata set to `http`, IPS obtains information about the `http` service running on the host from the RNA network map and associates that information with the traffic. The rule then compares its metadata to the traffic metadata, and processes the traffic even though it occurs on a non-standard port.

See Understanding and Writing Intrusion Rules on page 761 for more information.

Keysight, EX-1004
Page 701

Adaptive Profiles and RNA Recommended Rules

Like RNA recommended rules, adaptive profiles compare metadata in a rule to RNA host information to determine whether a rule should apply for a particular host. However, while RNA recommended rules provide recommendations for enabling or disabling rules using that information, adaptive profiles use the information to dynamically adjust application of specific rules for specific traffic.

RNA 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, RNA 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 RNA recommended rules within 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 RNA Rule State Recommendations on page 468 for more information.

Configuring Adaptive Profiles

Platform: DC

Component: IPS + RNA

To use RNA host information to determine which target-based profiles are used for IP defragmentation and TCP stream preprocessing and which rules are applied for service traffic from a host, 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 RNA network map and must be in the segment monitored by the detection engine where you apply the intrusion policy.

You can indicate the hosts in the RNA 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 371 for information on using variables.

You can use any of these addressing methods alone or in any combination as a list of IP addresses, CIDR blocks, or variables separated by commas, as shown in the following example:

```
192.168.1.101, 192.168.4.0/24, $HOME_NET
```


Keysight, EX-1004
Page 702

For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

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 RNA network map data is synced from the Defense Center to the 3D Sensor. IPS uses the data on the sensor to determine what profiles should be used when processing traffic.

To configure adaptive profiles:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

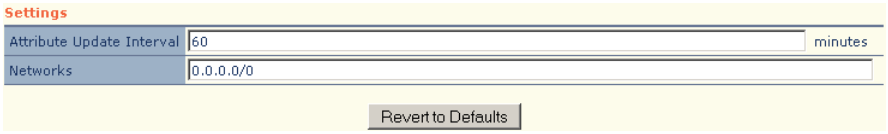
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



   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 494 for more information.

Keysight, EX-1004
Page 703

5. Optionally, in the **Attribute Update Interval** field, type the number of minutes that should elapse between synchronization of RNA network map data from the Defense Center to the 3D Sensor.

   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, CIDR block, or variable, or a list that includes any of these addressing methods separated by commas, to identify any host in the RNA network map for which you want to use adaptive profiles.

   See Managing RNA Rule State Recommendations on page 468 for information on configuring variables in an intrusion policy. See Creating RNA Detection Policies on page 900 for information on configuring the RNA network map.

7. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 704

Chapter 17
Using Global Rule Thresholding

You can use thresholds to limit the number of times the system logs and displays intrusion events. Thresholds cause IPS 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 705 and Configuring Global Thresholds on page 708.
- 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 439.

Understanding Thresholding

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 705

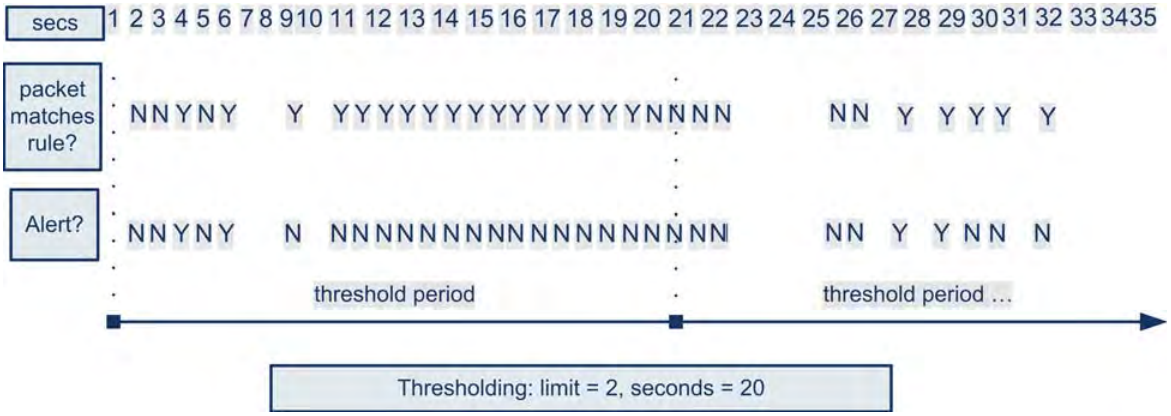
You can also override this threshold by setting individual thresholds on specific rules. For example, you might set a global limit threshold of five events every 60 seconds, but then set a specific threshold of ten events for every 60 seconds for SID 1315. All other rules generate no more than five events in each 60 second period, but IPS 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 439.

TIP! If you apply a policy containing a global or individual threshold to a detection engine running on a sensor with multiple CPUs, you may receive 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 IPS does not generate any more events during that period.



Understanding Thresholding Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 706

events. When you configure global thresholding, first, specify the thresholding type, as described in the following 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, IPS stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the count argument) trigger the rule during the specified time period. Note that the counter for the time restarts once you hit the threshold count of events and IPS 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. IPS 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, IPS logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to two, and **Seconds** to 10, the following event counts result: • If the rule is triggered once in 10 seconds, IPS does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, IPS generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, IPS generates one event (the threshold is met when the rule triggered the second time and following events are ignored)

Keysight, EX-1004
Page 707

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, IPS 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 IPS 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 705.

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 708

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

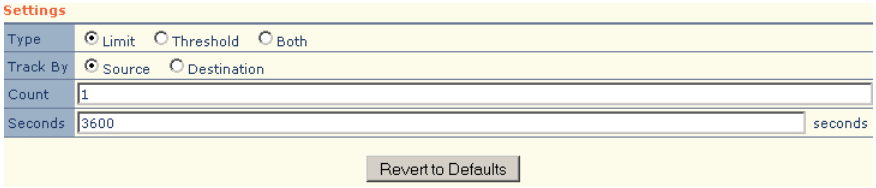
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



   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 494 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.

Keysight, EX-1004
Page 709

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.

     For a threshold threshold, 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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 710

Disabling the Global Threshold

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.



Keysight, EX-1004
Page 711

5. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 712

Chapter 18
Using Performance Settings in an Intrusion Policy

Sourcefire provides several features for improving IPS performance. See the following sections for more information.

- Event Queue Configuration on page 714 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 716 describes how you can balance security with the need to maintain sensor latency at an acceptable level with packet latency thresholding.
- Understanding Rule Latency Thresholding on page 720 describes how you can balance security with the need to maintain sensor latency at an acceptable level with rule latency thresholding.
- Performance Statistics Configuration on page 726 describes how you can configure the basic parameters of how 3D Sensors with IPS monitor and report on their own performance.
- Constraining Regular Expressions on page 728 describes how you can override default match and recursion limits on PCRE regular expressions.
- Rule Processing Configuration on page 731 describes how you can configure rule processing event queue settings.

Keysight, EX-1004
Page 713

Event Queue Configuration

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

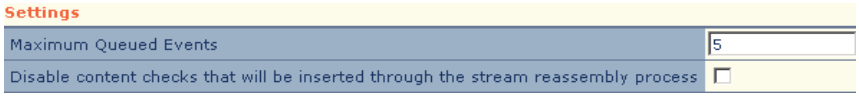
3. Click **Advanced Settings** in the navigation panel on the left.

   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 494 for more information.

Keysight, EX-1004
Page 714

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. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 715

Understanding Packet Latency Thresholding

Platform: DC/MDC/Standalone Sensor

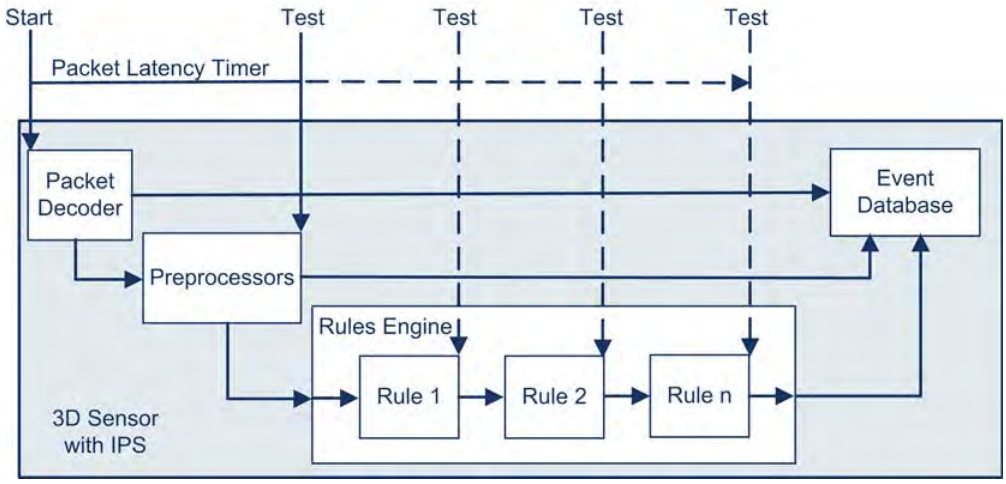
Component: IPS

You can balance security with the need to maintain sensor 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

Keysight, EX-1004
Page 716

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 435.

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 since processing simply moves to the next packet.

See the following sections for more information:

- Setting Packet Latency Thresholding Options on page 718.
- Configuring Packet Latency Thresholding on page 719.

Keysight, EX-1004
Page 717

Setting Packet Latency Thresholding Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 718 for recommended minimum threshold settings.

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

Keysight, EX-1004
Page 718

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can enable or disable packet latency thresholding and modify the latency threshold.

To configure packet latency thresholding:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 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.



5. You can modify either of the options on the Latency-Based Packet Handling page.

   See Setting Packet Latency Thresholding Options on page 718 for a full description of each available option. See the Minimum Packet Latency Threshold Settings table on page 718 for recommended minimum **Threshold** settings.

6. 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 492 for more information.

Keysight, EX-1004
Page 719

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding Rule Latency Thresholding

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can balance security with the need to maintain sensor 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.

Keysight, EX-1004
Page 720

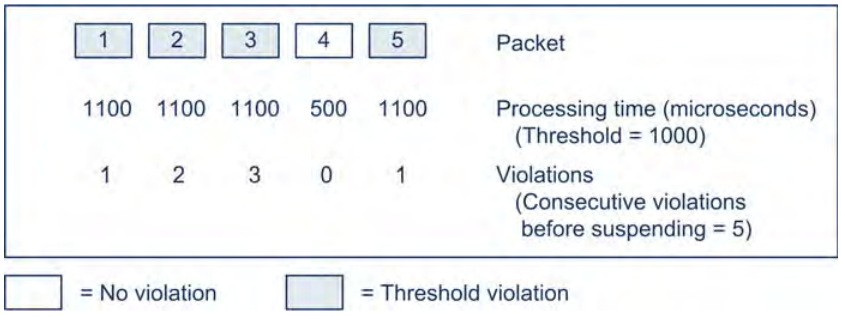
The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. However, rule latency thresholding gives you a tool you can use to balance security with connectivity.

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, IPS increments a counter. If the number of consecutive threshold violations reaches a specified number, IPS 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

IPS 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.

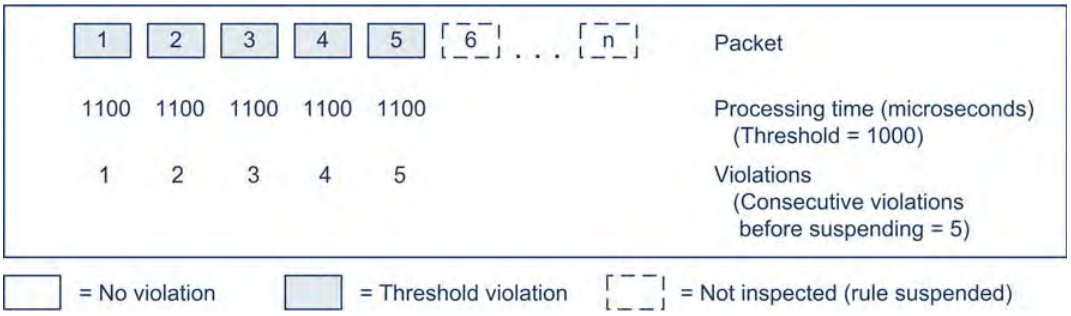
The example in the following illustration shows five consecutive rule processing times that do not result in rule suspension.



In the above example, the time required to process each of the first three packets violates the rule latency threshold of 1000 microseconds, and the violations counter increments with each violation. Processing of the fourth packet does not violate the threshold, and the violations counter resets to zero. The fifth packet violates the threshold and the violations counter restarts at one.

Keysight, EX-1004
Page 721

The example in the following illustration 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.

You can easily identify events triggered when IPS suspends or re-enables rules because IPS modifies the rule generator ID (GID) in intrusion event table views. An event for suspended rules uses a GID of 134 and a Signature ID (SID) of 1. An event for re-enabled rules uses a GID of 134 and a Signature ID (SID) of 2.

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 435.

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

Keysight, EX-1004
Page 722

sensor 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 723.
- Configuring Rule Latency Thresholding on page 724.

Setting Rule Latency Thresholding Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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**. 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 724 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

Keysight, EX-1004
Page 723

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

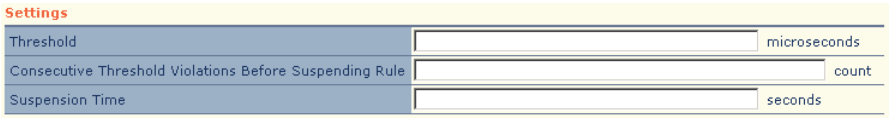
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

Keysight, EX-1004
Page 724

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 494 for more information.

5. You can modify any of the options on the Latency-Based Rule Handling page.

   See Setting Rule Latency Thresholding Options on page 723 for a full description of each available option. See the Minimum Rule Latency Threshold Settings table on page 724 for recommended minimum **Threshold** settings.

6. 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 492 for more information.

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 725

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Performance Statistics Configuration

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure the basic parameters of how 3D Sensors with IPS monitor and report on their own performance. This allows you to specify the intervals at which the system updates performance statistics on your 3D Sensors by configuring the following:

- number of seconds
- number of packets analyzed

WARNING! Do not apply a 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 726

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

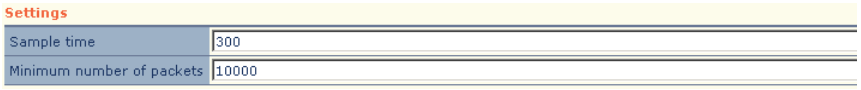
   The Policy Information page appears.

3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

4. Click **Edit**.

   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 494 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**.

WARNING! Do not apply a policy with the Performance Statistics troubleshooting **Log Session/Protocol Distribution** check box selected unless directed to do so by Sourcefire Support.

Keysight, EX-1004
Page 727

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Constraining Regular Expressions

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 800 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.

Keysight, EX-1004
Page 728

Note that when a rule that requires this feature is enabled in an intrusion policy, you must enable the feature or choose to allow IPS to enable it automatically before you can save the policy. For more information, see Automatically Enabling Advanced Settings on page 490.

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	Specifies whether to override **Match Limit**. You have the following options: • select **Default** to use the value configured for **Match Limit** • select **Unlimited** to permit an unlimited number of attempts. • select **Custom** to specify either a limit of 1 or greater for **Match Limit**, or to specify 0 to completely disable PCRE match evaluations
Match Limit	Specified the number of times to attempt to match a pattern defined in a PCRE regular expression.
Match Recursion Limit State	Specifies whether to override **Match Recursion Limit**. You have the following options: • select **Default** to use the value configured for **Match Recursion Limit** • select **Unlimited** to permit an unlimited number of recursions • select **Custom** to specify either a limit of 1 or greater for **Match Recursion Limit**, or to specify 0 to completely disable PCRE recursions Note that for **Match Recursion Limit** to be meaningful, it must be smaller than **Match Limit**.
Match Recursion Limit	Specifies the number of recursions when evaluating a PRCE regular expression against the packet payload.

To configure PCRE overrides:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

Keysight, EX-1004
Page 729

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

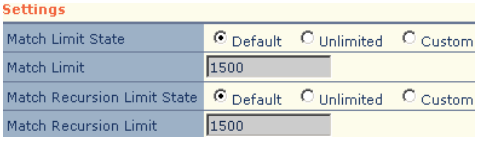
3. Click **Advanced Settings** in the navigation panel on the left.

   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 494 for more information.

5. You can modify any of the options in the Regular Expression Constraint Options table on page 729.

6. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 730

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Rule Processing Configuration

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 Procession Configuration Options table describes the options you can configure to determine how many events are logged per packet or stream.

Rule Procession 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: • 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.

Keysight, EX-1004
Page 731

To configure how many events are logged per packet or stream:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

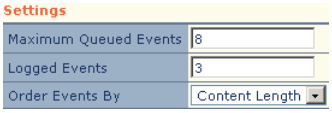
3. Click **Advanced Settings** in the navigation panel on the left.

   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.

   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 494 for more information.

5. You can modify any of the options on the Rule Processing Configuration page.

   See the Rule Procession Configuration Options table on page 731 for a full description of each available option.

Keysight, EX-1004
Page 732

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 733

Chapter 19
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. Additionally, you can specify Check Point™ OPSEC™ firewall responses for intrusion events generated when a standard text rule triggers. 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 439 for more information.

See the following sections for more information:

- Using Check Point OPSEC Responses on page 735 explains how to configure Check Point OPSEC firewalls to respond to triggered rules with defined responses.
- Using SNMP Responses on page 748 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.

Keysight, EX-1004
Page 734

- Using Syslog Responses on page 752 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 756 describes the options you can configure to send notifications of intrusion events by email.

Using Check Point OPSEC Responses

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure an intrusion policy to initiate Check Point™ OPSEC™ Suspicious Activity Monitor (SAM) firewall responses whenever specified intrusion rules are triggered. By applying the policy to a detection engine, you integrate with the Check Point firewall, causing the 3D Sensor to start a firewall response whenever rules to which you add OPSEC alerting in the policy trigger. See Adding OPSEC Alerts on page 460 for more information.

You need to configure both the Check Point firewall and the 3D Sensor, and you **must** set up the Check Point firewall first. For details on configuring your Check Point firewall for use with the 3D Sensor, see Creating the OPSEC SAM Application on page 736. For additional information, see your Check Point documentation or contact support for your firewall vendor.

If you are using the Defense Center to manage your 3D Sensor, you can use the Defense Center web interface to configure OPSEC settings for the sensor.

After configuring your Check Point firewall, you must configure your sensor's connection to the SAM server and configure a peer trust for the sensor.

You can enable and disable SAM responses within a policy on the Layer summary page. See Modifying Advanced Settings on page 478 for information on enabling and disabling advanced features in an intrusion policy.

IMPORTANT! OPSEC settings in intrusion policies applied to 3Dx800 sensors have no effect.

See the following sections for more information:

- Creating the OPSEC SAM Application on page 736
- Connecting the 3D Sensor and the SAM Server on page 739
- Selecting a SAM Server on page 743
- Clearing All OPSEC SAM Responses for a Detection Engine on page 745
- Adding and Deleting Firewall Objects on page 745
- Viewing Sourcefire SAM Client System Information on page 746
- Viewing OPSEC Debug Messages on page 747

Keysight, EX-1004
Page 735

Creating the OPSEC SAM Application

Platform: DC/MDC/Standalone Sensor

Component: IPS

Before setting up OPSEC properties in the web interface, you must first configure a Check Point OPSEC SAM application that communicates with the 3D Sensor. After creating the Check Point application, you configure the connection between the 3D Sensor and the SAM server, associate firewall responses with existing rules within a custom intrusion policy, and then apply that policy to a detection engine. See Connecting the 3D Sensor and the SAM Server on page 739 for detailed information about configuring connectivity and setting firewall responses for rules. If you are configuring multiple 3D Sensors with IPS for use with a Defense Center, repeat this procedure to create an application for each managed sensor. See Adding OPSEC Alerts on page 460 for information on associating firewall responses with existing rules.

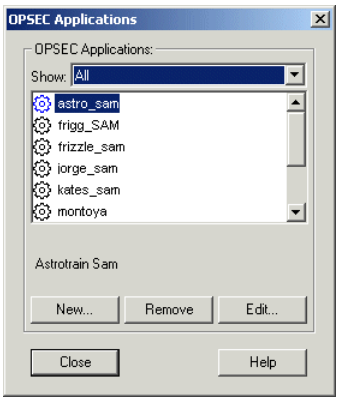
IMPORTANT! The default port number for an OPSEC SAM server is 18183/tcp. For more information on modifying the default port for an OPSEC SAM server, as well as more detailed information about using Check Point OPSEC SAM and the Check Point SmartDashboard™, see the Check Point documentation.

To create and configure an OPSEC SAM application connection to the 3D Sensor:

Access: P&R Admin/Admin

1. Open the Check Point SmartDashboard and, from the **Manage** menu, select **OPSEC Applications**.

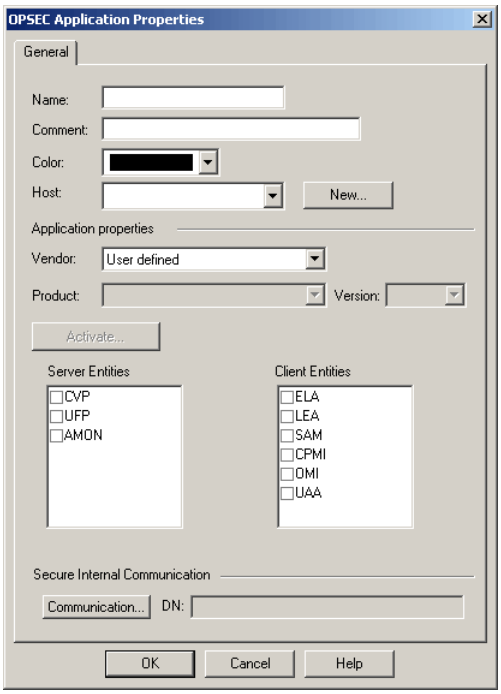
   The OPSEC Applications dialog box appears.



Keysight, EX-1004
Page 736

2. Click **New** and, from the choices that appear, select **OPSEC Application**.

   The OPSEC Application Properties dialog box appears.



3. Fill in the fields as follows:

   - In the **Name** field, type the name for the new OPSEC application that you will use to connect to the 3D Sensor (for example, you could use *sf_sam* as the application name).

     Note the name that you specify. You will need to enter it later on the 3D Sensor when configuring connectivity.

   - Optionally, enter a comment in the **Comment** text box and select a color from the **Color** list.

     The Comment and Color features are only used by Check Point management software and have no effect on the connection with the 3D Sensor.

Keysight, EX-1004
Page 737

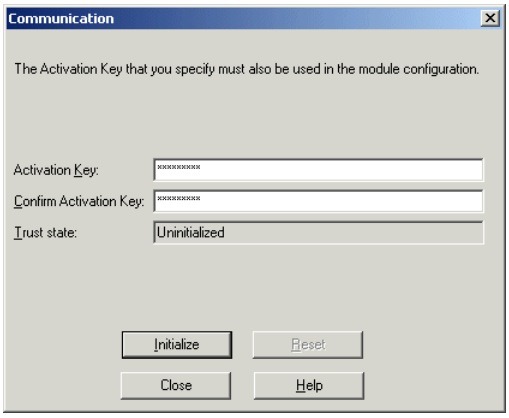
- In the **Host** list, select the 3D Sensor with which the firewall will communicate.

IMPORTANT! If you are managing your 3D Sensor with a Defense Center, you must configure the sensor, **not** the Defense Center, as the host. The Defense Center pushes the OPSEC configuration to managed 3D Sensors by applying policies that contain the OPSEC settings, but does so by managing the peer trust between the sensor and the firewall. Note that OPSEC settings in intrusion policies applied to 3Dx800 sensors have no effect.

  If the 3D Sensor does not appear in the list, click **New**, enter the sensor's fully qualified domain name and IP address and click **OK**.
- Make no changes to the **Vendor** list; it should be set to **User Defined**.
- Make no changes to the **Server Entities** box; all check boxes should remain disabled.
- In the **Client Entities** box, select the **SAM** check box to indicate that this application will implement Suspicious Activity Monitoring.

4. Click **Communication**.

   The Communication dialog box appears.



5. In the **Activation Key** and **Confirm Activation Key** text boxes, enter a password.

   The Communication dialog box allows you to specify an activation key to be used to ensure a secure transfer of authentication certificates between the appliance and the Check Point application.

IMPORTANT! Remember the Activation Key that you use. You must also enter it on the appliance where you are configuring the response) to ensure a secure transfer of authentication certificates between the 3D Sensor and the Check Point application.

Keysight, EX-1004
Page 738

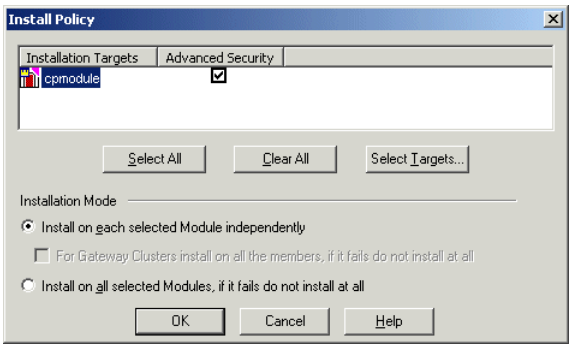
6. Click **Initialize** to create authentication certificates for the application, and then click **Close**.

   The Communication dialog box closes and you return to the OPSEC Application Properties page.

7. Click **OK** to close the OPSEC Application Properties dialog box, and then click **Close** to close the OPSEC Applications dialog box.

8. From the **Policy** menu on the Check Point Smart Dashboard, select **Install**.

   The Install Policy dialog box appears.



9. Select the module, enable **Install on each selected Module independently**, and click **OK**.

   The new policy is installed on the firewall.

10. Click **Close** to close the Installation Process dialog box.

    Your Check Point system is configured for OPSEC SAM communication with the configured sensor. Proceed to Connecting the 3D Sensor and the SAM Server on page 739 to configure the 3D Sensor connection to the Check Point OPSEC SAM.

TIP! If you are managing other 3D Sensors with a Defense Center, repeat this procedure with the other sensors where it is required for your deployment.

Connecting the 3D Sensor and the SAM Server

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure your 3D Sensor to initiate Check Point OPSEC Suspicious Activity Monitor (SAM) firewall responses when rules configured for IPS are triggered. To manage this interaction, you must configure how the 3D Sensor communicates with the specified OPSEC SAM server and define the appropriate

Keysight, EX-1004
Page 739

firewall responses. You can perform this procedure on the standalone sensor web interface or on the web interface of the Defense Center managing your sensors.

To set up communication between the 3D Sensor and the SAM server:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > OPSEC > Peer**.

   The Peer page appears.

   TIP! If you have not yet created a SAM server, you can also access the Peer page by clicking the link **No SAM servers are configured. You must configure a SAM server to use OPSEC**. on the OPSEC Configuration policy page.

2. Click **Add New Peer**.

   The OPSEC Peers section expands.

3. In the **Server** field, enter the IP address of the firewall server.

4. Locate your Check Point module distinguished name:

   - If you do not know the Check Point module name, continue with step 5 for instructions on locating it.
   - If you already know the Check Point module name, skip to step 9 to enter it.

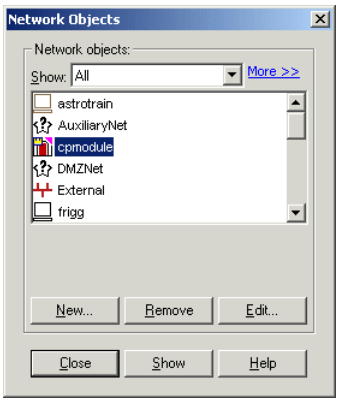
   WARNING! The Check Point module distinguished name is **not** the same as the Check Point application distinguished name that was created in Creating the OPSEC SAM Application on page 736. The following procedures provide detailed information about accessing the module distinguished name.

5. Open the Check Point Smart Dashboard.

Keysight, EX-1004
Page 740

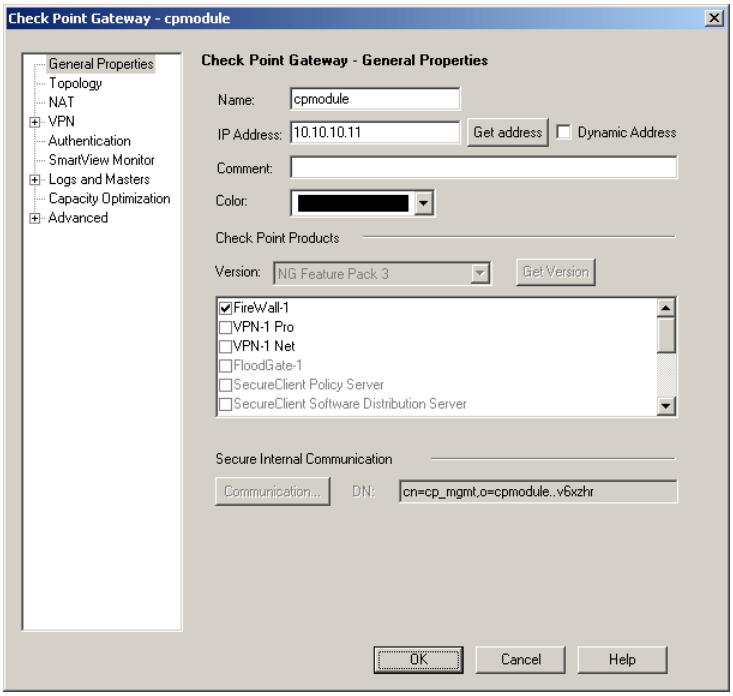
6. From the **Manage** menu, select **Network Objects**.

   The Network Objects dialog box appears.



7. Select the Check Point module (typically **cpmodule**), and click **Edit**.

   The General Properties page for the module appears.



8. Under Secure Internal Communication, select and copy the module's distinguished name.

   In this example, the Check Point module DN name is `cn=cp_mgmt,o=cpmodule..v6xzhr`. This will be different for your Check Point server, but the format should be similar.

Keysight, EX-1004
Page 741

9. In the **Check Point module DN** field on the Add New Peer page, enter the Check Point firewall server distinguished name (also known as the Secure Internal Communication, or SIC, name).

   WARNING! The Check Point module DN name is **not** the same as the distinguished name for the application. If you are unsure of the module name, follow steps 5 through 8 of this procedure to locate it.

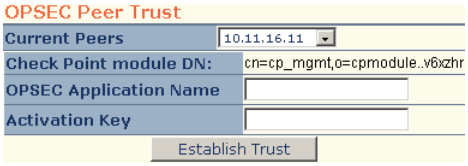
10. Click **Add**.

    The OPSEC Peers section collapses and refreshes to show the Peers drop-down list with the new peer added to the list.

    TIP! To delete a peer, select the IP address of the SAM server from the Peers list and click **Delete**. Note that you cannot delete a peer that you are using to respond to intrusion events within a custom policy; you must first either disable SAM responses (see Modifying Advanced Settings on page 478) or set your SAM server to "none" (see Adding OPSEC Alerts on page 460).

11. Click **Peer Trust**.

    The Trust page appears.



12. If you are using the Defense Center web interface to configure OPSEC settings, select the 3D Sensor that you want to set up peer trust for from the **Please select a sensor to configure Peer Trust** list.
13. Select the peer for which you want to configure trust from the **Current Peers** drop-down list.
14. In the **OPSEC Application Name** field, type the name of the OPSEC application with which you are integrating the sensor.

    TIP! This is the name of the OPSEC application you created in Creating the OPSEC SAM Application on page 736.

15. In the **Activation Key** field, type the Activation key that you specified in the Communication dialog box of your Check Point module.

Keysight, EX-1004
Page 742

16. Click **Establish Trust**.

    This connects to the Check Point module and transfers the authentication certificate. You can now configure firewall responses within a custom policy and apply it to a detection engine.

Selecting a SAM Server

Platform: DC/MDC/Standalone Sensor

Component: IPS

The OPSEC Configuration page allows you to select the Check Point™ OPSEC™ Suspicious Activity Monitor (SAM) firewall that responds when specific standard text rules within an intrusion policy trigger. If you have not yet created a SAM server, you can also access the OPSEC Peer page from OPSEC Configuration page to create a server.

To select a SAM server:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

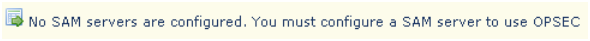
3. Click **Advanced Settings** in the navigation panel on the left.

   The Advanced Settings page appears.

4. You have two choices, depending on whether **OPSEC Configuration** under External Responses is enabled:

   - If the configuration is enabled, click **Edit**.
   - If the configuration is disabled, click **Enabled**, then click **Edit**.

   The OPSEC 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 494 for more information.

Keysight, EX-1004
Page 743

5. From the **Servers** list, select the firewall that you want to respond to triggered rules.

   TIP! If you have not yet created a SAM server, click the link **No SAM servers are configured. You must configure a SAM server to use OPSEC**. For more information, see Connecting the 3D Sensor and the SAM Server on page 739.

6. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.
   - You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
   - You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

     See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Keysight, EX-1004
Page 744

Clearing All OPSEC SAM Responses for a Detection Engine

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can clear applied or active OPSEC SAM responses from the Check Point OPSEC SAM server. For information on clearing or adding individual responses in an intrusion policy, see Adding OPSEC Alerts on page 460.

IMPORTANT! Clearing OPSEC SAM responses clears rule configurations from the firewall, but does **not** change the rule configuration on the detection engine.

To clear SAM rule configurations from a firewall:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > OPSEC > Response**.

   The Response page appears.

2. Select the detection engines where you want to clear OPSEC SAM responses from the firewalls with which they communicate and click **Clear SAM Responses**.

   WARNING! Clicking **Clear SAM Responses** clears all rules installed on the firewall and therefore may clear rules that were installed on your firewall by other detection engines using the same peer.

   Responses are cleared from the firewall communicating with the selected detection engine.

Adding and Deleting Firewall Objects

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can add or delete firewall objects that you configure when you add OPSEC responses for individual rules in an intrusion policy. See Adding OPSEC Alerts on page 460 for more information.

To add or delete firewall objects:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > OPSEC > Peer**.

   The Peer page appears.

Keysight, EX-1004
Page 745

2. Perform one of the following:
   - To add a new firewall object, type the name of the firewall object in the **Firewall Host** field and click **Add**. See your system administrator to identify the firewall objects in your Check Point deployment.
   - To delete a user-defined firewall object, select the firewall object you want to delete from the list and click **Delete**.

Viewing Sourcefire SAM Client System Information

Platform: DC/MDC/Standalone Sensor

Component: IPS

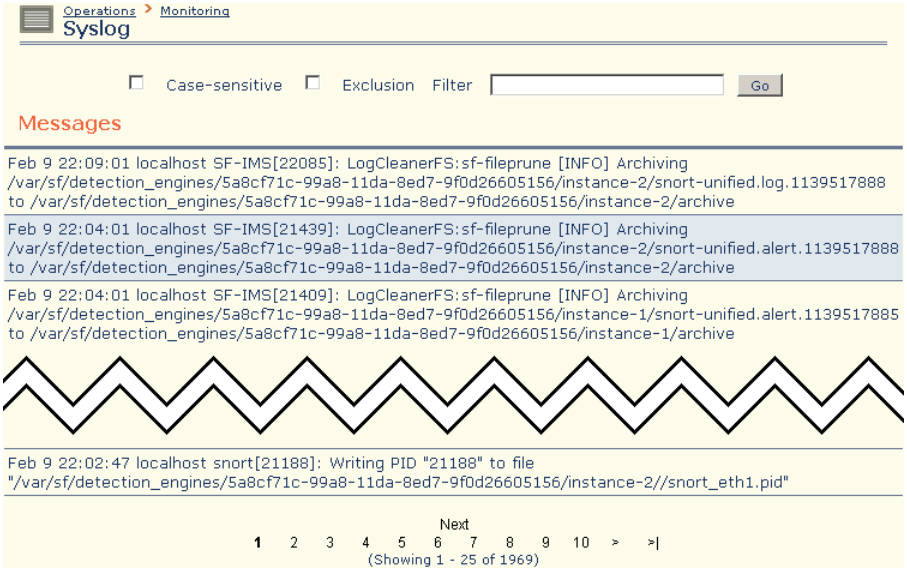
You can view system log messages for the Sourcefire SAM client from the web interface. The SAM Client module name is `SFReactD`.

To view the status associated with the SAM Client:

Access: Maint/Admin

1. Select **Operations > Monitoring > Syslog**.

   A list of system log messages appears.



2. In the **Filter** box, type `SFReactD` and click **Go**.

   Messages associated with the Sourcefire SAM client appear.

Keysight, EX-1004
Page 746

Viewing OPSEC Debug Messages

Platform: DC/MDC/Standalone Sensor

Component: IPS

To view OPSEC debug messages, you must be connected to the 3D Sensor via ssh and logged in using the `admin` user. The debug messages are sent to the screen rather than to the web interface.

To view OPSEC debug messages:

Access: Admin

1. Connect to the sensor via ssh and log in using the admin account.
2. Enter expert at the prompt.
3. Run the following script to determine the UUIDs for the detection engines where you set up and applied the intrusion policy:

   ```
   /var/sf/bin/de_info.pl
   ```

   The output looks like this:

   ```
   DE Name         : IS1 DE1
   DE Type         : ids
   DE Description  : DE1 on IS1
   DE Resources    : 1
   DE UUID         : ef0009fe-8d54-11da-000c-afa1eb00002c
   DE InterfaceSet : passive interfaces

   DE Name         : IS1 DE2
   DE Type         : ids
   DE Description  : DE2 on IS1
   DE Resources    : 1
   DE UUID         : a50006f0-8db2-11da-000c-94e75200004b
   DE InterfaceSet : passive interfaces
   ```

4. Use the following commands (where *DE_UUID* is the UUID that you determined in the previous step) to stop the `SFReactD` daemon and start it again in the foreground:

   ```
   pmtool disablebyid de_uuid-react
   ```

   and then, all on one line, enter:

   ```
   SFReactD --nodaemon
   -X /var/sf/detection_engines/de_uuid/SFReactD.pid
   -c /var/sf/de_uuid/SFReactD.conf
   ```

   The debug messages are directed to the screen.

Keysight, EX-1004
Page 747

5. Use the following command to restart the `SFReactD` daemon:

   ```
   pmtool enablebyid de_uuid-react
   ```

Using SNMP Responses

Platform: DC/MDC/Standalone Sensor

Component: IPS

An *SNMP trap* is a network management notification. You can configure the sensor 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 sensor that detected it
- the name of the detection engine that detected it
- the event data

Intrusion event response SNMP traps use the Message Digest 5 (MD5) authentication protocol and the Data Encryption Standard (DES) privacy protocol.

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 478.

TIP! If your network management system requires a management information base file (MIB), you can obtain it from the Defense Center or the 3D Sensor 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 Server	The server that will receive SNMP traps notification.
Community String	The community name.

Keysight, EX-1004
Page 748

SNMP v3 Options

For SNMP v3, you can specify the options described in the SNMP v3 Options table.

SNMP v3 Options

Option	Description
Trap Server	The server that will receive SNMP traps notification.
Authentication Password	The password required for authentication. If you specify an authentication password, authentication is enabled.
Private Password	The SNMP key for privacy. 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 749.

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.

Configuring SNMP Responses

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure SNMP alerting in an intrusion policy. After you apply the policy to a detection engine, the detection engine notifies you of any intrusion events it

Keysight, EX-1004
Page 749

detects via SNMP trap. For more details on SNMP alerting, see Using SNMP Responses on page 748.

IMPORTANT! You must select SNMP v2 **or** SNMP v3.

To configure SNMP alerting options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

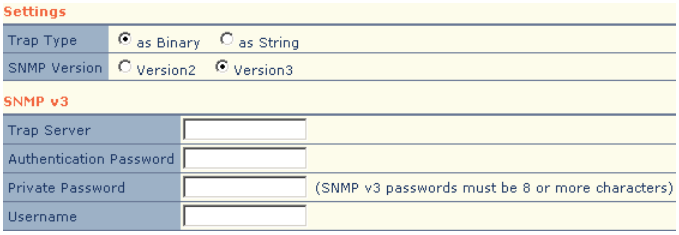
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



   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 494 for more information.

Keysight, EX-1004
Page 750

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 748.
   - 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 749 for more information.

IMPORTANT! When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format.

7. You have the following options:
   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.
   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.
   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 751

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Using Syslog Responses

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 policy to a detection engine, the detection engine 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

Keysight, EX-1004
Page 752

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.
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

Keysight, EX-1004
Page 753

Syslog Priority Levels (Continued)

Level	Description
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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure syslog alerting in an intrusion policy. After you apply the policy to a detection engine, the detection engine notifies you of any intrusion events it detects via the syslog. For more information on syslog alerting, see Using Syslog Responses on page 752.

To configure syslog alerting options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Intrusion Policy.**

   The Intrusion Policy page appears.

2. Click **Edit** next to the policy you want to edit.

   If you have unsaved changes in another policy, click **OK** to discard those changes and continue. To save the changes, click **Cancel**, open the other policy and commit your changes, then return to the beginning of this procedure. See Committing Intrusion Policy Changes on page 345 for information on saving unsaved changes in another policy.

   The Policy Information page appears.

Keysight, EX-1004
Page 754

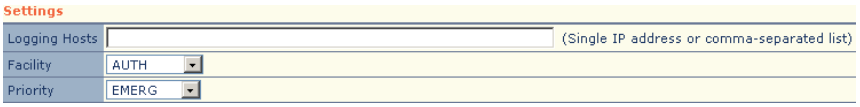
3. Click **Advanced Settings** in the navigation panel on the left.

   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.



   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 494 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 752 for details on facility and priority options.

7. You have the following options:

   - You can continue editing your policy.

     See Editing an Intrusion Policy on page 341 for descriptions of other intrusion policy settings you can modify.

   - You can remove all saved and unsaved changes on this page by reverting to the default configuration settings in the base policy. Click **Revert to Defaults**, then click **OK** at the prompt. See Understanding the Base Policy on page 361 for more information.

   - You can save your changes at this time. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Commit Changes**.

     The Intrusion Policy page appears.

     See Committing Intrusion Policy Changes on page 345 for descriptions of prompts you might encounter when you save your changes.

     You must apply your policy to the appropriate detection engines to put your changes into effect. See Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 755

- You can discard all unsaved changes. If necessary, click **Policy Information** in the navigation panel on the left side of the page to return to the Policy Information page, then click **Discard Changes**; 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.
- You can exit the policy, leaving your changes in the system cache. On exiting, click **OK** when prompted, or click **Cancel** to remain in the advanced editor.

  See Committing Intrusion Policy Changes on page 345 for information on how the system caches one policy per user.

Understanding Email Alerting

Platform: DC/MDC/Standalone Sensor

Component: IPS

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. You configure email alerting for any detection engine on the appliance rather than per detection engine. Your email alert settings are used regardless of the policy in place for each detection engine on the sensor.

Keysight, EX-1004
Page 756

The Email Alerting Parameters table describes the parameters you can set for email alerting.

Email Alerting Parameters

Parameter	Description
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

Keysight, EX-1004
Page 757

Email Alerting Parameters (Continued)

Parameter	Description
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 sensor 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, on a 3D Sensor:
	`2011-05-18 10:35:10 icmp 10.10.10.1:8 -> 10.2.1.3:0 snort_decoder: Unknown Datagram decoding problem! (116:108)`
	
	On a Defense Center:
	`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 759.

Keysight, EX-1004
Page 758

Configuring Email Alerting

Platform: DC/MDC/Standalone Sensor

Component: IPS

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:**

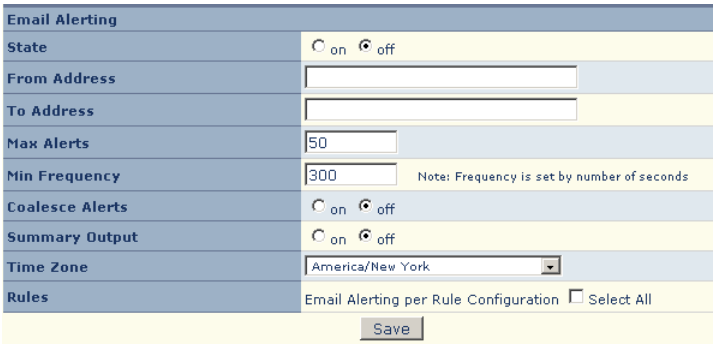
- configure your mail host to receive email alerts (see Configuring a Mail Relay Host and Notification Address on page 1751)
- make sure that both the sensor and the Defense Center can reverse resolve their own IP addresses

To configure email alerting options:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Email**.

   The Edit Email Alerting page appears.



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**.

Keysight, EX-1004
Page 759

8. 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.

9. In the **Time Zone** field, select your time zone from the drop-down list.
10. 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**.

11. 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.

      Note that shared object rules have a generator ID of 3 (for example, 3:1902) and standard text rules have a generator ID of 1 (for example, 1:1000062).
12. Click **Save**.

    The system saves your email alerting configuration. When applicable intrusion events occur, you receive email alerts.

Keysight, EX-1004
Page 760

Chapter 20
Understanding and Writing Intrusion Rules

An *intrusion rule* is a specified set of keywords and arguments on a 3D Sensor with the IPS component that detects attempts to exploit vulnerabilities on your network by analyzing network traffic to check if it matches the criteria in the rule. IPS 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 3D Sensor web interface or, in the case of a managed 3D Sensor, 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 can seriously affect the performance of your Sourcefire 3D System.

Note the following:

- For a *drop* rule in an inline deployment, IPS drops the packet and generates an event. For more information on drop rules, see Setting Rule States on page 435.
- Sourcefire provides two types of intrusion rules: shared object rules and standard text rules. The Sourcefire Vulnerability Research Team 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 are creating a standard text rule.

Keysight, EX-1004
Page 761

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 267.

See the following sections for more information.

- Understanding Rule Anatomy on page 762 describes the components, including the rule header and rule options, that make up a valid standard text rule.
- Understanding Rule Headers on page 764 provides a detailed description of the parts of a rule header.
- Understanding Keywords and Arguments in Rules on page 769 explains the usage and syntax of the intrusion rule keywords available in the Sourcefire 3D System.
- Constructing a Rule on page 861 explains how to build a new rule using the rule editor.
- Searching for Rules on page 869 explains how to search for existing rules.
- Filtering Rules on the Rule Editor Page on page 872 explains how to display a subset of rules to help you find specific rules.

Understanding Rule Anatomy

Platform: DC/MDC/Standalone Sensor

Component: IPS

All standard text rules contain two logical sections: the rule header and the rule options. The rule header contains:

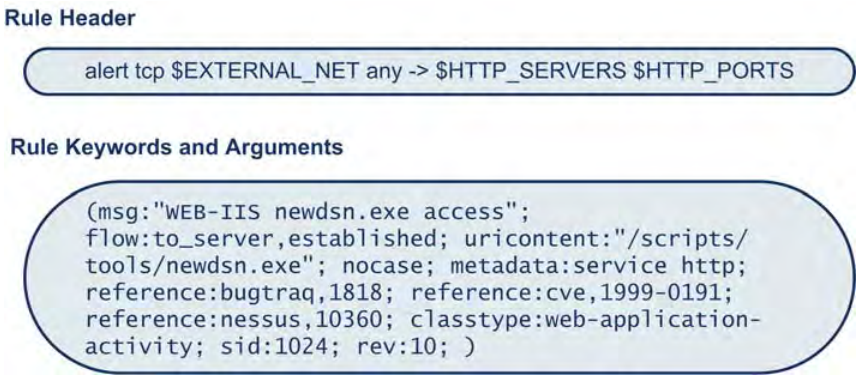
- the rule's action or type
- the protocol
- the source and destination IP addresses and netmasks
- direction indicators showing the flow of traffic from source to destination
- the source and destination ports

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

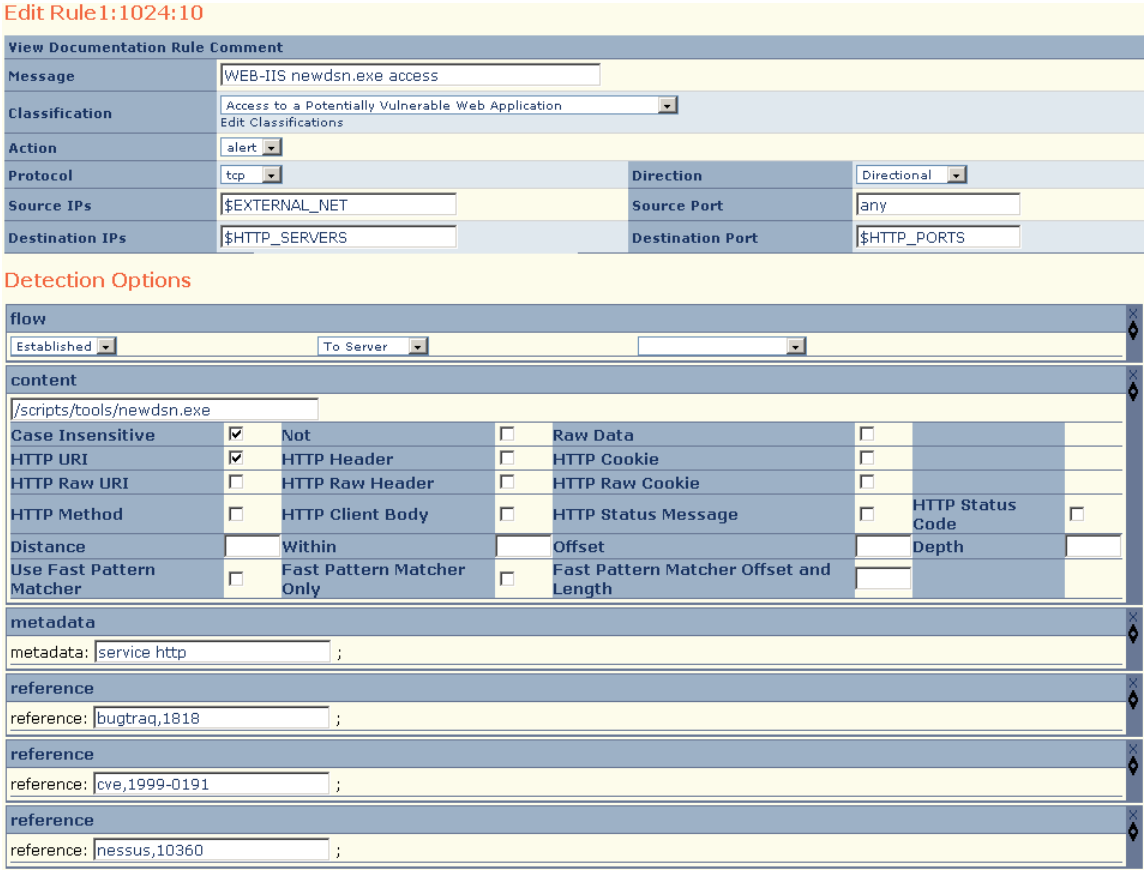
Keysight, EX-1004
Page 762

The following diagram illustrates the parts of a rule:



Note that the options section of a rule is the section enclosed in parentheses. The rule editor provides an easy-to-use interface to help you build standard text rules.

The following displays the same standard text rule within the rule editor:



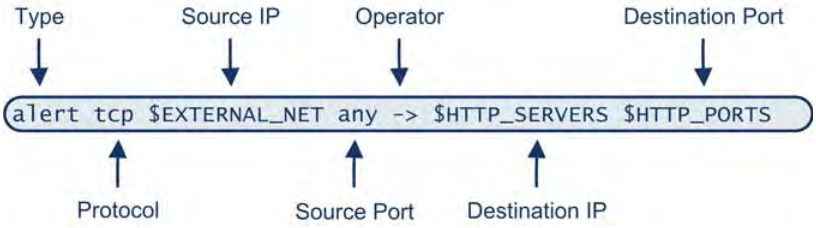
Keysight, EX-1004
Page 763

Understanding Rule Headers

Platform: DC/MDC/Standalone Sensor

Component: IPS

Every standard text rule and shared object rule has parameters and arguments that make up the rule header. For example, the following illustrates the 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 371 for more information about variables, what they mean, and how to configure them.

Keysight, EX-1004
Page 764

See the following sections for more information about rule header parameters:

- Specifying Rule Actions on page 765 describes rule types and explains how to specify the action that occurs when the rule triggers.
- Specifying Protocols on page 766 explains how to define the traffic protocol for traffic that the rule should test.
- Specifying Source and Destination IP Addresses on page 766 explains how to define the individual IP addresses and IP ranges in the rule header.
- Specifying Source and Destination Ports on page 768 explains how to define the individual ports and port ranges in the rule header.
- Specifying Direction on page 769 describes the available operators and explains how to specify the direction traffic must be traveling to be tested by the rule.

Specifying Rule Actions

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 435.

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 861.

Keysight, EX-1004
Page 765

Specifying Protocols

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 768.

- 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 861 for more information about the procedures you use to build a rule header using the rule editor.

Specifying Source and Destination IP Addresses

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 766

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	`192.168.1.1`
a list of IP addresses	brackets (`[]`) to enclose the IP addresses and commas to separate them	`[192.168.1.1,192.168.1.15]`
a range of IP addresses	CIDR block notation	`192.168.1.0/24`
anything except a specific IP address or set of addresses	the `!` character before the IP address or addresses you want to negate **IMPORTANT!** In the same rule, you cannot mix IP addresses that use the `!` character with IP addresses that do not use it. You should create another rule.	`!192.168.1.15`

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 Managing Variables on page 371 and IP Address Conventions on page 59.

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 861 for more information about the procedures you use to build a rule header using the rule editor.

Keysight, EX-1004
Page 767

Specifying Source and Destination Ports

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 766.

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 *any*, which if negated would indicate *no port*.	`!20`

TIP! You can currently substitute a colon (:) for the hyphen (-) in any port definition in an rule header, although eventually the colon will be deprecated.

Keysight, EX-1004
Page 768

You can list ports 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 401.

Within the rule editor, you specify source and destination ports in the **Source Port** and **Destination Port** fields. See Constructing a Rule on page 861 for more information about the procedures you use to build a rule header using the rule editor.

Specifying Direction

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 861 for more information about the procedures you use to build a rule header using the rule editor.

Understanding Keywords and Arguments in Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 combine any number of compatible keywords to create highly specific rules. This helps decrease the

Keysight, EX-1004
Page 769

chance of false positives and false negatives and focus the intrusion information you receive.

IMPORTANT! Requires: DC+IPS+RNA If you want IPS to dynamically adapt active rule processing for specific packets based on rule metadata and RNA host information, you can also enable adaptive profiles. For more information, see Using Adaptive Profiles on page 697.

See the following sections for more information:

- Defining Intrusion Event Details on page 771 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 776 describes how to use the `content` keyword to test the content of the packet payload.
- Constraining Content Matches on page 778 describes how to use modifying keywords for the `content` keyword.
- Replacing Content in Inline Deployments on page 792 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 794 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 800 describes how to use the `pcre` keyword to use Perl-compatible regular expressions in rules.
- Adding Metadata to a Rule on page 810 describes how to use the `metadata` keyword to add information to a rule.
- Inspecting IP Header Values on page 812 describes the syntax and use of keywords that test values in the packet's IP header.
- Inspecting ICMP Header Values on page 816 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 818 describes the syntax and use of keywords that test values in the packet's TCP header.
- Enabling and Disabling TCP Stream Reassembly on page 824 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 825 describes the use and syntax of keywords that extract version and state information from encrypted traffic.

Keysight, EX-1004
Page 770

- Reading Packet Data into Keyword Arguments on page 840 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 828 describes the use and syntax of keywords that test application layer protocol properties.
- Inspecting Packet Characteristics on page 837 describes the use and syntax of the `dsize`, `sameIP`, `isdataat`, `fragoffset`, and `cvs` keywords.
- Initiating Active Responses with Rule Keywords on page 844 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 849 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 850 describes how to log additional traffic for the host or session.
- Detecting Attacks That Span Multiple Packets on page 851 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 855 describes how to generate events on the type of encoding in an HTTP request or response URI, header, or cookie, including set-cookies, prior to normalization.
- Pointing to a Specific Payload Type on page 857 describes how to point to the beginning of the HTTP response entity body, SMTP payload, or Base64-encoded MIME email attachment.
- Decoding and Inspecting Base64 Data on page 859 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 771

See the following sections for more information about event-related keywords:

- Defining the Event Message on page 772
- Defining the Event Priority on page 772
- Defining the Event Classification on page 772
- Defining the Event Reference on page 776

Defining the Event Message

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 define almost any alphanumeric string, up to 63 characters long, as the event 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 861 for more information about using the rule editor to build rules.

Defining the Event Priority

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 861 for more information about using the rule editor to build rules.

Defining the Event Classification

Platform: DC/MDC/Standalone Sensor

Component: IPS

For each rule, you can specify an attack classification that appears in the packet display of the event.

Keysight, EX-1004
Page 772

The following lists the predefined classifications:

- A Client was Using an Unusual Port
- A Network Trojan was Detected
- A Suspicious Filename was Detected
- A Suspicious String was Detected
- A System Call was Detected
- A TCP Connection was Detected
- Access to a Potentially Vulnerable Web Application
- An Attempted Login Using a Suspicious Username was Detected
- Attempted Administrator Privilege Gain
- Attempted Denial of Service
- Attempted Information Leak
- Attempt to Login By a Default Username and Password
- Attempted User Privilege Gain
- Decode of an RPC Query
- Denial of Service
- Detection of a Denial Of Service Attack
- Detection of a Network Scan
- Detection of a Non-Standard Protocol or Event
- Executable Code was Detected
- Generic ICMP Event
- Generic Protocol Command Decode
- Information Leak
- Large Scale Information Leak
- Misc Activity
- Misc Attack
- Not Suspicious Traffic
- Pornography was Detected
- Potential Corporate Privacy Violation
- Potentially Bad Traffic
- Successful Administrator Privilege Gain
- Successful User Privilege Gain
- Unknown Traffic
- Unsuccessful User Privilege Gain
- Web Application Attack

Keysight, EX-1004
Page 773

To specify a classification in the rule editor, select a classification from the **Classification** list. See Writing New Rules on page 861 for more information on the rule editor.

Adding Custom Classifications

Platform: DC/MDC/Standalone Sensor

Component: IPS

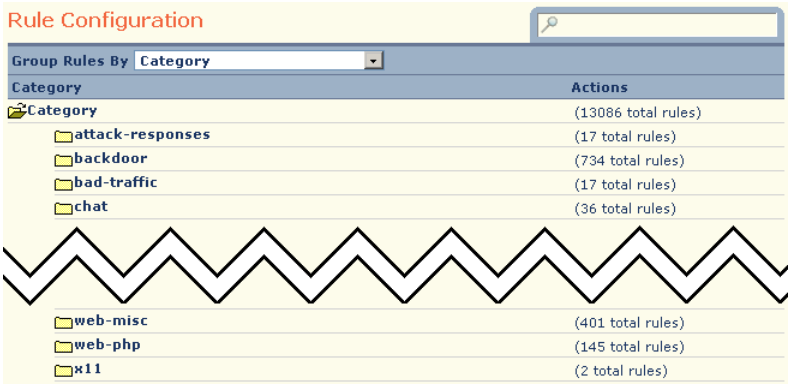
If you want more customized content for the packet display description of the events generated by a rule you define, create a custom classification.

To add classifications to the Classification list:

Access: P&R Admin/Admin

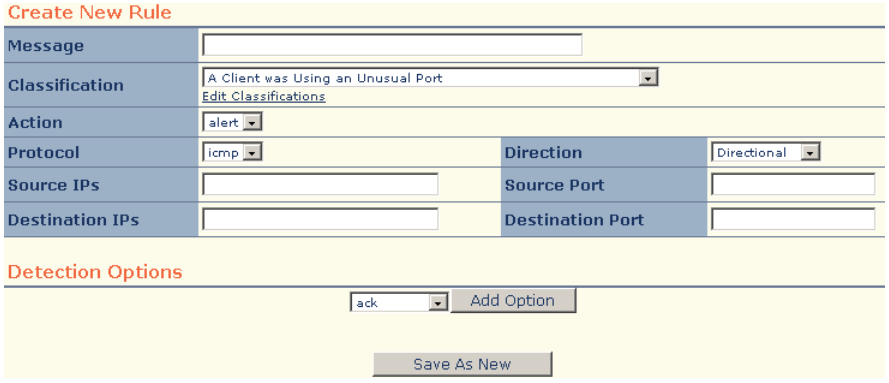
1. Select **Policy & Response > IPS > Rule Editor**.

   The Rule Editor page appears.



2. Click **Create Rule**.

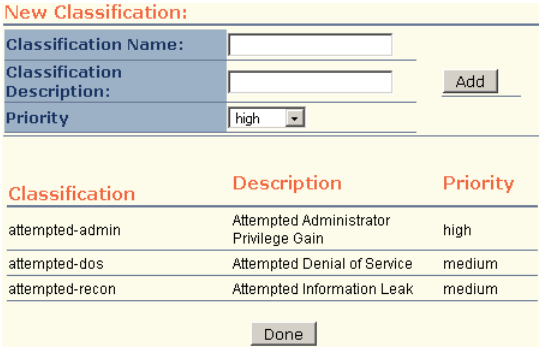
   The new rule Create page appears.



Keysight, EX-1004
Page 774

3. Under the **Classification** drop-down list, click **Edit Classifications**.

   A pop-up window appears.



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.

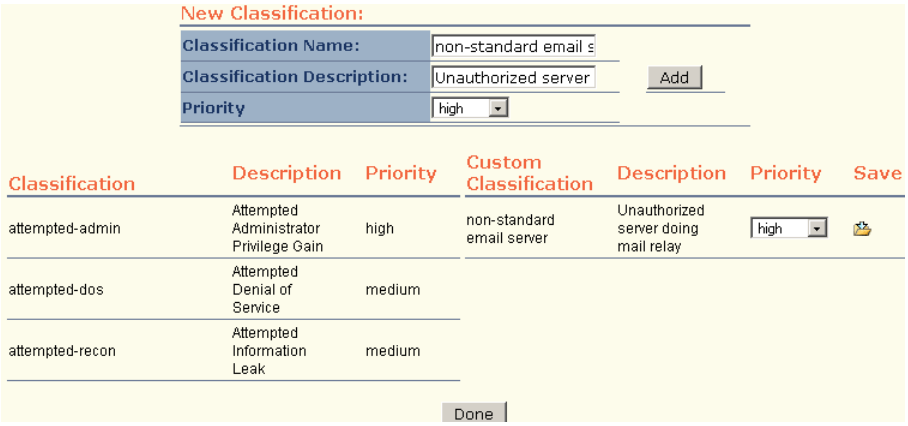
   You can use up to 255 alphanumeric characters and spaces.

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.



8. Click **Done**.

Keysight, EX-1004
Page 775

Defining the Event Reference

Platform: DC/MDC/Standalone Sensor

Component: IPS

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	URL Prefix	Example ID
Bugtraq	`http://www.securityfocus.com/bid/`	`8550`
CVE	`http://cve.mitre.org/cgi-bin/cvename.cgi?name=`	`CAN-2003-0702`
Arachnids	`http://www.whitehats.com/info/IDS`	`563`
McAfee	`http://vil.nai.com/vil/content/v_`	`100561.htm`
url	`http://`	`www.example.com?exploit=14`

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`

See Constructing a Rule on page 861 for more information about using the rule editor to build rules.

Searching for Content Matches

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 776

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.

IMPORTANT! 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.

To enter content to be matched:

Access: P&R Admin/Admin

1. In the `content` field, type the content you want to search for (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 779.

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 769. For more information on constraining the `content` keyword, see Constraining Content Matches on page 778.
3. Continue with creating or editing the rule. See Writing New Rules on page 861 or Modifying Existing Rules on page 864 for more information.

If you are using an inline detection engine, 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 792 for more information.

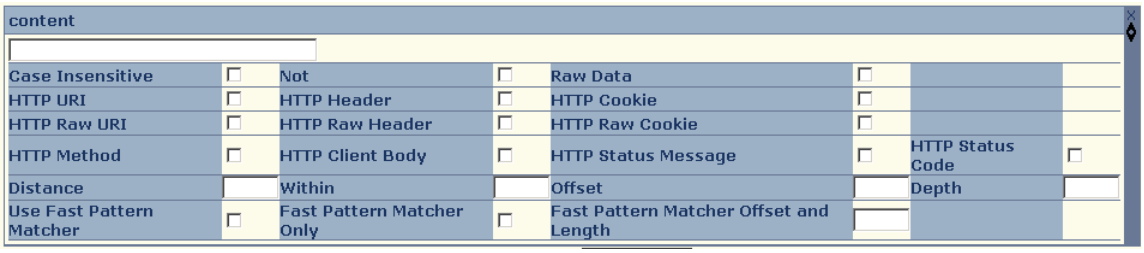
Keysight, EX-1004
Page 777

Constraining Content Matches

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 778
- Raw Data on page 779
- Not on page 779
- Distance on page 781
- Within on page 781
- Offset on page 782
- Depth on page 783
- HTTP Content Options on page 783
- Use Fast Pattern Matcher on page 789

Case Insensitive

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/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 776, Writing New Rules on page 861 or Modifying Existing Rules on page 864 for more information.

Keysight, EX-1004
Page 778

Raw Data

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 prior to 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 783 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 564.

To analyze raw data:

Access: P&R Admin/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 776, Writing New Rules on page 861, or Modifying Existing Rules on page 864 for more information.

Not

Platform: DC/MDC/Standalone Sensor

Component: IPS

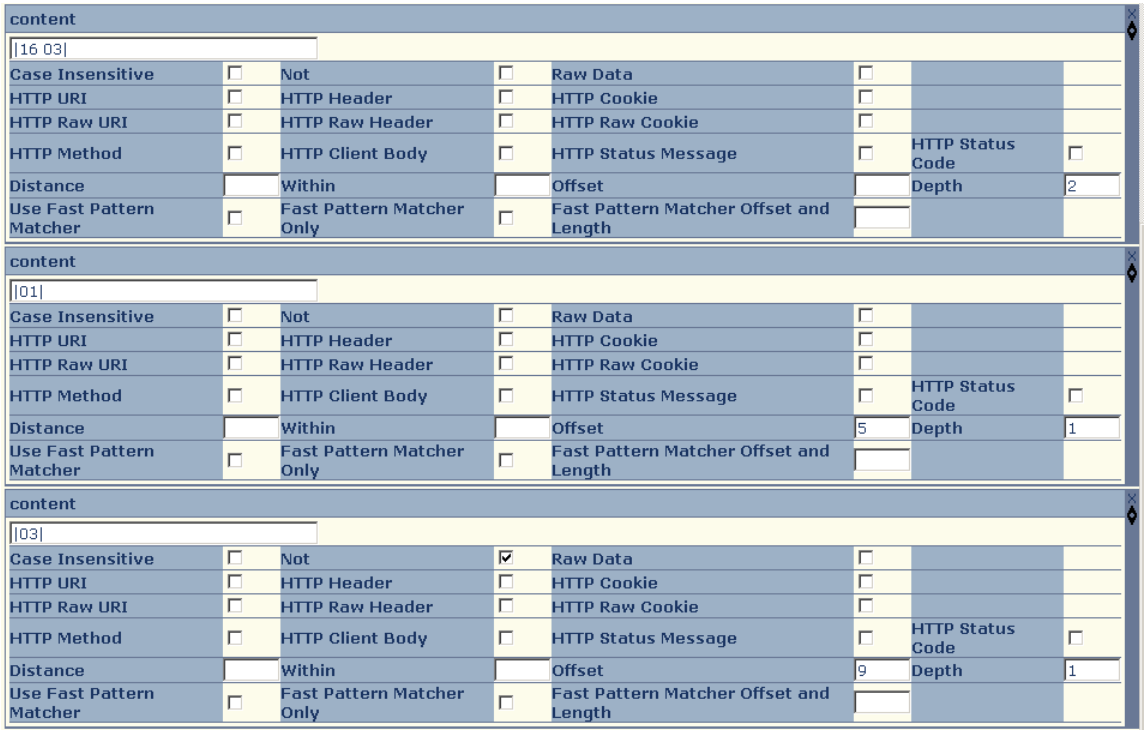
Select the **Not** option to search for content that does not match the specified content. If you create a rule that includes a `content` keyword with the **Not** option selected, you must also include in the rule at least one other `content` keyword without the **Not** option selected.

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 Creating Intrusion Policies on page 331.

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

Keysight, EX-1004
Page 779

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: P&R Admin/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 776, Writing New Rules on page 861, or Modifying Existing Rules on page 864 for more information.

Keysight, EX-1004
Page 780

Distance

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can instruct the rules engine to identify subsequent content matches that occur a specified number of bytes after the previous successful content match. The distance counter starts at byte 0, so calculate the distance value by subtracting 1 from the number of bytes you want to move forward from the last successful content match. For example, if you specify a **Distance** value of `4`, the rules engine starts searching for identical content matches five bytes after the previous content match.

You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.

Note that you should use this option only once in a rule. Note also that you cannot use this option in combination with the **Depth** or **Offset** option.

To specify a distance for a content match:

Access: P&R Admin/Admin

1. Type the distance value in the **Distance** field for the `content` keyword you are adding.

   You can specify a value of -65535 to 65535 bytes.

2. Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches on page 776, Writing New Rules on page 861 or Modifying Existing Rules on page 864 for more information.

Within

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.

Note that you should use this option only once in a rule. Note also that you cannot use this option in combination with the **Depth** or **Offset** option.

Keysight, EX-1004
Page 781

To specify a within value in the user interface:

Access: P&R Admin/Admin

1. Type the within value in the **Within** field for the `content` keyword you are adding.

   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.

2. Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches on page 776, Writing New Rules on page 861 or Modifying Existing Rules on page 864 for more information.

Offset

Platform: DC/MDC/Standalone Sensor

Component: IPS

Use the `offset` keyword to specify where the rules engine should start searching for content within a packet payload, measured in bytes (note that all packet payloads start at byte 0). Because the `offset` counter starts at byte 0, calculate the `offset` value by subtracting 1 from the number of bytes you want to move forward from the beginning of the packet payload. For example, if you added a content match with an `offset` value of 7, the rules engine starts searching for the content at the eighth byte in the packet payload (or byte 7).

Using an offset promotes more efficient searches by constraining the portion of packet payload that is searched and is useful in instances where you know that the matching content will not appear in the first part of the packet. Conversely, you should be sure not to set the value for `offset` too stringently, because the rules engine will not inspect the bytes that appear before the specified `offset` value.

You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.

Note that you should use this option only once in a rule. Note also that you cannot use this option in combination with the **Distance** or **Within** option.

To specify an offset for a content match:

Access: P&R Admin/Admin

1. Type the offset value in the **Offset** field for the `content` keyword you are adding.

   You can specify a value of -65535 to 65535 bytes.

2. Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches on page 776, Writing New Rules on page 861, or Modifying Existing Rules on page 864 for more information.

Keysight, EX-1004
Page 782

Depth

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can specify 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 can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.

Note that you should use this option only once in a rule. Note also that you cannot use this option in combination with the **Distance** or **Within** option.

To specify a depth for a content match:

Access: P&R Admin/Admin

1. Type the depth value in the **Depth** field for the `content` keyword you are adding.

   You can specify 0 or a value that is greater than or equal to the length of the specified content, up to a maximum of 65535 bytes.

2. Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches on page 776, Writing New Rules on page 861, or Modifying Existing Rules on page 864 for more information.

HTTP Content Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 783

Five options search normalized fields in HTTP requests, responses, or both, as appropriate (see the HTTP content Keyword Options table on page 785 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 the HTTP content Keyword Options table on page 785 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).

Keysight, EX-1004
Page 784

- 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 Status Message**, or **HTTP Status Code**

  unless you also select one or more of the following normalized fields:

  HTTP URI, **HTTP Header**, **HTTP Method**, or **HTTP Client Body**

- 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 789 for more information.

The above restrictions are reflected in the description of each option in the HTTP content Keyword Options table on page 785.

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 490.

The HTTP content Keyword Options table describes the HTTP content options.

HTTP content Keyword Options

Select this option...	To search for content matches in...
HTTP URI	the normalized request URI field. Note that you cannot use this option and **Raw Data** or **HTTP Raw URI** in the same `content` keyword. See Raw Data on page 779. Note also 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 806 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.

Keysight, EX-1004
Page 785

HTTP content Keyword Options (Continued)

Select this option...	To search for content matches in...
HTTP Raw URI	the raw request URI field. Note that you cannot use this option and **HTTP URI**, **Raw Data**, or **Use Fast Pattern Matcher** in the same `content` keyword. See Raw Data on page 779 and Use Fast Pattern Matcher on page 789 for more information. Note also that you cannot use this option in combination with the `pcre` keyword HTTP raw URI (I) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table on page 806 for more information. **IMPORTANT!** A pipelined HTTP request packet contains multiple URIs. When **HTTP Raw 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	the request method field, which identifies the action such as GET, PUT, CONNECT, and so on to take on the resource identified in the URI.
HTTP Header	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 and **HTTP Raw Header** or **Raw Data** in the same `content` keyword. See Raw Data on page 779 for more information. Note also 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 806 for more information.
HTTP Raw Header	the raw header field, except for cookies, in HTTP requests; also in responses when the HTTP Inspect preprocessor I**nspect HTTP Responses** option is enabled. Note that you cannot use this option and **HTTP Header**, **Raw Data**, or **Use Fast Pattern Matcher** in the same `content` keyword. See Raw Data on page 779 and Use Fast Pattern Matcher on page 789 for more information. Note also 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 806 for more information.

Keysight, EX-1004
Page 786

HTTP content Keyword Options (Continued)

Select this option...	To search for content matches in...
HTTP Cookie	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 IPS 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 564 for more information. Note that you cannot use this option and **Raw Data**, **HTTP Raw Cookie**, or **Use Fast Pattern Matcher** in the same `content` keyword. See Raw Data on page 779 and Use Fast Pattern Matcher on page 789 for more information. Note also that 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 806 for more information.
HTTP Raw Cookie	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 IPS 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 564 for more information. Note that you cannot use this option and **Raw Data**, **HTTP Cookie**, or **Use Fast Pattern Matcher** in the same `content` keyword. See Raw Data on page 779 and Use Fast Pattern Matcher on page 789 for more information. Note also that 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 806 for more information.

Keysight, EX-1004
Page 787

HTTP content Keyword Options (Continued)

Select this option...	To search for content matches in...
HTTP Client Body	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 564 for more information. Note also that you cannot use this option and **Raw Data** in the same `content` keyword. See Raw Data on page 779 for more information.
HTTP Status Code	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 564 for more information. Note that you cannot use this option and **Raw Data** or **Use Fast Pattern Matcher** in the same `content` keyword. See Raw Data on page 779 and Use Fast Pattern Matcher on page 789 for more information.
HTTP Status Message	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 564 for more information. Note that you cannot use this option and **Raw Data** or **Use Fast Pattern Matcher** in the same `content` keyword. See Raw Data on page 779 and Use Fast Pattern Matcher on page 789 for more information.

Keysight, EX-1004
Page 788

To specify an HTTP content option when doing a content search of TCP traffic:

Access: P&R Admin/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 778, Searching for Content Matches on page 776, Writing New Rules on page 861, or Modifying Existing Rules on page 864 for more information.

Use Fast Pattern Matcher

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS uses the fast pattern matcher to quickly determine 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.

Keysight, EX-1004
Page 789

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**, H**TTP Raw Cookie**, **HTTP Cookie**, **HTTP Status Message**, or **HTTP Status Code**

  unless you also select one or more of the following normalized fields:

  HTTP URI, **HTTP Header**, **HTTP Method**, or **HTTP Client Body**

  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 783 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 F**ast 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 859 for more information.
- You cannot use the fast pattern matcher when using the file_data keyword with its **MIME** argument to point to Base64-encoded MIME data; see Pointing to a Specific Payload Type on page 857 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 re-evaluate 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.

Keysight, EX-1004
Page 790

Note the following conditions when using this option:

- The specified content is location-independent; that is, it can 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.

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: P&R Admin/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.

Keysight, EX-1004
Page 791

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 778, Searching for Content Using PCRE on page 800, Writing New Rules on page 861, or Modifying Existing Rules on page 864 for more information.

Replacing Content in Inline Deployments

Platform: DC/MDC/Standalone Sensor

Component: IPS

When you apply an intrusion policy set to **Drop when Inline** to a detection engine using an inline or inline with failopen interface set, you can use the `replace` keyword 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 IPS 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 re-sends the dropped packet until it bypasses

Keysight, EX-1004
Page 792

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, IPS automatically updates the packet checksums so that the destination host can receive the packet without error.

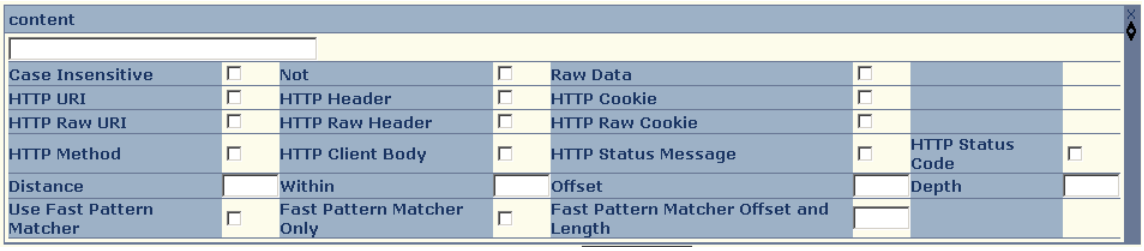
Note that you cannot use the `replace` keyword in combination with HTTP request message `content` keyword options. See Searching for Content Matches on page 776 and HTTP Content Options on page 783 for more information.

To replace content in an inline deployment:

Access: P&R Admin/Admin

1. On the new rule Create 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.

Keysight, EX-1004
Page 793

Using Byte_Jump and Byte_Test

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 514 and DCE/RPC Keywords on page 832 for more information.

See the following sections for more information:

- byte_jump on page 794
- byte_test on page 797

byte_jump

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to jump forward from the beginning of the packet payload or the last successful content match. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.

Keysight, EX-1004
Page 794

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.
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.

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.

Keysight, EX-1004
Page 795

Endianness Arguments (Continued)

Argument	Description
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_jump` keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 514 for more information. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type**, **Endian**, and **From Beginning** arguments do not apply. When you enable this argument, you can also use `byte_jump` in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords on page 832 for more information. 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 490.

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

Keysight, EX-1004
Page 796

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

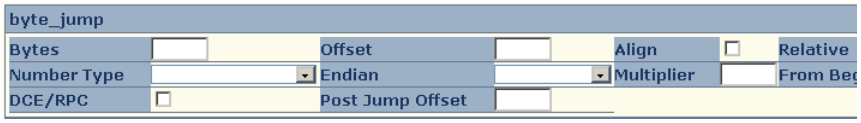
the rules engine calculates the number described in the four bytes that appear 13 bytes after the beginning of the packet. Then, the engine multiplies that number by two to obtain the total number of bytes to skip. For instance, if the four calculated bytes in a specific packet were `00 00 00 1F`, the 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: P&R Admin/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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 797

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 <, >, =, !, &, ^, !>, !<, !=, !&, or !^. For example, if you specify `!1024`, `byte_test` would convert the specified number, and if it did not equal 1024, it would generate an event (if all other keyword parameters matched). Note that ! and != are equivalent. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.
Offset	The number of bytes into the payload to start processing. The `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to count forward from the beginning of the packet payload or the last successful content match. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 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.

Keysight, EX-1004
Page 798

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	Specifies a `byte_test` keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 514 for more information. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** argument do not apply. When you enable this argument, you can also use `byte_test` in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords on page 832 for more information. 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 490.

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.

Keysight, EX-1004
Page 799

For example, if the value for `byte_test` is specified as the following:

- Bytes = 4
- Operator and Value > 128
- Offset = 8
- Relative enabled

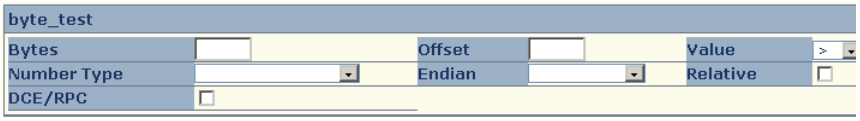
the rules engine calculates the number described in the four bytes that appear 9 bytes away from (relative to) the last successful content match, and, if the calculated number is larger than 128 bytes, the rule is triggered.

To use byte_test:

Access: P&R Admin/Admin

1. On the new rule Create page, select `byte_test` in the drop-down list and click **Add Option**.

   The `byte_test` section appears beneath the last keyword you selected.



Searching for Content Using PCRE

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `pcre` keyword allows you to use Perl-compatible regular expressions (PCRE) to inspect packet payloads for specified content. You can use PCRE to avoid writing multiple rules to match slight variations of the same content.

Regular expressions are useful when searching for content that could be displayed in a variety of ways. The content may have different attributes that you want to account for in your attempt to locate it within a packet's payload.

Note that the regular expression syntax used in intrusion rules is a subset of the full regular expression library and varies in some ways from the syntax used in commands in the full library. When adding a `pcre` keyword using the rule editor, enter the full value in the following format:

```
!/pcre/ ismxAEGRBUIPHDMCKSY
```

where:

- ! is an optional negation (use this if you want to match patterns that **do not** match the regular expression).
- */pcre/* is a Perl-compatible regular expression.
- `ismxAEGRBUIPHDMCKSY` is any combination of modifier options.

Keysight, EX-1004
Page 800

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*`? ismxAEGRBUIPHDMCKSY` where *regex* is your Perl-compatible regular expression and `ismxAEGRBUIPHDMCKSY` is any combination of modifier options. See Perl-Compatible Regular Expression Basics on page 802 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 802 describes the common syntax used in Perl-compatible regular expressions.
- PCRE Modifier Options on page 804 describes the options you can use to modify your regular expression.
- Example PCRE Keyword Values on page 808 gives example usage of the `pcre` keyword in rules.

Keysight, EX-1004
Page 801

Perl-Compatible Regular Expression Basics

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 `s` is used as a modifying option, it also includes newline characters.	`abc.` matches `abcd`, `abc1`, `abc#`, and so on.
`*`	Matches zero or more occurrences of a character or expression.	`abc*` matches `abc`, `abcc`, `abccc`, `abccccc`, and so on.
`?`	Matches zero or one occurrence of a character or expression.	`abc?` matches `abc`.
`+`	Matches one or more occurrences of a character or expression.	`abc+` matches `abc`, `abcc`, `abccc`, and so on.
`()`	Groups expressions.	`(abc)+` matches `abc`, `abcabc`, `abcabcabc` and so on.
`{}`	Specifies a limit for the number of matches for a character or expression. If you want to set a lower and upper limit, separate the lower limit and upper limit with a comma.	`a{4,6}` matches `aaaa`, `aaaaa`, or `aaaaaa`. `(ab){2}` matches `abab`.

Keysight, EX-1004
Page 802

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 803 for more information about using character classes in PCRE.

Character Classes

Platform: DC/MDC/Standalone Sensor

Component: IPS

Character classes include alphabetic characters, numeric characters, alphanumeric characters, and whitespace characters. While you can create your own character classes within brackets (see Metacharacters on page 802), 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.

Keysight, EX-1004
Page 803

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]
\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 whitespace characters, including spaces, carriage returns, tabs, newlines, and form feeds.	[\r\t\n\f]
\S	Matches anything that is not a whitespace character.	[^ \r\t\n\f]

PCRE Modifier Options

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use modifying options after you specify regular expression syntax in the `pcre` keyword’s value. These modifiers perform Perl, PCRE, and Snort-specific processing functions. Modifiers always appear at the end of the PCRE value, and appear in the following format:

`/pcre/ismxAEGRBUIPHDMCKSY`

where `ismxAEGRBUPHMC` can include any of the modifying options that appear in the following tables.

TIP! Optionally, you can surround the regular expression and any modifying options with quotes, for example, `“/pcre/ismxAEGRBUIPHDMCKSY”`. The option of using quotes accommodates experienced users accustomed to previous versions when quotes were required instead of optional. The rule editor does not display quotation marks when you display a rule after saving it.

Keysight, EX-1004
Page 804

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 whitespace 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.

Keysight, EX-1004
Page 805

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 490.

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).

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 the HTTP content Keyword Options table on page 785 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 the HTTP content Keyword Options table on page 785 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 the HTTP content Keyword Options table on page 785 for more information.

Keysight, EX-1004
Page 806

Snort-Specific Post Regular Expression Modifiers (Continued)

Option	Description
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 the HTTP content Keyword Options table on page 785 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 the HTTP content Keyword Options table on page 785 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 the HTTP content Keyword Options table on page 785 for more information.
C	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the normalized content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-cookie data. Note that cookies included in the message body are treated as body content. Note also that you cannot use this option in combination with the `content` keyword **HTTP Cookie** option to search the same content. See the HTTP content Keyword Options table on page 785 for more information.
K	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the raw content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-–cookie data. Note that cookies included in the message body are treated as body content. Note also that you cannot use this option in combination with the `content` keyword **HTTP Raw Cookie** option to search the same content. See the HTTP content Keyword Options table on page 785 for more information.
S	Searches the 3-digit status code in an HTTP response. See the `content` keyword **HTTP Status Code** option in the HTTP content Keyword Options table on page 785 for more information.

Keysight, EX-1004
Page 807

Snort-Specific Post Regular Expression Modifiers (Continued)

Option	Description
Y	Searches the textual description that accompanies the status code in an HTTP response. See the `content` keyword **HTTP Status Message** option in the HTTP content Keyword Options table on page 785 for more information.

Example PCRE Keyword Values

Platform: DC/MDC/Standalone Sensor

Component: IPS

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`

Keysight, EX-1004
Page 808

- **`/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.

Keysight, EX-1004
Page 809

Adding Metadata to a Rule

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

IPS 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

Note the following character restrictions:

- Do not use a semi-colon (;) or colon (:) in a `metadata` keyword.
- Be aware when using commas that IPS 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 IPS interprets these characters as separators between *key* and *value*. For example:

  ```
  key value
  key=value
  ```


Keysight, EX-1004
Page 810

All other characters are permitted.

Avoiding Reserved Metadata

Platform: DC/MDC/Standalone Sensor

Component: IPS

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
service
soid
```

IMPORTANT! Contact Sourcefire Support for assistance in adding restricted metadata to local rules that might not otherwise function as expected. For example, a pass rule for passing SNMP traffic might not be effective if adaptive profiles are enabled and the rule does not include service metadata. See Using Adaptive Profiles on page 697 and Importing Local Rule Files on page 1857 for more information.

Searching for Rules with Metadata

Platform: DC/MDC/Standalone Sensor

Component: IPS

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`*.

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.

Keysight, EX-1004
Page 811

Note that regardless of the format you use to add metadata, IPS 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, IPS 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 IPS would interpret as a single *key value* statement:

```
ab ef
```

For more information, see Searching for Rules on page 869.

Setting a Red Impact Flag

Platform: DC/MDC/Standalone Sensor

Component: IPS

On a Defense Center with RNA and IPS deployed on the same network segment, 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 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 307 a for more information.

Inspecting IP Header Values

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 813
- Inspecting the IP Header Identification Value on page 814
- Identifying Specified IP Options on page 814

Keysight, EX-1004
Page 812

- Identifying Specified IP Protocol Numbers on page 815
- Inspecting a Packet's Type of Service on page 815
- Inspecting a Packet's Time-To-Live Value on page 815

Inspecting Fragments and Reserved Bits

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 on all specified bits.
asterisk (*****)	The packet can match on 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.

Keysight, EX-1004
Page 813

Inspecting the IP Header Identification Value

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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
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.

Keysight, EX-1004
Page 814

Identifying Specified IP Protocol Numbers

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

A packet's time-to-live (ttl) value indicates how many hops it can make before it is dropped. You can use the `ttl` keyword to test the packet's IP header ttl value against the value, or range of values, you specify as the keyword's argument. It may be helpful to set the `ttl` keyword parameter to a low value such as 0 or 1, as low time-to-live values are sometimes indicative of a traceroute or intrusion evasion attempt. (Note, though, that the appropriate value for this keyword

Keysight, EX-1004
Page 815

depends on your sensor 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 816
- Inspecting the ICMP Message Type on page 817
- Inspecting the ICMP Message Code on page 817

Identifying Static ICMP ID and Sequence Values

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 816

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

Use the `itype` keyword to look for packets with specific ICMP message type values. You can specify either a valid ICMP type value (see http://www.iana.org/assignments/icmp-parameters or http://www.faqs.org/rfcs/rfc792.html for a full list of ICMP type numbers) or an invalid ICMP type value to test for different types of traffic. For example, attackers may set ICMP type values out of range to cause denial of service and flooding attacks.

You can specify a range for the `itype` argument value using less than (<) and greater than (>).

For example:

- <35
- >36
- 3<>55

TIP! See http://www.iana.org/assignments/icmp-parameters or http://www.faqs.org/rfcs/rfc792.html for a full list of ICMP type numbers.

Inspecting the ICMP Message Code

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.)

Keysight, EX-1004
Page 817

You can use the `icode` keyword to identify packets with specific ICMP code values. You can choose to specify either a valid ICMP code value or an invalid ICMP code value to test for different types of traffic.

You can specify a range for the `icode` argument value using less than (<) and greater than (>).

For example:

- to find values less than 35, specify `<35.`
- to find values greater than 36, specify `>36.`
- to find values between 3 and 55, specify `3<>55.`

TIP! You can use the `icode` and `itype` keywords together to identify traffic that matches both. For example, to identify ICMP traffic that contains an ICMP Destination Unreachable code type with an ICMP Port Unreachable code type, specify an `itype` keyword with a value of 3 (for Destination Unreachable) and an `icode` keyword with a value of 3 (for Port Unreachable).

Inspecting TCP Header Values and Stream Size

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 818
- Inspecting TCP Flag Combinations on page 819
- Applying Rules to a TCP or UDP Client or Server Flow on page 820
- Identifying Static TCP Sequence Numbers on page 822
- Identifying TCP Windows of a Given Size on page 822
- Identifying TCP Streams of a Given Size on page 823

Inspecting the TCP Acknowledgement Value

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 818

Inspecting TCP Flag Combinations

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 820 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.
R1	An ECN congestion window has been reduced.
R2	ECN echo.

TIP! For more information on Explicit Congestion Notification (ECN), see the information provided at: http://www.faqs.org/rfcs/rfc3168.html.

Keysight, EX-1004
Page 819

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 `Urg` and `all` to specify that a packet must contain the Urgent flag and may contain any other flags.
any	The packet can contain any of the specified flags.	Select `Ack`, `Psh`, and `any` to specify that either or both the `Ack` and `Psh` flags must be set to trigger the rule, and that other flags may also be set on a packet.
not	The packet must **not** contain the specified flag set.	Select `Urg` and `not` to specify that the Urgent flag is **not** set on packets that trigger this rule.

Applying Rules to a TCP or UDP Client or Server Flow

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 820

To specify flow, select the `flow` keyword from the **Detection Options** list on the new rule Create 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 `flow` 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 `flow` 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`.

Keysight, EX-1004
Page 821

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, you must enable TCP or UDP stream preprocessing 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 627 and Reassembling TCP Streams on page 637 for information about using TCP stream preprocessing. See Using UDP Stream Preprocessing on page 645 for information about using UDP stream preprocessing. See Automatically Enabling Advanced Settings on page 490 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 822

helpful in identifying attacks and network scans that use generated packets with static TCP window sizes.

Identifying TCP Streams of a Given Size

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 (,).

You must enable TCP stream preprocessing to use the `stream_size` keyword in a rule. See Using TCP Stream Preprocessing on page 627 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 490 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 `both, >, 200` would trigger when traffic from the client is greater than 200 bytes AND traffic from the server is greater than 200 bytes.
either	triggers on traffic from either the client or the server matching the specified stream size, whichever occurs first. For example, the argument `either, >, 200` would trigger when traffic from the client is greater than 200 bytes OR traffic from the server is greater than 200 bytes.

Keysight, EX-1004
Page 823

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 824

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 490.

To use stream_reassemble:

- On the new rule Create page, select `stream_reassemble` in the drop-down list and click **Add Option**.

  The `stream_reassemble` section appears.



Extracting SSL Information from a Session

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

IMPORTANT! The `ssl_state` and `ssl_version` keywords can only invoke the SSL preprocessor on sensors where the SSL preprocessor is enabled in the policy applied to the sensor.

For more information, see the following sections:

- ssl_state on page 826
- ssl_version on page 827

Keysight, EX-1004
Page 825

ssl_state

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `ssl_state` keyword can be used to match against state information for an encrypted session. 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, IPS evaluates them using the OR operator. For example, if you specify `client_hello` and `server_hello` as arguments, IPS 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 490.

The `ssl_state` keyword takes the following identifiers as arguments:

`ssl_state` **Arguments**

Argument	Purpose
`client_hello`	Matches against a handshake message with `ClientHello` as the message type, where the client requests an encrypted session.
`server_hello`	Matches against a handshake message with `ServerHello` as the message type, where the server responds to the client's request for an encrypted session.
`client_keyx`	Matches against a handshake message with `ClientKeyExchange` as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.

Keysight, EX-1004
Page 826

`ssl_state` **Arguments (Continued)**

Argument	Purpose
`server_keyx`	Matches against a handshake message with `ServerKeyExchange` as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
`unknown`	Matches against any handshake message type.

ssl_version

Platform: DC/MDC/Standalone Sensor

Component: IPS

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, IPS 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 490.

The `ssl_version` keyword takes the following SSL/TLS version identifiers as arguments:

`ssl_version` **Arguments**

Argument	Purpose
`sslv2`	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 2.
`sslv3`	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 3.

Keysight, EX-1004
Page 827

`ssl_version` **Arguments (Continued)**

Argument	Purpose
`tls1.0`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.0.
`tls1.1`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.1.
`tls1.2`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.2.

Inspecting Application Layer Protocol Values

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 828
- ASN.1 on page 829
- urilen on page 830
- DCE/RPC Keywords on page 832

RPC

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 828

rpc Keyword Arguments (Continued)

Argument	Description
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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 829

asn.1 Keyword Arguments (Continued)

Argument	Description
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 or `byte_test`/`byte_jump`. To decode an ASN.1 sequence right after the content "foo", set Relative Offset to 0, and do not set an Absolute Offset. Relative Offset may be positive or negative. (Remember that the offset counter starts at 0.)

For example, there is a known vulnerability in the Microsoft ASN.1 Library that creates a buffer overflow, allowing an attacker to exploit the condition with a specially crafted authentication packet. When the system decodes the asn.1 data, exploit code in the packet could execute on the host with system-level privileges or could cause a DoS condition. The following rule uses the `asn1` keyword to detect attempts to exploit this vulnerability:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 445
(flow:to_server, established; content:"|FF|SMB|73|"; nocase;
offset:4; depth:5;
asn1:bitstring_overflow,double_overflow,oversize_length
100,relative_offset 54;)
```

The above rule generates an event against TCP traffic traveling from any IP address defined in the $EXTERNAL_NET variable, from any port, to any IP address defined in the $HOME_NET variable using port 445. In addition, it only executes the rule on established TCP connections to servers. The rule then tests for specific content in specific locations. Finally, the rule uses the `asn1` keyword to detect bitstring encodings and double ASCII encodings and to identify asn.1 type lengths over 100 bytes in length starting 55 bytes from the end of the last successful content match. (Remember that the `offset` counter starts at byte 0.)

urilen

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 830

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 627 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 490.
- The rule protocol is always TCP. See Specifying Protocols on page 766 for more information.
- Target ports are always HTTP ports. See Specifying Source and Destination Ports on page 768 and Understanding Existing Variables on page 372 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;)
```


Keysight, EX-1004
Page 831

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The three DCE/RPC keywords described in the DCE/RPC Keywords table allow you to monitor DCE/RPC session traffic for exploits. When IPS processes rules with these keywords, it invokes the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 514 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 490.

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 794 and Reading Packet Data into Keyword Arguments on page 840.

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.

Keysight, EX-1004
Page 832

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 776 and Use Fast Pattern Matcher on page 789 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 833
- dce_opnum on page 835
- dce_stub_data on page 836

dce_iface

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use the `dce_iface` keyword to identify a specific DCE/RPC service.



Optionally, you can also use `dce_iface` in combination with the `dce_opnum` and `dce_stub_data` keywords to further limit the DCE/RPC traffic to inspect. See dce_opnum on page 835 and dce_stub_data on page 836 for more information.

Keysight, EX-1004
Page 833

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 490.

A fixed, sixteen-byte Universally Unique Identifier (UUID) identifies the application interface assigned to each DCE/RPC service. For example, the UUID 4b324fc8-670-01d3-1278-5a47bf6ee188 identifies the DCE/RPC lanmanserver service, also known as the srvsvc service, which provides numerous management functions for sharing peer-to-peer printers, files, and SMB named pipes. The DCE/RPC preprocessor uses the UUID and associated header values to track DCE/RPC sessions.

The interface UUID is comprised of five hexadecimal strings separated by hyphens:

```
<4hexbytes>-<2hexbytes>-<2hexbytes>-<2hexbytes>-<6hexbytes>
```

You specify the interface by entering the entire UUID including hyphens, as seen in the following UUID for the netlogon interface:

```
12345678-1234-abcd-ef00-01234567cffb
```

Note that you must specify the first three strings in the UUID in big endian byte order. Although published interface listings and protocol analyzers typically display UUIDs in the correct byte order, you might encounter a need to rearrange the UUID byte order before entering it. Consider the following messenger service UUID shown as it might sometimes be displayed in raw ASCII text with the first three strings in little endian byte order:

```
f8 91 7b 5a 00 ff d0 11 a9 b2 00 c0 4f b6 e6 fc
```

You would specify the same UUID for the `dce_iface` keyword by inserting hyphens and putting the first three strings in big endian byte order as follows:

```
5a7b91f8-ff00-11d0-a9b2-00c04fb6e6fc
```

Although a DCE/RPC session can include requests to multiple interfaces, you should include only one `dce_iface` keyword in a rule. Create additional rules to detect additional interfaces.

DCE/RPC application interfaces also have interface version numbers. You can optionally specify an interface version with an operator indicating that the version equals, does not equal, is less than, or greater than the specified value.

Both connection-oriented and connectionless DCE/RPC can be fragmented in addition to any TCP segmentation or IP fragmentation. Typically, it is not useful to associate any DCE/RPC fragment other than the first with the specified interface, and doing so can result in a large number of false positives. However, for flexibility you can optionally evaluate all fragments against the specified interface.

Keysight, EX-1004
Page 834

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.
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 on 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 can result in false positives.

dce_opnum

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 490.

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.

Keysight, EX-1004
Page 835

You should precede a `dce_opnum` keyword with a `dce_iface` keyword to identify the application service for the operation, as shown in the following example. See dce_iface on page 833 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

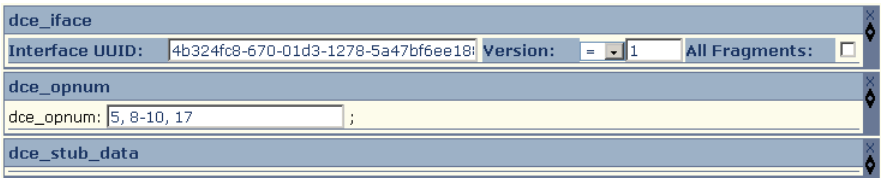
Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

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 490.

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.



Keysight, EX-1004
Page 836

The `dce_stub_data` keyword has no arguments. See dce_iface on page 833 and dce_opnum on page 835 for more information.

Inspecting Packet Characteristics

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 837
- isdataat on page 837
- sameip on page 839
- fragoffset on page 839
- cvs on page 839

dsize

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `isdataat` keyword instructs the rules engine to verify that data resides at a specific location in the payload.

Keysight, EX-1004
Page 837

The isdataat Arguments table lists the arguments you can use with the `isdataat` keyword.

isdataat Arguments

Argument	Type	Description
Offset	Required	The specific location in the payload. For example, to test that data appears at byte 50 in the packet payload, you would specify `50` as the offset value. A `!` modifier negates the results of the `isdataat` test; it alerts if a certain amount of data is not present within the payload. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments on page 840 for more information.
Relative	Optional	Makes the location relative to the last successful content match. If you specify a relative location, note that the counter starts at byte 0, so calculate the location by subtracting 1 from the number of bytes you want to move forward from the last successful content match. For example, to specify that the data must appear at the ninth byte after the last successful content match, you would specify a relative offset of `8`.
Raw Data	Optional	Specifies that the data is located in the original packet payload before decoding or application layer normalization by any Sourcefire 3D System preprocessor. You can use this argument with **Relative** if the previous content match was in the raw packet data.

For example, in a rule searching for the content `foo`, if the value for `isdataat` is specified as the following:

- Offset = !10
- Relative = enabled

IPS alerts if the rules engine does not detect 10 bytes after `foo` before the payload ends.

To use isdataat:

1. On the new rule Create page, select `isdataat` in the drop-down list and click **Add Option**.

   The `isdataat` section appears.



Keysight, EX-1004
Page 838

sameip

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `sameip` keyword tests that a packet's source and destination IP addresses are the same. It does not take an argument.

fragoffset

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `cvs` keyword tests Concurrent Versions System (CVS) traffic for malformed CVS entries. An attacker can use a malformed entry to force a heap overflow and execute malicious code on the CVS server. This keyword can be used to identify attacks against two known CVS vulnerabilities: CVE-2004-0396 (CVS 1.11.x up to 1.11.15, and 1.12.x up to 1.12.7) and CVS-2004-0414 (CVS 1.12.x through 1.12.8, and 1.11.x through 1.11.16). The `cvs` keyword checks for a well-formed entry, and generates alerts when a malformed entry is detected.

Your rule should include the ports where CVS runs. In addition, any ports where traffic may occur should be added to the list of ports for stream reassembly in

Keysight, EX-1004
Page 839

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` service (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 638.

To detect malformed CVS entries:

Access: P&R Admin/Admin

- Add the `cvs` option to a rule and type `invalid-entry` as the keyword argument.

Reading Packet Data into Keyword Arguments

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 840

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 841 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**, IPS 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 841 for more information.

Keysight, EX-1004
Page 841

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	Specifies a `byte_extract` keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic on page 514 for more information. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** arguments do not apply. When you enable this argument, you can also use `byte_extract` in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords on page 832 for more information. 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 490.

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.

Keysight, EX-1004
Page 842

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

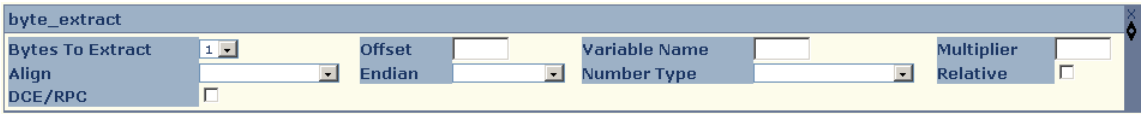
Keyword	Argument
content	Depth, Offset, Distance, Within See Constraining Content Matches on page 778 for more information.
byte_jump	Offset See byte_jump on page 794 for more information.
byte_test	Offset, Value See byte_test on page 797 for more information.
isdataat	Offset See isdataat on page 837 for more information.

To use byte_extract:

Access: P&R Admin/Admin

1. On the new rule Create page, select t `byte_extract` in the drop-down list and click **Add Option**.

   The `byte_extract` section appears beneath the last keyword you selected.



Keysight, EX-1004
Page 843

Initiating Active Responses with Rule Keywords

Platform: DC/MDC/Standalone Sensor

Component: IPS

IPS 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, IPS 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, IPS 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 IPS 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, IPS does not allow TCP resets to initiate TCP resets; this prevents an unending sequence of active responses. IPS 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 629 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 478 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 845
- Sending an HTML Page Before a TCP Reset on page 846
- Setting the Active Response Reset Attempts and Interface on page 847

Keysight, EX-1004
Page 844

Initiating Active Responses by Type and Direction

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 766 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.
reset_both	Directs a TCP reset packet to both the sending and receiving endpoints. Alternatively, you can specify `rst_all`, which is supported for backward compatibility.
icmp_net	Directs an ICMP network unreachable message to the sender.
icmp_host	Directs an ICMP host unreachable message to the sender.

Keysight, EX-1004
Page 845

resp Arguments (Continued)

Argument	Description
icmp_port	Directs an ICMP port unreachable message to the sender. This argument is used to terminate UDP traffic.
icmp_all	Directs the following ICMP messages to the sender: • network unreachable • host unreachable • port unreachable

For example, to configure a rule to reset both sides of a connection when a rule is triggered, use `reset_both` as the value for the `resp` keyword.

You can use a comma-separated list to specify multiple arguments as follows:

`argument``,`*`argument`*`,`*`argument`*

See Setting the Active Response Reset Attempts and Interface on page 847 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: P&R Admin/Admin

1. On the new rule Create 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 845 in the **resp** field; use a comma-separated list to specify multiple arguments.

Sending an HTML Page Before a TCP Reset

Platform: DC/MDC/Standalone Sensor

Component: IPS

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, IPS 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`

Keysight, EX-1004
Page 846

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 762 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 847 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: P&R Admin/Admin

1. On the new rule Create page, select **react** in the drop-down list and click **Add Option.**

   The `react` keyword appears.

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 629 for more information.

Keysight, EX-1004
Page 847

You use the **config response** command by inserting it on a separate line in a custom detection-engine-specific variable with the reserved name USER_CONF. See Understanding Custom Variables on page 391 and Configuring Custom Variables on page 392 for information on using a detection-engine-specific USER_CONF variable.

WARNING! **Do not** use the reserved variable `USER_CONF` to configure an IPS feature unless you are instructed to do so in the feature description or by Sourcefire Support. Conflicting or duplicate configurations will halt IPS.

To specify active response reset attempts, the active response interface, or both:

Access: P&R Admin/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 detection-engine-specific 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, IPS inserts reset packets directly into the stream in place of triggering packets. IPS sends only 1 ICMP reachable active response.

    dev is an alternate interface where you want IPS to send active responses in a passive deployment or insert active responses in an inline deployment.

  Remember to apply your intrusion policy to the detection engine where you configure the detection-engine-specific USER_CONF variable. See Applying an Intrusion Policy on page 347 for more information.

Keysight, EX-1004
Page 848

Filtering Events

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 IPS 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 IPS tracks event instances.

detection_filter Track Arguments

Argument	Description
by_src	Detection criteria count by source IP address.
by_dst	Detection criteria count by destination IP address.

The `count` argument specifies the number of packets that must trigger the rule for the specified IP address within the specified time before the rule generates an event.

The `seconds` argument specifies the number of seconds within which the specified number of packets must trigger the rule before the rule generates an event.

Consider the case of a rule that searches packets for the content `foo` and uses the `detection_filter` keyword with the following arguments:

```
track by_src, count 10, seconds 20
```

In the example, the rule will not generate an event until it has detected `foo` in 10 packets within 20 seconds from a given source IP address. If IPS 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.

Keysight, EX-1004
Page 849

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. See Configuring Event Thresholding on page 439, Configuring Suppression Per Intrusion Policy on page 447, and Setting a Dynamic Rule State on page 454 for more information.

Evaluating Post-Attack Traffic

Platform: DC/MDC/Standalone Sensor

Component: IPS

Use the `tag` keyword to tell IPS 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`).

Keysight, EX-1004
Page 850

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 can 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

Use the `flowbits` keyword to assign state names to sessions. By analyzing subsequent packets in a session according to the prior state, IPS 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

Keysight, EX-1004
Page 851

content to help distinguish malicious packets from those you do not want to alert on. 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 853 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 mutually exclusive group of states. When a rule triggers and sets a state in the group, IPS unsets any other currently set state in the same group. This prevents unrelated rules from triggering false positives. Note that a state can belong to only one group. 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 Using optional_group_name on page 854 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.

To specify an argument and its related value for flowbits without specifying a group name, use the following syntax:

```
argument,state_name
```

where *`argument`* is one of the arguments listed in the table below, and *`state_name`* is the state name you want to provide to the argument.

The *`set`* and *`toggle`* arguments accept an optional, additional group name and use the following syntax:

```
argument,state_name,optional_group_name
```

where *`optional_group_name`* specifies a group of mutually exclusive states.

Note that 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 IPS to automatically enable the required TCP or UDP stream preprocessing. See Automatically Enabling Advanced Settings on page 490 for more information.

Keysight, EX-1004
Page 852

You can define actions for specified state names by using `flowbits` with the arguments listed in the flowbits Arguments table.

flowbits Arguments

Argument	Description
set,*state_name*	Sets the specified state for a particular flow, and unsets all other flowbit states in the same group when the group name is specified.
unset,*state_name*	Unsets the defined state for a particular flow.
toggle,*state_name*	Sets the specified state if it is currently unset, and unsets all other flowbit states in the same group when the group name is specified; unsets that state if it is currently set.
isset,*state_name*	Determines whether the specified state is set for the flow.
isnotset,*state_name*	Determines if the specified state is not set for the flow.
noalert	Indicates that regardless of the rest of the detection options specified (such as setting or unsetting), no events are generated.
reset	Resets all defined states for a given session.

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;)
```


Keysight, EX-1004
Page 853

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 an 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;)
```

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 Using optional_group_name

Including different state names that are set in different rules in a mutually exclusive 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 examples illustrate first how you can get false positives when you do not include multiple state names in a mutually exclusive group and, then, how including state names in a group can prevent false positives.

Consider first the case where the following three rule fragments trigger in the order shown during a single session:

```
(msg:"JPEG transfer"; content:"image/";pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fp?jpe?g/smi";
flowbits:set,http.jpeg; flowbits:noalert;)
```

The `content` and `pcre` keywords in the first rule fragment match a JPEG file download, `flowbits:set,http.jpeg` sets the `http.jpeg` flowbit 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 flowbit state so one or more companion rules can test for the state name in combination with malicious content and generate events when malicious content is detected.

The next rule fragment detects a GIF file download subsequent to the JPEG file download above:

```
(msg:"GIF transfer"; content:"image/"; pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
flowbits:set,http.gif,image_downloads; flowbits:noalert;)
```

The `content` and `pcre` 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.

Keysight, EX-1004
Page 854

The third rule fragment is a companion to the first rule fragment:

```
(msg:"JPEG exploit";
flowbits:isset,http.jpeg;content:"|FF|"; pcre:"
/\xFF[\xE1\xE2\xED\xFE]\x00[\x00\x01]/";)
```

In the third rule fragment, `flowbits:isset,http.jpeg` determines that the now-irrelevant `http.jpeg` state is set, and `content` and `pcre` match content that would be malicious in a JPEG file but not in a GIF file. The third rule fragment results in a false positive event for a nonexistent exploit in a JPEG file.

Now consider the same case except that the first two rules now include their two different state names in the same state group.

```
(msg:"JPEG transfer"; content:"image/";pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fp?jpe?g/smi";
flowbits:set,http.jpeg,image_downloads; flowbits:noalert;)
```

When the first rule fragment detects a JPEG file download, the `flowbits:set,http.jpeg,image_downloads` keyword sets the flowbit state to `http.jpeg` and includes the state in the mutually exclusive `image_downloads` group. Note that setting the rule state to `http.jpeg` unsets any other currently-set states in the group `image_downloads`, and vice versa.

The next rule then detects a subsequent GIF file download:

```
(msg:"GIF transfer"; content:"image/"; pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
flowbits:set,http.gif,image_downloads; flowbits:noalert;)
```

When the second rule fragment matches the GIF download, the `flowbits:set,http.gif,image_downloads` keyword sets the `http.gif` flowbit state and, because `http.jpeg` is also in the `image_downloads` group, unsets `http.jpeg`.

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]/";)
```

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can use the `http_encode` keyword to generate events on the type of encoding in an HTTP request or response prior to 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.

Keysight, EX-1004
Page 855

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 IPS 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 490.

You must also configure the preprocessor to inspect HTTP responses and HTTP cookies to return matches for these. See Decoding HTTP Traffic on page 560 and Selecting Server-Level HTTP Normalization Options on page 564 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 570 for more information.

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. For example, detects base36 encoding but base36 encoding is not enabled for decoding by the HTTP Inspect preprocessor.
base36	Detects Base 36 encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
uencode	Detects Microsoft %u encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
bare_byte	Detects bare byte encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.

Keysight, EX-1004
Page 856

To identify the HTTP encoding type and location in an intrusion rule:

Access: P&R Admin/Admin

1. Add the `http_encode` keyword to a rule.
2. From the **Encoding Location** dropdown 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, base36, 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 URl`
   - **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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The `file_data` keyword provides a pointer that serves as a reference for the positional arguments available for such other keywords as `content`, `byte_jump`,

Keysight, EX-1004
Page 857

`byte_test`, and `pcre`.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, you must enable the HTTP Inspect preprocessor and configure the preprocessor to inspect HTTP responses. See Decoding HTTP Traffic on page 560 and **Inspect HTTP Responses** in Selecting Server-Level HTTP Normalization Options on page 564 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, you must enable the HTTP Inspect preprocessor and 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 560, and the **Inspect HTTP Responses** and **Inspect Compressed Data** options in Selecting Server-Level HTTP Normalization Options on page 564. The `file_data` keyword matches if the HTTP Inspect preprocessor detects uncompressed gzip data in the HTTP response body.

- SMTP payload

  To inspect the SMTP payload, you must enable the SMTP preprocessor. See Configuring SMTP Decoding on page 585 for more information. The `file_data` keyword matches if the SMTP preprocessor detects SMTP data.

- Base64-encoded MIME email attachment

  To inspect MIME email attachments, you must enable the SMTP preprocessor and configure the preprocessor to decode MIME data. See Configuring SMTP Decoding on page 585 for more information; see also the **Enable MIME Decoding** and **Maximum MIME Decoding Depth** options in Understanding SMTP Decoding on page 582. The `file_data` keyword with the **MIME** argument matches if the SMTP preprocessor detects a MIME email attachment.

Note that you can use the fast pattern matcher in a `content` keyword that follows the file_data keyword except when you use the **MIME** argument to inspect Base64-encoded MIME data. See Constraining Content Matches on page 778 and Use Fast Pattern Matcher on page 789 for more information.

To point to the beginning of a specific payload type:

Access: P&R Admin/Admin

1. On the new rule Create page, select **file_data** from the drop-down list and click **Add Option.**

   The `file_data` keyword appears.



Keysight, EX-1004
Page 858

2. Optionally, select **MIME** from the `file_data` drop-down list to point to the beginning of a Base64-encoded MIME email attachment detected by the SMTP preprocessor.

Decoding and Inspecting Base64 Data

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 859
- base64_data on page 860

base64_decode

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 860 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

Keysight, EX-1004
Page 859

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: P&R Admin/Admin

1. On the new rule Create 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 860.

base64_data

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 860

Note the following when inspecting Base64 data:

- You cannot use the fast pattern matcher; see Use Fast Pattern Matcher on page 789 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 783 for more information.

To inspect decoded Base64 data:

Access: P&R Admin/Admin

1. On the new rule Create page, select **base64_data** from the drop-down list and click **Add Option.**

   The `base64_data` keyword appears.



Constructing a Rule

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 861
- Modifying Existing Rules on page 864
- Adding Comments to Rules on page 867
- Deleting Custom Rules on page 868

Writing New Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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

Keysight, EX-1004
Page 861

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 Objects on page 2009. See also, Exporting an Intrusion Policy on page 2001.

To write a custom standard text rule using the rule editor:

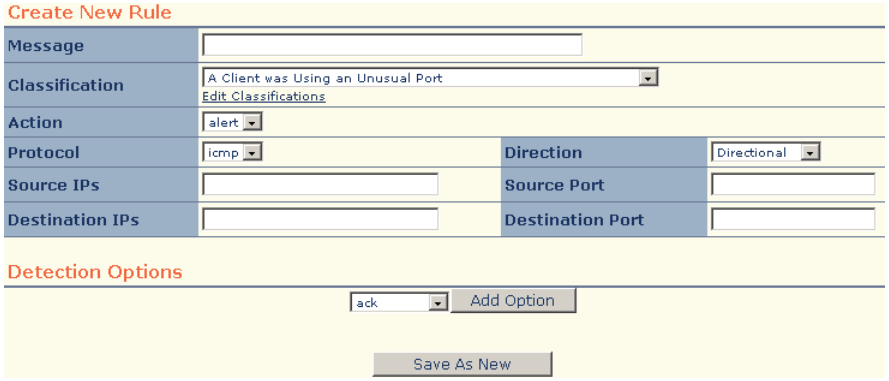
Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Rule Editor**.

   The Rule Editor page appears.

2. Click **Create Rule**.

   The Create page appears.



Keysight, EX-1004
Page 862

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 772.

   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 Event Classification on page 772.

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 766.

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 399.

Keysight, EX-1004
Page 863

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 401.

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 769.

    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 policy. See Applying an Intrusion Policy on page 347 for more information.

Modifying Existing Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 864

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 764and Deleting Custom Rules on page 868 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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Rule Editor**.

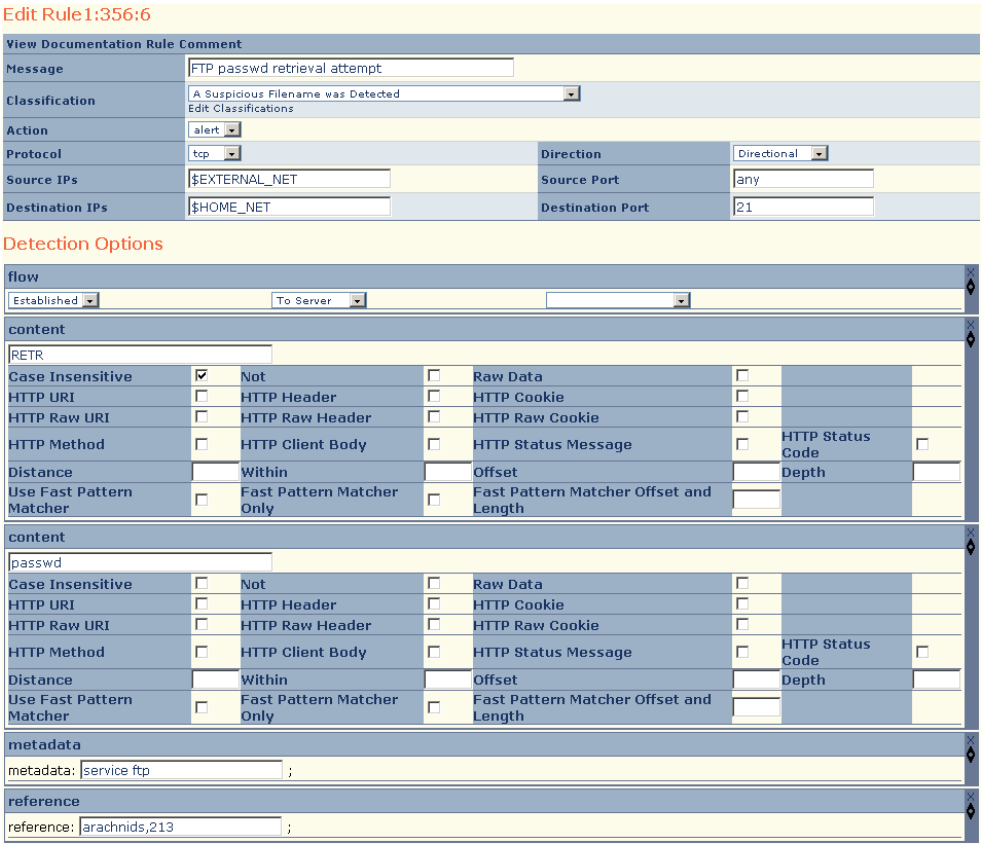
   The Rule Editor page appears.

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 **Edit** 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 869 for more information.

Keysight, EX-1004
Page 865

- 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 right of the rule list. Navigate to the rule you want and click **Edit** next to the rule. See Filtering Rules on the Rule Editor Page on page 872 for more information.

The rule editor opens, displaying the rule you selected.



Note that if you select a shared object rule, the rule editor displays only the rule header information. A shared object rule can be identified on the Rule Editor page by a listing that begins with the number 3 (the GID), for example, 3:1000004.

Keysight, EX-1004
Page 866

3. Make any modifications to the rule (see Writing New Rules on page 861 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 435 and activate the local rule.

4. Activate the intrusion policy as described in Applying an Intrusion Policy on page 347 to apply your changes.

Adding Comments to Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > Rule Editor**.

   The Rule Editor page appears.

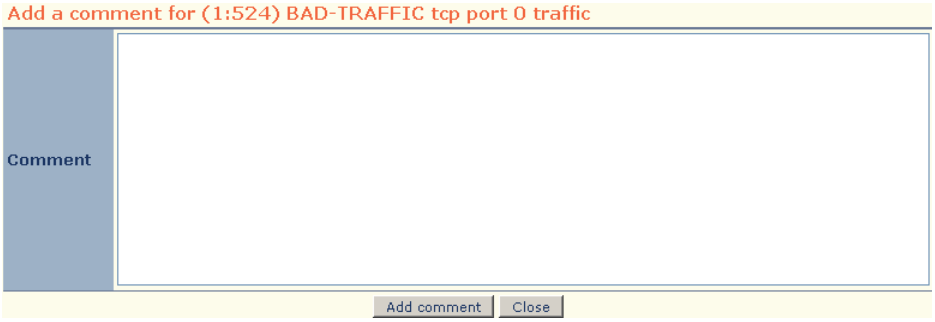
2. Locate the rule you want to annotate. You have the following options:
   - To locate a rule by browsing rule categories, navigate through the folders to the rule you want and click **Edit** next to the rule.
   - To locate a rule by searching for it, enter the search criteria (most simply, the SID) for the rule you want and click **Search**. Click the rule returned by the search as appropriate. See Searching for Rules on page 869 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 right of the rule list. Navigate to the rule you want and click **Edit** next to the rule. See Filtering Rules on the Rule Editor Page on page 872 for more information.

   The rule editor appears.

Keysight, EX-1004
Page 867

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 293.

Deleting Custom Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can delete custom rules that are not currently enabled in an intrusion policy. You cannot delete either standard text rules or shared object rules rules provided by 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 864 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 Import SEU page. See, for example, Using One-Time SEU Imports on page 1851.

See the following sections for more information:

- For information on creating custom rules, see Writing New Rules on page 861.
- For information on importing local rules, see Importing SEUs and Rule Files on page 1848.
- For information on setting rule states, see Setting Rule States on page 435.

Keysight, EX-1004
Page 868

To delete custom rules:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Rule Editor**.

   The Rule Editor page appears.

2. Navigate through the folders to the local rule category. You have two choices:

   - Click **Delete All**, 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.

   - Click on the local rule category to expand it, then click **Delete** 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! IPS 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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.

Keysight, EX-1004
Page 869

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 `overflow`.
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.

Keysight, EX-1004
Page 870

Rule Search Criteria (Continued)

Option	Description
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:

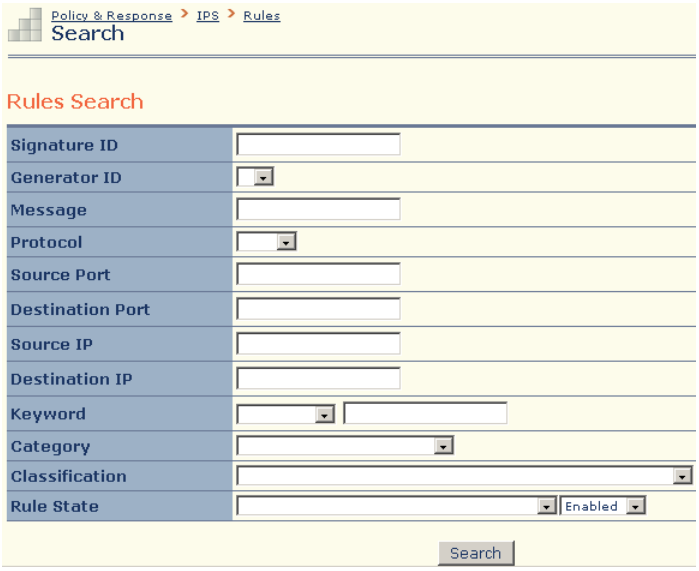
Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > Rule Editor**.

   The Rule Editor page appears.

2. Click **Search** on the toolbar.

   The Search page appears.



3. Add search criteria using any of the fields described in the Rule Search Criteria table on page 870.

IMPORTANT! You must specify at least one search criterion to search for rules.

Keysight, EX-1004
Page 871

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 769.
   - 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 864 for detailed information about editing rules.

Filtering Rules on the Rule Editor Page

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 an error 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.

Keysight, EX-1004
Page 872

See the following sections for more information:

- Using Keywords in a Rule Filter on page 873
- Using Character Strings in a Rule Filter on page 875
- Combining Keywords and Character Strings in a Rule Filter on page 875
- Filtering Rules on page 876

Using Keywords in a Rule Filter

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 873 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 875 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
`arachnids`	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 776 for more information.	`arachnids:181`
`bugtraq`	Returns one or more rules based on all or part of the Bugtraq ID in a rule reference. See Defining the Event Reference on page 776 for more information.	`bugtraq:2120`

Keysight, EX-1004
Page 873

Rule Filter Keywords (Continued)

Keyword	Description	Example
`cve`	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 776 for more information.	`cve:2003-0109`
`gid`	The argument `1` returns standard text rules. The argument `3` returns shared object rules. See Reading Preprocessor Generator IDs on page 488 and the Rule Types table on page 405 for more information.	`gid:3`
`mcafee`	Returns one or more rules based on all or part of the McAfee ID in a rule reference. See Defining the Event Reference on page 776 for more information.	`mcafee:10566`
`msg`	Returns one or more rules based on all or part of the rule Message field, also known as the event message. See Defining the Event Message on page 772 for more information.	`msg:chat`
`nessus`	Returns one or more rules based on all or part of the Nessus ID in a rule reference. See Defining the Event Reference on page 776 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 776 and Defining the Event Message on page 772 for more information.	`ref:MS03-039`
`sid`	Returns the rule with the exact Signature ID. See Reading Preprocessor Generator IDs on page 488 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 776 for more information.	`url:faqs.org`

Keysight, EX-1004
Page 874

Using Character Strings in a Rule Filter

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 772. For information on rule SIDs and GIDs, see Reading Preprocessor Generator IDs on page 488.

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

Platform: DC/MDC/Standalone Sensor

Component: IPS

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`

Keysight, EX-1004
Page 875

Filtering Rules

Platform: DC/MDC/Standalone Sensor

Component: IPS

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: P&R Admin/Admin

1. Select **Policy & Response > IPS > 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 864 for more information.

2. Optionally, select a different grouping method from the Group Rules By list.

   TIP! Filtering can 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 right 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 872 for more information.

   The filter clearing icon (⊠) appears when you type the first character and disappears if you delete all characters.

   The page refreshes to display any group that contains at least one matching rule.

Keysight, EX-1004
Page 876

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 864.

   To put any changes you make into effect, apply the intrusion policy to the appropriate detection engines as described in Applying an Intrusion Policy on page 347.

Keysight, EX-1004
Page 877

Chapter 21
Introduction to Sourcefire RNA

Sourcefire RNA allows your organization to confidently monitor and protect your network using a combination of forensic analysis, behavioral profiling, and built-in alerting and remediation. 3D Sensors with RNA passively observe your organization's network traffic and analyze it to provide you with a complete, up-to-the-minute profile of your network.

By default, RNA is installed on most 3D Sensors. (The 3D9800 does not support RNA.) Sourcefire also makes key components of RNA available in installation packages for Crossbeam Systems X-Series Platforms. However, to control how network intelligence is gathered and to view the resulting information, you must manage 3D Sensors with RNA with a Defense Center. In addition, to enable RNA functionality, that Defense Center must have an RNA host license installed.

IMPORTANT! You **cannot** forward RNA data to a Master Defense Center.

To supplement the data gathered by RNA, you can use records generated by NetFlow-enabled devices, Nmap active scans, and the host input feature.

The following sections explain more about RNA and how to include RNA in your Sourcefire 3D System deployment:

- Understanding RNA on page 879
- Creating RNA Detection Policies on page 900
- Managing RNA Detection Policies on page 910

Keysight, EX-1004
Page 878

Understanding RNA

Platform: DC

Component: RNA

RNA allows your organization to confidently monitor and protect your network using a combination of forensic analysis, behavioral profiling, and built-in alerting and remediation. 3D Sensors with RNA passively observe your organization's network traffic and analyze it to provide you with a complete, up-to-the-minute profile of your network.

RNA analyzes network traffic using what are called detection engines, which reside on 3D Sensors with RNA and tie together computing resources and a set of network interfaces. RNA detection engines analyze the traffic on the network segments to which their interfaces are connected.

After you have physically deployed your Sourcefire 3D System appliances (and, optionally, NetFlow-enabled devices), you can begin to configure the Sourcefire 3D System to gather RNA data so that you can begin your analysis.

The most critical piece of your RNA deployment is the RNA detection policy. Once applied to your RNA detection engines, the RNA detection policy controls how RNA events and flow data are collected, which network segments are monitored by 3D Sensors with RNA and which are monitored with NetFlow-enabled devices, and whether traffic that travels from or to specific ports is excluded from monitoring.

One of the major decisions you must make when creating an RNA detection policy is whether you want to use subnet detection, which allows RNA to make recommendations indicating the optimal detection engines to analyze the traffic on the various network segments in your organization.

As RNA continues to monitor your network, it may be able to refine its subnet recommendations, especially if your network configuration has been altered through routing or host changes. Choosing which subnets to monitor with which detection engines is an iterative process that you can either revisit from time to time, or configure the Sourcefire 3D System to automatically handle.

For more information on RNA, see the following sections:

- What Data Does RNA Collect? on page 880
- What Information Does RNA Provide? on page 880
- How Can You Use RNA? on page 886
- What are the Prerequisites for RNA Data Collection? on page 887
- What is an RNA Detection Policy? on page 889
- How Do I Manage my RNA Deployment with Subnet Detection? on page 897

Keysight, EX-1004
Page 879

For more information on NetFlow, Nmap active scans, and the host input feature, see the following sections:

- Introduction to NetFlow on page 1140
- Configuring Active Scanning on page 1418
- Importing Host Input Data on page 1406

What Data Does RNA Collect?

Platform: DC

Component: RNA

RNA passively monitors the traffic that travels through your network, including data exported by the NetFlow-enabled devices in your deployment. RNA collects and reports various information about your network, including:

- the number and types of hosts (including network devices such as bridges, routers, load balancers, and NAT devices) running on your network
- basic network topology data, indicating which detection engines are closest to which monitored hosts based on the number of hops
- the operating systems and client applications running on monitored hosts
- the active services and open ports on monitored hosts
- the vulnerabilities and exploits to which monitored hosts may be susceptible
- flow data, which are records of active sessions involving monitored hosts

What Information Does RNA Provide?

Platform: DC

Component: RNA

RNA analyzes the data it collects, comparing specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine what types of hosts are communicating on the network and which operating systems and services are running on the hosts.

RNA also determines which of your hosts are closest to which RNA detection engines, and uses that information to recommend which detection engines should be primarily responsible for monitoring specific subnets.

RNA events, which are triggered by the changes detected on your network, and flow events, which are records of sessions involving monitored hosts, alert you to the activity on your network and provide you with the information you need to respond appropriately. You can also use the flow data that your 3D Sensors with RNA and your NetFlow-enabled devices collect to create and view a profile of your normal network traffic.

Keysight, EX-1004
Page 880

When you install sensors with IPS and RNA on the same network segment and then view the intrusion events on a Defense Center that manages both sensors, you can gain more insight into the risk, or impact, that the events have on your network.

RNA compares the data it collects and analyzes with its vulnerability database to determine the potential vulnerabilities on the detected host, and also maps and displays network and vulnerability data that is easily viewable and searchable from the Defense Center web interface.

You can use Nmap active scans to add information about operating systems, services, and vulnerabilities, supplementing the data gathered by RNA. You can also use the host input feature to supplement the information that RNA gathers, either by configuring a third-party application to interact with the Sourcefire 3D System via an API, or by manually adding data. RNA reconciles the collected data about operating system and service identities and determines each identity based on fingerprint source priority values, identity conflict resolution settings, and time of collection.

For more information, see the following sections:

- Subnet Monitoring Recommendations on page 881
- Network Map on page 882
- Host Profiles and Host Attributes on page 882
- RNA Events on page 883
- Flow Data and Traffic Profiles on page 884
- Impact Flags on page 885
- Host Vulnerabilities on page 886
- Configuring Active Scanning on page 1418
- Importing Host Input Data on page 1406

Subnet Monitoring Recommendations

If your deployment includes multiple 3D Sensors with RNA (or one 3D Sensor using different interface sets to monitor multiple networks), it is important to configure each RNA detection engine to monitor hosts in close proximity. In other words, you should configure RNA detection policies so each RNA detection engine is configured as the reporting detection engine for the hosts that are closest to it from a network hop standpoint. (The reporting detection engine for a particular network segment reports most of the information about the hosts on that network, for example, the operating systems and services running on the hosts.)

Unfortunately, you may not always be informed of network configuration changes. A network administrator may modify a network configuration through routing or host changes without informing you, which can make it challenging to stay on top of proper RNA policy configurations.

Keysight, EX-1004
Page 881

To mitigate this problem, RNA can automatically determine the closest subnets to each RNA detection engine, than make recommendations about which detection engines should be the reporting detection engines for specific subnets. The Sourcefire 3D System can automatically update the RNA detection policy with this information, or it can notify you so that you can make the changes manually.

With subnet detection, you do not need insight into network topology, nor must you continually update RNA detection policies, to monitor hosts closest to each 3D Sensor. This saves time and maximizes the performance of RNA running on 3D Sensors.

Note that subnet detection is **not** supported in the following cases:

- if your deployment includes only one RNA detection engine; with only one RNA detection engine, the only choice is automatically the best choice
- if your deployment includes load-balanced Crossbeam–based software sensors in multi-member RNA VAP groups
- for networks you are monitoring with NetFlow-enabled devices

For more information, see:

- Monitor the Appropriate Network Segments on page 891
- How Do I Manage my RNA Deployment with Subnet Detection? on page 897
- Managing Subnet Detection on page 911

Network Map

RNA aggregates data for each host and service into the *network map*, which is an up-to-date representation of your network assets. The network map is continuously updated as network traffic indicates a change or a new host. RNA performs its analysis passively, which means that, unlike active scanners, it does not interact with or scan your hosts directly. It continuously evaluates your network traffic to ensure that your network map stays current. You can quickly navigate through the network map to access other information relevant to your analysis.

Note that although you can configure the RNA detection policy to add hosts and services to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts and services is limited.

For more information, see:

- Using the Network Map on page 923
- How is NetFlow Data Different from RNA Flow Data? on page 1142

Host Profiles and Host Attributes

RNA also provides *host profiles* for the hosts it detects. A host profile provides a complete view of all the information available for a single host. You can access general host information, such as the host name and operating system, through

Keysight, EX-1004
Page 882

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, and include annotations you can make for any monitored host, as well as an indicator of business criticality (called *host criticality*) that establishes the host's importance. This allows you to easily track changes to servers that are of the highest importance to your organization. You can also create, modify, or delete your own host attributes and apply them either manually or automatically (based on IP address). You can also set host attribute values in response to compliance rule violations.

For more information, see Using Host Profiles on page 944.

RNA Events

RNA events are triggered by changes detected in your monitored network.

The events generated by RNA can alert you to the activity on your network and provide you with the information you need to respond appropriately. As a simple example, you may have conference rooms or a series of spare work spaces where visiting employees connect to your network. You would expect to see New Host events generated by the detection engines monitoring 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, you can escalate your response accordingly.

RNA 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. You can also set up RNA to detect the services, network protocols, client applications, and potential vulnerabilities running on your network hosts. In addition, you can track RNA-specific features such as changes to host criticality, user-defined host attributes, and user-initiated interactions with host vulnerabilities.

Host input events, a subset of RNA events, allow you track changes to your network map made using the host input feature.

RNA provides a set of predefined workflows that you can use to analyze the events that are generated for your network. You can use these predefined workflows, or you can create custom workflows that display only the information that matches your specific needs.

For more information, see Working with RNA Events on page 994.

Keysight, EX-1004
Page 883

Flow Data and Traffic Profiles

RNA generates flow events when connections between monitored hosts and any other hosts are terminated. Flow events include information about the collected traffic, including:

- the flow type (whether the flow was detected by a 3D Sensor with RNA, or by a NetFlow-enabled device), and if the flow was detected by a NetFlow-enabled device, the IP address of the device
- the timestamps of the first and last packet of the transaction
- the source IP address, destination IP address, and the ports and protocol used by the flow
- any TCP flags detected in the flow (for flows exported by NetFlow-enabled devices)
- the user logged into the source and destination hosts (if RUA is enabled and the hosts are on the monitored network)
- the number of packets and bytes sent and received by the monitored host
- the client application and URL involved in the transaction, if applicable

Note that not only does some of the information in a flow event depends on the flow type, but also the number of flow events generated for each flow depends on how the flow was detected:

- For flows detected by 3D Sensors, RNA generates one bidirectional flow event.
- NetFlow-enabled devices export unidirectional flow data, so RNA generates at least two flow events for each flow detected by NetFlow-enabled devices, depending on whether you configured the device to output records at a fixed interval even if a flow is still ongoing.
- If you are monitoring the same network segment with NetFlow-enabled devices and 3D Sensors with RNA, RNA generates at least three flow events for each detected flow—one that represents the flow as detected by the sensor, and at least two for the flow as detected by NetFlow-enabled device.

As with RNA events, you can view flow data using a set of predefined workflows, or you can create custom workflows. You can choose to view flow data in a table format or as graphs. You can also use the flow data to create and view a profile of your normal network traffic, called a *traffic profile*.

In the RNA detection policy, there are two ways of specifying how flow data is collected. The first way is by setting the flow data mode, which affects all

Keysight, EX-1004
Page 884

detection engines where you apply the RNA detection policy. You can set the flow mode so that RNA:

- transmits individual flows from RNA detection engines to the Defense Center
- transmits flow data to the Defense Center as aggregated data over five-minute intervals, called flow summaries (the default)
- entirely disables flow data collection

The data collection mode, which you set on a network-by-network basis, also affects how flow data is collected. For each network segment you monitor with an RNA detection engine, RNA can:

- collect data that is used to populate host profiles for the hosts in the network, as well as flow data for those hosts
- collect only host data for the hosts in the network (the default)
- collect neither host data nor flow data for the hosts in the network

Note that disabling or otherwise limiting flow data collection can reduce the traffic between your 3D Sensors and your Defense Center as well as reduce the space required to store flow data on the Defense Center. However, it can also prevent or limit you from taking advantage of features that rely on flow data.

For more information, see:

- Filter Flows as Needed on page 893
- Working with Flow Data and Traffic Profiles on page 1071
- Introduction to NetFlow on page 1140

Impact Flags

When you install sensors with IPS and RNA on the same network segment and then view the intrusion events on a Defense Center that manages both sensors, you can gain more insight into the risk, or impact, that the events have on your network.

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 detection data, RNA events, and vulnerability information.

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 red (Vulnerable) impact flags for intrusion events involving those hosts, unless you use the host input feature to manually set the hosts' operating system identity.

Keysight, EX-1004
Page 885

For more information, see:

- Using Impact Flags to Evaluate Events on page 307
- How is NetFlow Data Different from RNA Flow Data? on page 1142

Host Vulnerabilities

RNA also correlates the operating system and services detected on each host with a vulnerability database to help you determine whether changes to a particular host increase your risk of network compromise. When host information changes, RNA updates the vulnerability assessment. For example, if you upgrade a host's operating system to a new version, RNA detects the change by analyzing the network traffic the host produces, and automatically generates an updated list of known vulnerabilities specific to the new operating system.

You can manage vulnerabilities individually on a host, mark vulnerabilities as invalid by applying a patch that addresses the vulnerabilities, or configure fixes for an operating system to mark all addressed vulnerabilities as invalid on all hosts where you add that fix.

If RNA cannot identify the operating system of a host on your network, you can use RNA's custom fingerprinting feature to create custom client or server fingerprints. RNA uses these fingerprints to identify new hosts. You can map fingerprints to systems in the vulnerability database to allow RNA to display appropriate vulnerability information whenever a host is identified using the custom fingerprint.

Because there is no operating system information available for hosts added to the network map based on NetFlow data, you cannot use RNA to evaluate the vulnerabilities for those hosts, unless you use the host input feature to manually set the hosts' operating system identity.

For more information, see:

- Working with the Vulnerabilities Network Map on page 932
- Working with Vulnerabilities in the Host Profile on page 978
- Working with Sourcefire Vulnerabilities on page 1057
- Using Custom Fingerprinting on page 1363
- How is NetFlow Data Different from RNA Flow Data? on page 1142

How Can You Use RNA?

Platform: DC

Component: RNA

You can use RNA not only to view a complete, up-to-the-minute profile of your network, but you can also use the powerful policy and response feature to build *compliance policies* that let your Defense Center respond in real time to threats on your network.

Keysight, EX-1004
Page 886

Compliance policies use the data and events gathered and generated by 3D Sensors with RNA, with IPS, or with both to describe the type of activity that constitutes a policy violation. Policy violations can occur when:

- a specific intrusion event, RNA event, host input event, flow event, or RUA event occurs that meets specific criteria as characterized in a *compliance rule*
- your network traffic deviates from your normal network traffic pattern as characterized in an existing *traffic profile*
- RNA detects that a host on your network is running a forbidden operating system, client application, service, protocol, or combination thereof, as characterized in a *compliance white list*

Using policy and response, you can configure RNA to respond to suspicious activity in real-time by using compliance policies to send *alerts,* which can include email, syslog, and SNMP notifications.

You can also configure the Defense Center to respond to policy violations by launching *remediations*, which are programs that the Defense Center runs. The Defense Center ships with predefined remediations, which allow you to block suspicious traffic using a Cisco PIX firewall, a Cisco IOS router, Check Point FireWall-1, or use Nmap to scan hosts producing or receiving suspicious traffic. You can also write custom remediations to perform actions that are more specific to your organization's remediation or integration requirements.

For more information, see:

- Configuring Compliance Policies and Rules on page 1212
- Working with Flow Data and Traffic Profiles on page 1071
- Using RNA as a Compliance Tool on page 1156
- Configuring Responses for Compliance Policies on page 1281

What are the Prerequisites for RNA Data Collection?

Platform: DC

Component: RNA

Before you can use RNA, you must configure various settings to specify not only the kind of information that RNA collects and stores, but also other kinds of RNA behavior.

System Settings

Use the system settings to apply an RNA host license to your Defense Center. A host license allows you to collect and store RNA data. For more information, see Managing Feature Licenses on page 1787.

Keysight, EX-1004
Page 887

System Policy

Use the system policy (see Managing System Policies on page 1732) to configure several RNA-related settings:

- Database settings control the maximum number of various types of events that you want to store on the Defense Center. To work with RNA-related events, you must make sure that RNA event storage is enabled. For more information, see Configuring Database Event Limits on page 1747.
- Detection policy preferences control whether you must confirm your action when you apply an RNA detection policy. For more information, see Configuring RNA Settings on page 1756.
- RNA settings control the kinds of RNA data stored in the database, and therefore determine the data that other parts of the Sourcefire 3D System can use. RNA settings also control various timeout settings, how flow data is stored, how duplicate events and service and operating system identity conflicts are handled, subnet detection settings, the priority of active RNA data sources such as scanners, and how the Sourcefire 3D System performs impact flag correlation. For more information, see Configuring RNA Settings on page 1756.

RNA Detection Policy

Use the RNA detection policy to configure:

- how RNA events and flow data are collected
- which network segments are monitored by 3D Sensors with RNA and which are monitored with NetFlow-enabled devices
- whether traffic that travels from or to specific ports is excluded from monitoring

For more information, see What is an RNA Detection Policy? on page 889.

NetFlow Prerequisites

If you are planning to use NetFlow data to supplement the data gathered by 3D Sensors with RNA, see What are the Prerequisites for NetFlow Data Collection? on page 1146 for more information.

Keysight, EX-1004
Page 888

What is an RNA Detection Policy?

Platform: DC

Component: RNA

The key step in deploying RNA is creating and applying an RNA detection policy. Once applied to your RNA detection engines, the RNA detection policy controls the following aspects of the Sourcefire 3D System:

- how RNA events and flow data are collected
- which network segments are monitored by 3D Sensors with RNA and which are monitored with NetFlow-enabled devices
- whether traffic that travels from or to specific ports is excluded from monitoring

When you create a detection policy, it is important to keep several points in mind, as described in the following sections:

- Create and Apply One RNA Detection Policy on page 889
- Configure One Detection Engine per 3D Sensor on page 890
- Do Not Overlap Detection Engine Coverage on page 890
- Monitor the Appropriate Network Segments on page 891
- Exclude Load Balancers and NAT Devices as Needed on page 892
- Exclude Ports as Needed on page 893
- Filter Flows as Needed on page 893
- Special Considerations: NetFlow Monitoring on page 894
- Special Considerations: Crossbeam-Based Software Sensors on page 896
- Special Considerations: Clustering on page 896

Create and Apply One RNA Detection Policy

Rather than creating a separate detection policy for each detection engine or network segment, in most cases you should create one policy that maps all the detection engines in your deployment to the appropriate network segments.

When you apply a single RNA detection policy to multiple detection engines, the reporting detection engine for a particular network segment is responsible for reporting most of the information about the hosts on that network, for example, the operating systems and services running on the hosts.

Any other detection engines where you apply the same RNA detection policy report secondary information about the host, such as MAC address data and the number of hops away from the sensor where the detection engine is running. This secondary information can help RNA better determine the topology of your network. If you were to apply multiple detection policies to multiple detection engines, you would lose this secondary information.

Keysight, EX-1004
Page 889

In addition, and more important, applying one detection policy to all of your detection engines can significantly reduce the load on your Defense Center if you enable flow data collection. For example, if you have two policies, each of which is monitoring a separate network segment using separate detection engines, each detection engine generates a flow event when RNA detects that a connection is terminated between a monitored host on one of the networks and a monitored host on the other network. On the other hand, if you use one policy to monitor both networks, only the reporting detection engine for the flow initiator generates a flow event.

If for any reason you need to use two RNA detection policies, and thus RNA generates duplicate flow events, you can configure the RNA settings in the system policy so that the Defense Center drops these events. However, this can degrade performance. For more information, see Configuring RNA Settings on page 1756.

Configure One Detection Engine per 3D Sensor

You should configure only one RNA detection engine per 3D Sensor, unless you are using separate detection engines to monitor separate network interfaces, or are dedicating an RNA detection engine to processing NetFlow data.

Many sensors have only one detection resource to use for RNA and therefore support only one RNA detection engine, although certain models can support more. For more information, see Understanding Detection Resources and 3D Sensor Models on page 202.

Do Not Overlap Detection Engine Coverage

Sourcefire recommends that in most cases you do not assign more than one detection engine as the reporting detection engine for the same network segment.

For example, consider a pair of 3D Sensors deployed in different physical locations. Now, assume that the RNA detection policy on the sensors' managing Defense Center has configured the detection engines on both sensors to monitor the same network segment.

If there is some kind of device between a host and one of the sensors that modifies the fingerprint of the packets transmitted by a host, but not between the host and the other sensor, the sensors can gather conflicting information about the host. This can degrade performance as the Defense Center attempts to resolve the conflicts.

In addition, overlapping detection engine coverage can result in duplicate flow events. Although you can configure the RNA settings in the system policy so that

Keysight, EX-1004
Page 890

the Defense Center drops these duplicate events, this can also degrade performance. For more information, see Configuring RNA Settings on page 1756.

IMPORTANT! The system policy allows you to drop duplicate flow events for flows detected by 3D Sensors with RNA. It also allows you to drop duplicate flows that are exported by NetFlow-enabled devices. However, if you are monitoring the same network segment with NetFlow-enabled devices and 3D Sensors with RNA, RNA will generate flow events for flows detected by both devices. You **cannot** drop these duplicate events.

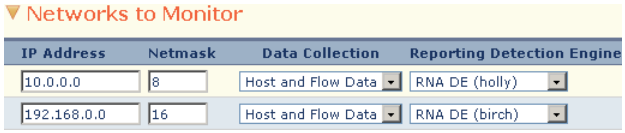
Monitor the Appropriate Network Segments

Optimally, your RNA detection policy should specify that each RNA detection engine is configured as the reporting detection engine for the hosts that are closest to it. This helps the Sourcefire 3D System report accurate information about the hosts on your monitored network segments, for example, the operating systems and services running on the hosts.

If you know which networks the sensing interfaces on your 3D Sensors are physically connected to, monitoring the appropriate segments is a straightforward task. You can simply configure the RNA detection policy so that the reporting detection engine for a subnet is the detection engine that uses the interface set connected to that subnet.

As a simple example, consider a pair of 3D Sensors named `holly` and `birch`, where `holly`'s sensing interface is physically connected to the 10.0.0.0/8 network, and `birch`'s sensing interface is physically connected to the 192.168.0.0/16 network.

The following graphic shows an RNA detection policy where the reporting detection engines for the 10.0.0.0/8 network resides on `holly` and the reporting detection engine for the 192.168.0.0/16 network resides on `birch`.



However, if you are not completely sure where and how your 3D Sensors are deployed, you can set the reporting detection engines for a subnet to **Auto-detect**. This allows RNA to use a process called *subnet detection* to evaluate your network and determine which detection engines should monitor which subnets. After RNA determines the optimal detection assignments, it presents you with recommendations that you can accept or reject. For information on the recommendations that RNA generates, see Understanding RNA-Generated Subnet Recommendations on page 911.

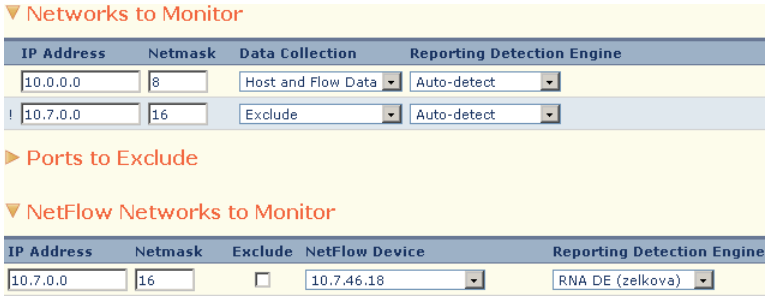
As RNA continues to monitor your network traffic, it may be able to refine its recommendations, especially if your network configuration has been altered

Keysight, EX-1004
Page 891

through routing or host changes. Choosing which subnets to monitor with which detection engines is an iterative process, and you may want to revisit the subnet recommendations from time to time to adjust your detection engine assignments. Alternately, you can configure the Sourcefire 3D System to automatically generate and apply new subnet recommendations on a daily basis, or to notify you via email so that you can make the changes manually. For information about the process of subnet detection, see How Do I Manage my RNA Deployment with Subnet Detection? on page 897.

For a scenario that involves NetFlow, consider monitoring the 10.0.0.0/8 supernet, but using a NetFlow-enabled device to monitor the 10.7.0.0/16 subnet. The following graphic shows an RNA detection policy where the reporting detection engine for the 10.0.0.0/8 supernet is set to auto-detect. However, the 10.7.0.0/16 subnet has been excluded from monitoring so that a NetFlow-enabled device (in this example, 10.7.46.18) can monitor it without duplicating flow data.

Because subnet detection is not supported for NetFlow monitoring, you must explicitly assign a detection engine on a 3D Sensor to process exported NetFlow data. This scenario uses the detection engine on `zelkova`.



Exclude Load Balancers and NAT Devices as Needed

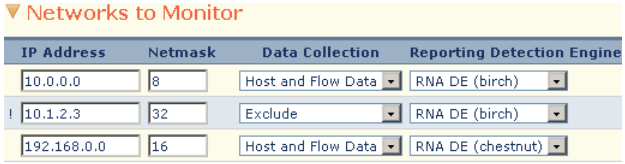
You should exclude load balancers (or specific ports on load balancers; see Exclude Ports as Needed on page 893) and NAT devices from monitoring. These devices can create excessive and misleading events, filling the database and overloading the Defense Center. For example, a monitored NAT device might exhibits 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 when you create your RNA detection policy.

TIP! RNA 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 RNA detection policy, wait for RNA to populate the network map, then perform a search of RNA hosts constraining on host type. For more information, see Searching for Hosts on page 1027.

Keysight, EX-1004
Page 892

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 RNA 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 detection policy to monitor that IP address again. For more information, see Fingerprinting Servers on page 1370.

The following graphic shows an RNA detection policy that excludes the host with an IP address of 10.1.2.3 from monitoring. This host does not appear in the network map and no events are reported for it.



Exclude Ports as Needed

Just you can exclude hosts from RNA monitoring (see Exclude Load Balancers and NAT Devices as Needed on page 892), you can exclude specific ports from monitoring by RNA detection engines. You cannot configure NetFlow-enabled devices to exclude ports from monitoring.

For example, load balancers can report multiple services on the same port in a short period of time. You can configure your RNA detection 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 application that uses a specific range of ports. If the traffic from this application generates excessive and misleading events, you can exclude those ports from monitoring. Similarly, you may decide that you do not want to use RNA to monitor DNS traffic. In that case, you could configure your detection policy so that it does not monitor port 53.

Filter Flows as Needed

For every network you monitor with an RNA detection engine (flow filtering is not supported for NetFlow data), you must set a data collection mode. This affects how flow data is collected. You can:

- collect data that is used to populate host profiles for the hosts in the network, as well as flow data for those hosts
- collect only host data for the hosts in the network (the default)
- collect neither host data nor flow data for the hosts in the network

In other words, when RNA detects a flow between two hosts, the information that RNA records in the database depends on the data collection mode you set for the networks where the hosts reside.

Keysight, EX-1004
Page 893

If you set the same data collection mode for two networks, RNA collects data for hosts in those networks consistent with the collection mode. For example, if you set the data collection mode for Networks A and B to host only, and then RNA detects a flow between a host in Network A (Host A) and a host in Network B (Host B), RNA does not log a record of the flow, but it does record general host information about the two hosts, such as host type and operating system data.

On the other hand, the following table explains what kind of data RNA records if the data collection modes for two networks differ.

Flow Filtering Consequences

For Host A, if you collect...	And for Host B, if you collect...	Then, if there is a flow between the two hosts, RNA collects...
host and flow data	host data only	host and flow data for both hosts
host and flow data	no data	host and flow data for Host A no data for Host B
host data only	no data	host data for Host A no data for Host B

Disabling or otherwise limiting flow data collection for certain networks can reduce the traffic between your 3D Sensors and your Defense Center as well as reduce the space required to store flow data on the Defense Center. You may want to only turn on flow collection for subnets that are most critical to your operations, while continuing to collect host data for your entire network.

On the other hand, for networks where you have disabled flow collection, you cannot take advantage of any feature that relies on flow data, such as the Flow Summary dashboard, traffic profiles, compliance rules based on flows or traffic profile changes, and flow trackers in compliance rules.

If you allow host data collection only, you can still use the network map, view host profiles, perform vulnerability and intrusion event impact assessment, as well as use compliance rules and white lists. Excluding a network excludes it from monitoring altogether.

Special Considerations: NetFlow Monitoring

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

Keysight, EX-1004
Page 894

through the devices. You can use this flow data to supplement the flow data gathered directly by your 3D Sensors with RNA. This is useful if you have NetFlow-enabled devices deployed on networks that your sensors cannot monitor, because you can use NetFlow data to monitor those networks.

However, there are several points you must keep in mind when configuring your RNA detection policy using NetFlow-enabled devices:

- Although you can use NetFlow-enabled devices exclusively to monitor your network, your deployment must include at least one 3D Sensor with RNA. This sensor must be deployed such that one of its sensing interfaces is connected to the destination network where your NetFlow-enabled device exports flow records. This ensures that the RNA detection engines on the sensor can collect and analyze the NetFlow data.
- Subnet detection is not supported for networks you are monitoring with NetFlow-enabled devices. If you want to use NetFlow-enabled devices to monitor all or part of your network, you must know where and how they are physically deployed so you can monitor the appropriate network segments.
- Flow filtering, that is, setting a data collection mode that affects how flow data is collected, is not supported for networks you are monitoring with NetFlow-enabled devices. You can, however, entirely exclude networks from monitoring.
- The port exclusion feature only affects data collected by 3D Sensors with RNA. You cannot configure NetFlow-enabled devices to exclude ports from monitoring.

For more information on using NetFlow-enabled devices to monitor your network, including the prerequisites for NetFlow data collection and the differences between the NetFlow data and the flow data detected directly by detection engines on 3D Sensors with RNA, see Introduction to NetFlow on page 1140.

Keysight, EX-1004
Page 895

Special Considerations: Crossbeam-Based Software Sensors

As described in the *Sourcefire 3D Sensor Software for X-Series Installation Guide*, you can take advantage of the load-balancing and redundancy benefits of the Crossbeam Systems X-Series Platform by deploying Defense Center with RNA in multi-member VAP groups, with each VAP running its own instance of the Defense Center.

If you choose to deploy RNA in this way, you must use the Defense Center to configure each Crossbeam-based software sensor with the same number of detection engines using the same interface sets.

You must also create separate RNA detection policies, one to apply to the detection engines on each software sensor. The RNA detection policies should be identical in almost all respects; you should configure each RNA detection policy to have identical detection policy settings, networks to monitor, and ports to exclude. The only difference between the policies should be the reporting detection engines you are using to monitor your network.

For example, if you want to use two load-balanced software sensors (in this example, RNA_A and RNA_B) to monitor the 10.0.0.0/8 network, you must create two nearly identical RNA detection policies. In one policy, use the detection engine on RNA_A to monitor the network; in the other, use the detection engine on RNA_B.

WARNING! RNA subnet detection is **not** supported for load-balanced multi-member RNA VAP groups. Do **not** set reporting detection engines to auto-detect when configuring RNA detection policies for load-balanced software sensors.

For more information on configuring Crossbeam-based software sensors, see the *Sourcefire 3D Sensor Software for X-Series Installation Guide*.

Special Considerations: Clustering

If you apply an RNA detection policy to an RNA detection engine on a 3D9900 sensor and then establish clustering with that sensor, you **must** edit and reapply your detection policy on the slave sensor.

Also, you cannot apply an RNA detection policy to the combined detection engines on a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). If you used the system policy to configure your Defense Center to automatically apply subnet recommendations, the auto-apply does not occur for clustered pairs with mismatched versions.

For more information, see Using Clustered 3D Sensors on page 250.

Keysight, EX-1004
Page 896

How Do I Manage my RNA Deployment with Subnet Detection?

Platform: DC

Component: RNA

After you have physically deployed your Sourcefire 3D System appliances, you can begin to configure the Sourcefire 3D System to gather RNA data so that you can begin your analysis.

One of the major decisions you must make when configuring RNA data collection is whether you want to use subnet detection in your RNA detection policy. Using subnet detection allows RNA to automatically determine the closest subnets to each RNA detection engine in your deployment. RNA can then make recommendations about which detection engines should be the reporting detection engines for specific subnets.

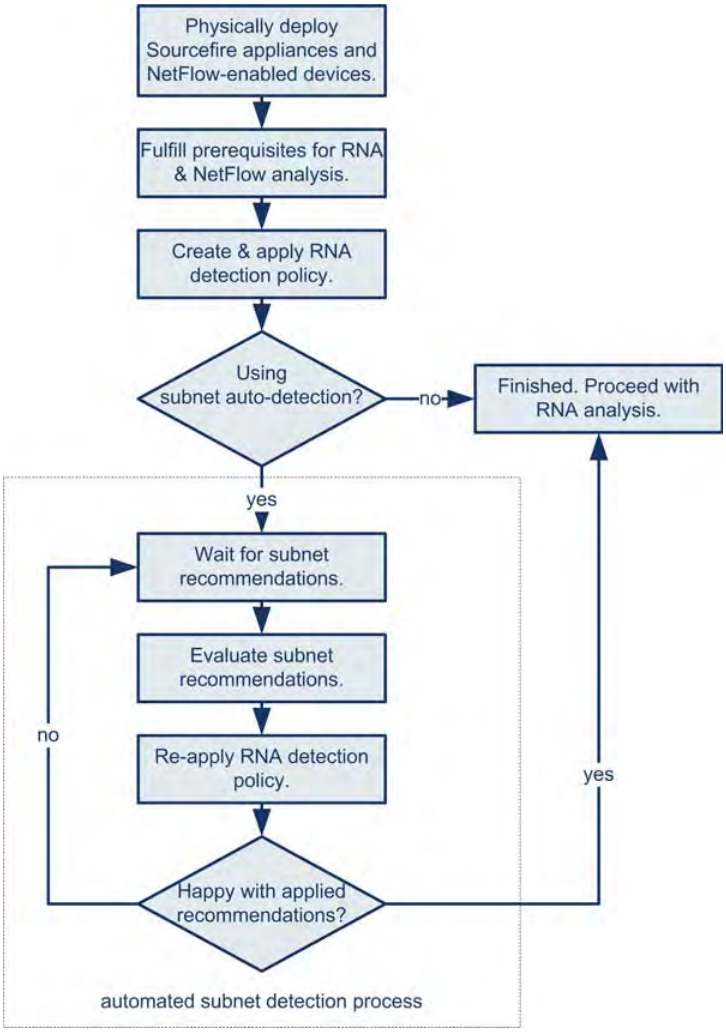
Note that subnet detection is **not** supported in the following cases:

- if your deployment includes only one RNA detection engine; with only one RNA detection engine, the only choice is automatically the best choice
- if your deployment includes load-balanced Crossbeam-based software sensors in multi-member RNA VAP groups
- for networks you are monitoring with NetFlow-enabled devices

Choosing which subnets to monitor with which detection engines is an iterative process. As the Sourcefire 3D System continues to monitor your network, it may be able to refine its recommendations, especially if your network configuration has been altered through routing or host changes. You may want to revisit the subnet recommendations from time to time to adjust your detection engine assignments. Alternately, you can configure the Sourcefire 3D System to automatically generate and apply new recommendations on a daily basis.

Keysight, EX-1004
Page 897

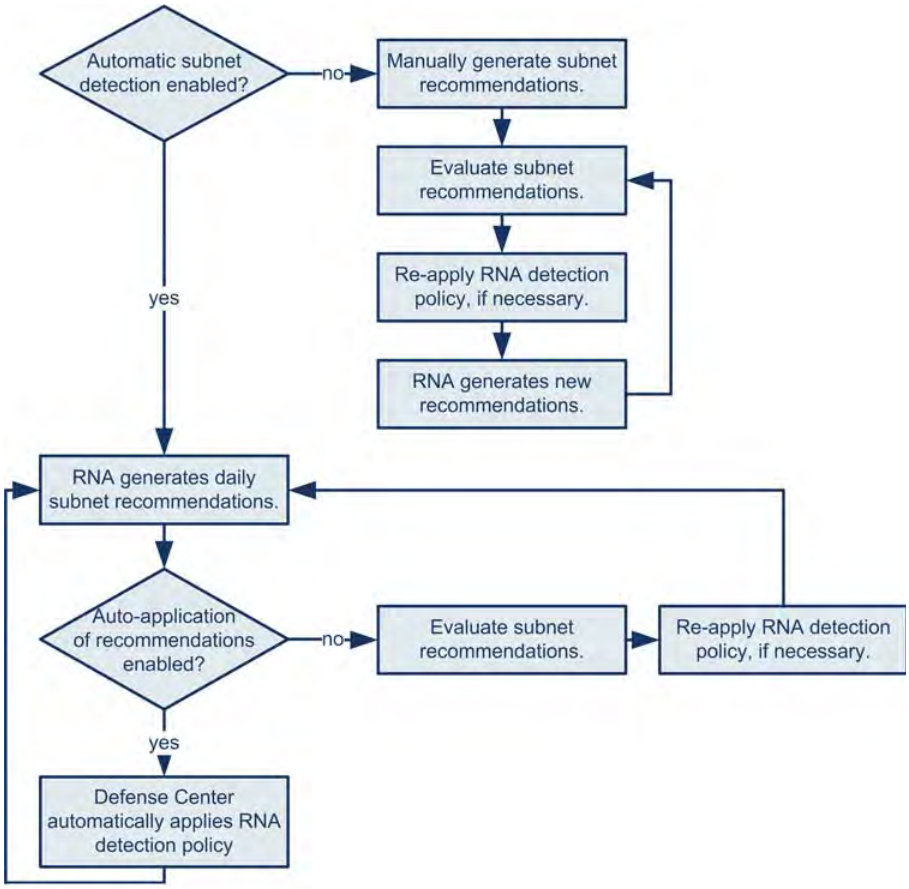
The following diagram outlines the process of managing your Sourcefire 3D System RNA deployment, from physical deployment through configuring and applying an RNA detection policy.



The following diagram illustrates the automated subnet detection process in more detail. Note that you can configure the Defense Center to notify you of subnet recommendations via email so that you can make the changes manually,

Keysight, EX-1004
Page 898

or, if you configured the Defense Center to automatically apply recommendations, to notify you of any changes made.



Keysight, EX-1004
Page 899

The following table explains where to find detailed information on the steps outlined in the above diagrams.

For More Information on Managing Your RNA Deployment

To learn more about...	Look in...
RNA and NetFlow deployment scenarios	• the installation guides for your appliances • Planning your NetFlow Deployment on page 1148
fulfilling prerequisites for your RNA deployment	• What are the Prerequisites for RNA Data Collection? on page 887 • What are the Prerequisites for NetFlow Data Collection? on page 1146
creating and applying an RNA detection policy, including the initial subnet detection configuration	• What is an RNA Detection Policy? on page 889 • Creating RNA Detection Policies on page 900 • Applying an RNA Detection Policy on page 919
managing the iterative subnet detection process	• Managing Subnet Detection on page 911

Creating RNA Detection Policies

Platform: DC

Component: RNA

After you have read and understood the information in What is an RNA Detection Policy? on page 889 and How Do I Manage my RNA Deployment with Subnet Detection? on page 897, you can begin creating an RNA detection policy.

TIP! Instead of creating a new policy, you can export an RNA detection 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 1997.

After you create an RNA detection policy, apply it to your RNA detection engines to start monitoring the hosts in the networks you specified. If you are using subnet detection, applying the policy also allows RNA to begin generating subnet recommendations.

Keysight, EX-1004
Page 900

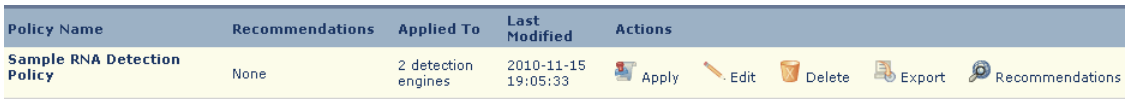
Note that if you did not use the system policy to configure your Defense Center to automatically apply subnet recommendations, you should revisit your detection policy after you apply it for the first time so that you can manually evaluate the recommendations. Explicitly assigning RNA detection engines to monitor your network ensures that RNA collects detailed information about the hosts in your network, including operating system and service identity data, flow data, and so on. RNA gathers only secondary information (hops and MAC address data) about hosts in subnets that are set to auto-detect.

To create an RNA detection policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears, listing any existing detection policies along with those policies' recommendation status, which detection engines the policies are applied to, and the time that the policies were last modified.



2. Click **Create Policy**.

   The Create Policy page appears.

3. Provide basic policy information, such as the policy name and description, and configure the detection policy settings; see Configuring RNA Detection Policy Settings on page 902.

4. Optionally, specify the networks you want to monitor with 3D Sensors; see Specifying Networks to Monitor with 3D Sensors on page 905.

   Note that if you choose not to monitor your network with 3D Sensors, you **must** use NetFlow-enabled devices.

5. Optionally, exclude traffic that travels from or to specific ports from monitoring by RNA; see Specifying Ports to Exclude on page 907.

6. Optionally, specify the networks you want to monitor with NetFlow-enabled devices; see Specifying Networks to Monitor with NetFlow Devices on page 909.

   Note that if you choose not to monitor your network with NetFlow-enabled devices, you **must** use 3D Sensors.

Keysight, EX-1004
Page 901

7. Click **Save Policy**.

   If you set subnets to auto-detect but have not configured the Defense Center to automatically apply subnet recommendations, confirm the message that appears. You must revisit your detection policy after you apply it for the first time so that you can manually evaluate the recommendations.

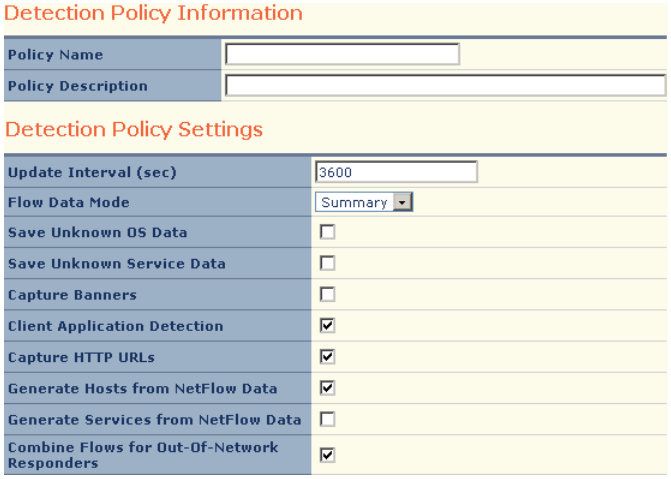
   The RNA detection policy is saved and the Detection Policy page appears. You can now apply the policy to begin monitoring your network and generating subnet recommendations. For more information, see Applying an RNA Detection Policy on page 919 and Managing Subnet Detection on page 911.

Configuring RNA Detection Policy Settings

Platform: DC

Component: RNA

You must give each RNA detection policy a name, and, optionally, a short description. You must also configure the detection policy settings, which help control how RNA collects data.



The following fields comprise the RNA detection policy.

Update Interval

The interval at which RNA updates information (such as when a host was last seen, when a service was used, or the number of hits for a service). 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.

Keysight, EX-1004
Page 902

Flow Data Mode

Indicates how flow data (including NetFlow data) is collected.

- Select **Enabled** to configure your 3D Sensors to transmit individual flows to the Defense Center.
- Select **Summary** (the default) to configure your 3D Sensors to transmit flow summaries, which are sets of flow data aggregated over five-minute intervals; see Understanding Flow Summary Data on page 1076.
- Select **Disabled** to disable the collection of flow data.

Note that disabling or limiting flow data collection can reduce the traffic between your 3D Sensors and your Defense Center as well as reduce the space required to store flow data on the Defense Center, but it can also prevent or limit you from taking advantage of any feature that relies on flow data, such as the Flow Summary dashboard, traffic profiles, compliance rules based on flows or traffic profile changes, and flow trackers in compliance rules.

Save Unknown OS Data

Select this check box if you want to save data related to unidentified operating systems. This data is saved in packet capture (pcap) format on the appliance should you need to send this information to Sourcefire Support.

Save Unknown Service Data

Select this check box if you want to save events related to unidentified services. This data is saved in packet capture (pcap) format on the appliance should you need to send this information to Sourcefire Support.

Capture Banners

Select this check box if you want RNA to store header information from network traffic that advertises service vendors and versions ("banners"). This information can provide additional context to the information gathered by RNA. You can access service banners collected for hosts by accessing service details.

Client Application Detection

Select this check box if you want RNA to generate events related to client application detection. Note that flows exported by NetFlow-enabled devices do not contain any information on client applications involved in the monitored session.

Keysight, EX-1004
Page 903

Capture HTTP URLs

Select this check box if you want RNA to capture HTTP URLs and include them in flow events, where applicable. Note that flows exported by NetFlow-enabled devices do not contain any URL information.

Generate Hosts from NetFlow Data

Select this check box if you want RNA to add hosts to the network map based on flow data exported by NetFlow-enabled devices. You cannot select this check box if you have not applied a NetFlow feature license to your Defense Center.

Note that there is no operating system information available for hosts added to the network map based on NetFlow data, unless you use the host input feature to manually set the hosts' operating system identity.

Generate Services from NetFlow Data

Select this check box if you want RNA to add services to the network map based on flow data exported by NetFlow-enabled devices. You cannot select this check box if you have not applied a NetFlow feature license to your Defense Center.

Note that the Sourcefire 3D System is unable to identify the names, version, and vendors of services running on hosts detected by your NetFlow-enabled devices, unless you use the host input feature to manually set those values.

Combine Flows for Out-of-Network Responders

Select this check box if:

- you kept the flow data mode as **Summary** (that is, your 3D Sensors will transmit aggregated flows instead of individual flows to the Defense Center)
- you want to combine flow summaries involving external hosts on the sensors, **before** they transmit the data to the Defense Center

Enabling this option treats flow summary data from IP addresses that are not in your list of monitored networks as coming from a single host. Event views, graphs, and reports use `external` to indicate the hosts outside your monitored network, instead of an individual IP address.

Your 3D Sensors will combine flow summaries involving a host on your monitored network and one or more external hosts if the flows use the same port, protocol, service, and if they were detected by the same detection engine (for flows detected by 3D Sensor) or were exported by the same NetFlow-enabled device and were processed by the same detection engine.

This can reduce the number of events sent to the Defense Center. It can also reduce the space required to store flow data and can also speed up the rendering of flow data graphs.

Keysight, EX-1004
Page 904

Note that this option has no effect unless the flow data mode is **Summary**, which prevents you from taking advantage of any feature that requires data from individual flows. However, you can use the system policy to configure the Defense Center to compress flow summary data **after** your 3D Sensors transmit it. For more information, see Combining Flow Summaries from External Responders on page 1077 and Configuring RNA Settings on page 1756.

To configure detection policy options:

Access: P&R Admin/Admin

1. On the Create Policy page, in the **Policy Name** field, type a name for the policy.
2. Optionally, in the **Policy Description** field, type a description for the policy.
3. Configure the policy options under Detection Policy Settings.

TIP! If you want to use the settings from an existing RNA detection policy, click **Copy Settings** and, in the pop-up window, select the policy you want to use and click **Load**.

4. Continue with the procedure in the next section, Specifying Networks to Monitor with 3D Sensors.

Specifying Networks to Monitor with 3D Sensors

Platform: DC

Component: RNA

RNA detection policies control which networks are monitored with 3D Sensors with RNA.

If you know which networks the sensing interfaces on your 3D Sensors are physically connected to, configure the RNA detection policy so that the reporting detection engine for a subnet is the detection engine that uses the interface set connected to that subnet. The reporting detection engine for a particular network segment is responsible for reporting most of the information about the hosts on that network, for example, the operating systems and services running on the hosts.

However, if you are not completely sure where and how your 3D Sensors are deployed, use subnet detection to evaluate your network and automatically determine which detection engines should monitor which subnets. For more information, see:

- Monitor the Appropriate Network Segments on page 891
- How Do I Manage my RNA Deployment with Subnet Detection? on page 897
- Managing Subnet Detection on page 911.

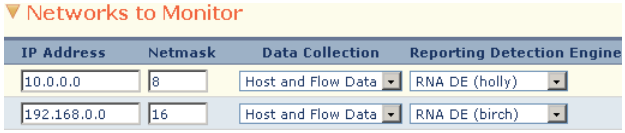
Keysight, EX-1004
Page 905

For each network you monitor, you must also set a data collection mode. You can:

- collect data that is used to populate host profiles for the hosts in the network as well as flow data for those hosts
- collect only host data for the hosts in the network (the default)
- collect neither host data nor flow data for the hosts in the network

In other words, when RNA detects a flow between two hosts, the information that RNA records in the database depends on the data collection mode you set for the networks where the hosts reside. For more information, see Filter Flows as Needed on page 893.

The following graphic shows an RNA detection policy where the reporting detection engine for the 10.0.0.0/8 network resides on `holly` and the reporting detection engine for the 192.168.0.0/16 network resides on `birch`. RNA is configured to collect host and flow data for both subnets.



To specify networks to monitor with 3D Sensors:

Access: P&R Admin/Admin

1. Under Networks to Monitor, click **Add Network**.
2. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in CIDR notation) that represents the network segment you want to monitor or exclude from monitoring.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
3. Select a data collection mode for the network segment.
   - Select **Host and Flow Data** to collect both host data and flow data.
   - Select **Host Data Only** to collect only host data.
   - Select **Exclude** to exclude the network from monitoring.

Keysight, EX-1004
Page 906

4. Under **Reporting Detection Engine**, select the detection engine that you want to monitor this network:
   - Select a specific detection engine if you are sure where its 3D Sensor is deployed and therefore you know that the detection engine is closest to the network you want to monitor.
   - Leave the default, **Auto-detect**, if you want RNA to recommend subnet assignments for your detection engines.

   Note that if your deployment only includes one RNA detection engine, you must select it as the reporting detection engine. You cannot select **Auto-detect**.

   WARNING! Do **not** select **Auto-detect** if your deployment includes load-balanced Crossbeam-based software sensors in multi-member RNA VAP groups. RNA subnet detection is **not** supported for that case. For more information, see Special Considerations: Crossbeam-Based Software Sensors on page 896.

5. To add additional networks, repeat steps 1 through 4.

   TIP! To remove a network, click **Delete** next to the network you want to remove.

6. Continue with the procedure in the next section, Specifying Ports to Exclude on page 907, to (optionally) exclude traffic that travels from or to specific ports from monitoring by RNA.

Specifying Ports to Exclude

Platform: DC

Component: RNA

You can configure the RNA detection policy to exclude ports from monitoring for specific IP addresses. For example, load balancers can exhibit multiple services on the same port in a short period of time. Similarly, you might want to exclude TCP traffic on port 80 on a load balancer that handles a web farm.

IMPORTANT! The port exclusion feature only affects data collected by 3D Sensors with RNA. You cannot configure NetFlow-enabled devices to exclude ports from monitoring.

Keysight, EX-1004
Page 907

To exclude ports from monitoring:

Access: P&R Admin/Admin

1. Under Ports to Exclude, click **Add Port**.
2. 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.
3. In the **Protocol** field, specify the protocol of the traffic you want to exclude.

   You can exclude either TCP or UDP traffic.
4. In the **Src/Dest Port** field, specify how you want to exclude traffic based on its direction:
   - Select **Source** to exclude traffic on the port where the hosts you specify are the source of the traffic.
   - Select **Destination** to exclude traffic on the port where the hosts you specify are the destination of the traffic.
   - Select **Source/Destination** to exclude all traffic on the port for the hosts you specify.
5. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in CIDR notation) that represents the network segment where you want to exclude port traffic.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

   WARNING! If you use the port exclusion feature to exclude from monitoring the ports on a NetFlow-enabled devices (that is, you enter the IP address of one of your NetFlow-enabled devices in the **IP Address** field), RNA will not process any data exported by that device.
6. To exclude additional ports, repeat steps 1 through 5.

   TIP! To remove a port, click **Delete** next to the port you want to remove.
7. Continue with the procedure in the next section, Specifying Networks to Monitor with NetFlow Devices on page 909, to (optionally) specify the networks you want to monitor with NetFlow-enabled devices.

Keysight, EX-1004
Page 908

Specifying Networks to Monitor with NetFlow Devices

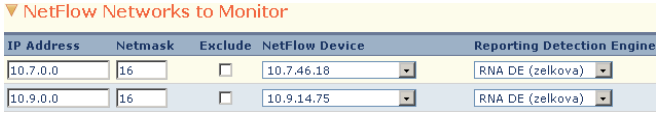
Platform: DC

Component: RNA + NetFlow

RNA detection policies control which networks are monitored with NetFlow-enabled devices.

Because subnet detection is not supported for NetFlow monitoring, you must explicitly configure at least one RNA detection engine to analyze the data exported by your NetFlow-enabled devices. The 3D Sensor where the detection engine resides must be deployed such that one of its sensing interfaces is connected to the destination network where your NetFlow-enabled device exports flow records. Otherwise, the detection engine will have no NetFlow data to process and will not transmit any NetFlow-based flows to the Defense Center for your analysis. The detection engine that analyzes the NetFlow data for a network segment is the reporting detection engine for that segment.

For example, the following graphic shows that the RNA detection engine on the sensor named `zelkova` is the reporting detection engine for the 10.7.0.0/16 and 10.9.0.0/16 subnets, which are being monitored by two separate NetFlow-enabled devices: 10.7.46.18 and 10.9.14.75, respectively.



Note that on the Defense Center, you cannot add NetFlow networks to monitor unless you have first applied a NetFlow feature license to your Defense Center and used the system settings to specify the NetFlow-enabled devices you plan to use. (On the Master Defense Center, the NetFlow-enabled devices you added to your managed Defense Centers appear in the **NetFlow Device** drop-down list.)

For more information on using NetFlow-enabled devices to monitor your network, including the prerequisites for NetFlow data collection and the differences between the NetFlow data and the flow data detected directly by detection engines on 3D Sensors with RNA, see Introduction to NetFlow on page 1140.

To specify networks to monitor with NetFlow-enabled devices:

Access: P&R Admin/Admin

1. Under NetFlow Networks to Monitor, click **Add Network**.
2. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in CIDR notation) that represents the network segment you want to monitor or exclude from monitoring.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
3. To exclude the network from monitoring, select **Exclude**.

Keysight, EX-1004
Page 909

4. Under **NetFlow Device**, select the NetFlow-enabled device that you want to use to monitor the network you specified.

   The NetFlow-enabled devices listed are those you specified in the system settings.

5. Under **Reporting Detection Engine**, select the RNA detection engine that you want to use to analyze the data exported by the NetFlow-enabled device you chose.

   You must make sure that the 3D Sensor where the detection engine resides can see the data exported by the NetFlow-enabled device.

6. To add additional networks, repeat steps 1 through 5.

   TIP! To remove a network, click **Delete** next to the network you want to remove.

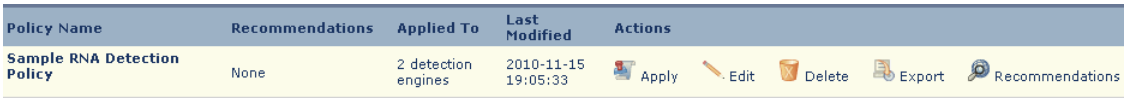
7. Continue with step 7 of the procedure in Creating RNA Detection Policies on page 900 to save the policy.

Managing RNA Detection Policies

Platform: DC

Component: RNA

Manage RNA detection policies using the Detection Policy page. You can create, modify, delete, export, and apply individual policies, as well as view subnet recommendations.



The page display how many detection engines each policy is applied to, as well as the time that the policy was last modified. In addition, the recommendations status indicates whether RNA has subnet recommendations for the policy or if RNA is in the process of generating new recommendations.

For more information, see the following sections:

- Managing Subnet Detection on page 911
- Applying an RNA Detection Policy on page 919
- Editing an RNA Detection Policy on page 921
- Deleting an RNA Detection Policy on page 921

Keysight, EX-1004
Page 910

Managing Subnet Detection

Platform: DC

Component: RNA

Optimally, your RNA detection policy specifies that each RNA detection engine is configured as the reporting detection engine for the hosts that are closest to it.

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 can make it challenging to stay on top of proper RNA policy configurations. Subnet detection allows RNA to make recommendations about which are the best detection engines to analyze the traffic on the various network segments in your organization.

As RNA continuously monitors your network, it may be able to refine its subnet recommendations, especially if your network configuration has been altered through routing or host changes. Choosing which subnets to monitor with which detection engines is an iterative process that you can either revisit from time to time, or configure the Sourcefire 3D System to automatically handle. For information on this process, see How Do I Manage my RNA Deployment with Subnet Detection? on page 897.

Note that subnet detection is **not** supported in the following cases:

- if your deployment includes only one RNA detection engine; with only one RNA detection engine, the only choice is automatically the best choice
- if your deployment includes load-balanced Crossbeam-based software sensors in multi-member RNA VAP groups
- for networks you are monitoring with NetFlow-enabled devices

For more information, see the following sections:

- Understanding RNA-Generated Subnet Recommendations on page 911
- Manually Generating Subnet Recommendations on page 915
- Manually Accepting or Rejecting Subnet Recommendations on page 916
- Automatically Generating and Applying Subnet Recommendations on page 918

Understanding RNA-Generated Subnet Recommendations

Platform: DC

Component: RNA

A subnet recommendation is RNA's determination of which detection engine you should use to monitor a specific network, based on the proximity of the hosts on your monitored network to that detection engine.

RNA makes recommendations for subnet mask lengths of 0, 4, 8, 12, 16, 20, 24, and 28 bits, as well as for any subnet mask lengths that you provide. For example, if you configure RNA to provide recommendations for the 10.0.0.0/11 supernet,

Keysight, EX-1004
Page 911

RNA can distribute the monitoring of that network among multiple detection engines; each of the smaller subnets will have a subnet mask length of 11, 12, 20, 24, or 28.

RNA generates subnet recommendations at the following times:

- when you apply an RNA detection policy that has at least one monitored network set to auto-detect
- when you manually generate recommendations for a specific detection policy
- if you configured automatic daily recommendation generation, at the time you specified in the system policy

RNA provides three different kinds of recommendation, as described in the following table.

RNA Subnet Detection Recommendation Types

Recommendation Type	Description
New	New recommendations indicate that you should monitor new subnetworks that are not explicitly specified in your detection policy. For example, if you specify in your detection policy that you want to monitor the 10.0.0.0/8 network, but your detection engines only see traffic from the 10.8.0.0/16 network, RNA can issue a new recommendation for the explicit monitoring of 10.8.0.0/16.
Modified	Modified recommendations indicate that you should change the reporting detection engine for a network that you are explicitly monitoring in your detection policy. For example, consider a situation where you want to use auto-detection to evaluate which detection engines to use to monitor the 10.0.0.0/8 network. If a detection engine sees traffic from that network, RNA can issue a modified recommendation that you switch from auto-detection to explicit monitoring of that entire network. As another example, RNA may initially recommend that you split the monitoring of a supernet among several detection engines, then later determine that using a single detection engine is best.
Revert to Auto-Detection	RNA recommends that you revert to auto-detection if RNA is not seeing sufficient network traffic using the currently assigned detection engine. Reverting to auto-detection allows RNA to restart the subnet detection process for those networks.

Keysight, EX-1004
Page 912

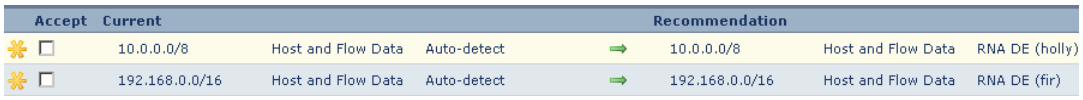
For each recommendation that RNA provides, you have two options:

- accept the recommendation, which incorporates the change into your RNA detection policy
- reject the recommendation, which does not incorporate the change into your detection policy

Note that even if you choose to reject some of RNA's recommendations, RNA may present the same recommendations the next time they are generated. In other words, you can never dismiss recommendations entirely.

As a simple example of subnet recommendations, consider a pair of 3D Sensors named `holly` and `fir`, where `holly`'s sensing interface is physically connected to the 10.0.0.0/8 network, and `fir`'s sensing interface is physically connected to the 192.168.0.0/16 network.

If you set both those networks to auto-detect in the RNA detection policy (rather than configuring `holly`'s RNA detection engine to monitor 10.0.0.0/8 and `fir`'s to monitor 192.168.0.0/16), RNA would generate the recommendations shown in the following graphic.



Note that recommendations marked with a yellow asterisk ( ) are recommendations that you have not evaluated. That is, you have neither accepted nor rejected that recommendation.

As another example, consider needing to monitor the 10.0.0.0/8 supernet. You have two 3D Sensors with RNA (`fir` and `holly`), but you are not sure where exactly they are deployed, so you decide to use subnet detection. RNA can then provide a list of smaller subnets that are closest, as shown in the following graphic.



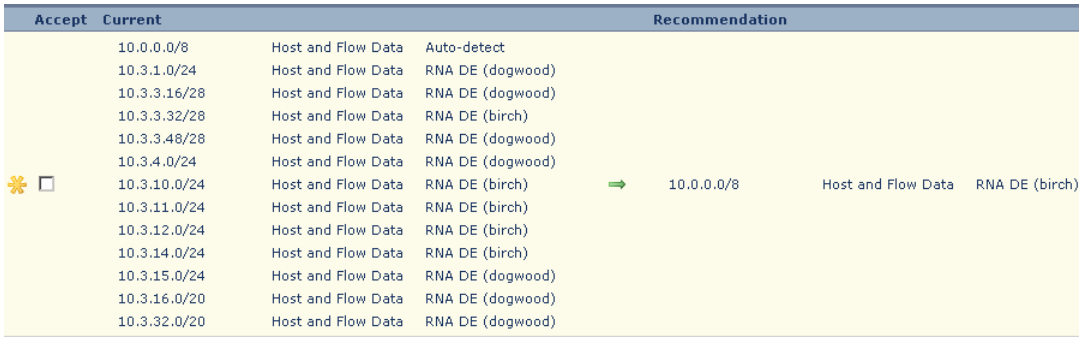
These recommendations say:

- the 10.3.0.0/16 network should be monitored by the detection engine on `holly`
- the 10.6.0.0/16 network should be monitored by the detection engine on `fir`

Keysight, EX-1004
Page 913

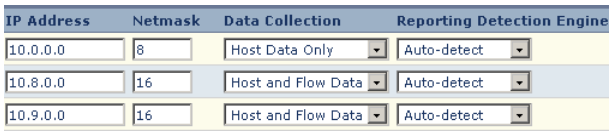
- the 10.9.0.0/20 network can be monitored by your choice of the detection engines on either sensor; when RNA presents you with a choice of detection engines to monitor a subnet, it means that RNA has determined that both detection engines are equally close to the subnet
- the rest of the 10.0.0.0/8 network remains as **Auto-detect** because RNA was not able to gather enough data to recommend a detection engine

As RNA continues to monitor network traffic, it may be able to refine its recommendations, especially if your network configuration has been altered through routing or host changes. For example, RNA may initially recommend that you split the monitoring of a network among several detection engines, then later determine that using a single detection engine is best, as shown in the following recommendation.

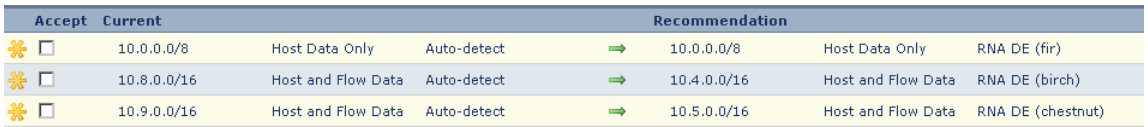


Note that for your convenience, the recommendations display the data collection (flow filtering) mode you configured for each network. Knowing the data collection mode that applies to a particular recommendation can help you distinguish among and evaluate recommendations for more complex RNA deployments.

For example, you may configure your detection policy so that you enable flow collection for subnets that are most critical to your operations, while collecting host data only for your entire network, as shown in the following graphic.



RNA then recommends that you use the detection engine on one sensor to monitor the 10.0.0.0/8 supernet and collect host data only, but use the detection engines on different sensors to monitor the 10.8.0.0/16 and 10.9.0.0/16 subnetworks, collecting host and flow data.



Keysight, EX-1004
Page 914

For information on how to accept or reject subnet recommendations, see Manually Accepting or Rejecting Subnet Recommendations on page 916. You can also configure RNA to automatically apply recommendations; see Automatically Generating and Applying Subnet Recommendations on page 918.

Manually Generating Subnet Recommendations

Platform: DC

Component: RNA

You can manually generate subnet recommendations at any time, with the following exceptions:

- The policy for which you want to generate recommendations must be applied to at least one RNA detection engine. You cannot generate recommendations for policies that are not in use.
- You cannot generate recommendations for a policy if RNA is already in the process of generating recommendations for that policy. Both the task queue (**Operations > Monitoring > Task Status**) and the Detection Policy page (**Policy & Response > RNA > Detection Policy**) indicate whether recommendations are currently being generated.

Note that the Defense Center does **not** automatically apply manually generated recommendations, even if you configured your system policy to automatically generate and apply recommendations on a daily basis. The automatic process only occurs once per day at the time you specified. After you manually generate recommendations, you must then evaluate and, optionally, apply them.

To manually generate subnet recommendations:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears.

2. Next to the policy for which you want to generate recommendations, click **Recommendations**.

   The Detection Engine Auto-Detection Recommendations page appears.

3. Click **Regenerate Recommendations**.

   RNA begins generating recommendations. You can monitor the process in the task queue (**Operations > Monitoring > Task Status**).

4. After RNA finishes generating recommendations, continue with the next section, Manually Accepting or Rejecting Subnet Recommendations, to evaluate and, optionally, apply them.

Keysight, EX-1004
Page 915

Manually Accepting or Rejecting Subnet Recommendations

Platform: DC

Component: RNA

If you have not configured RNA to automatically apply subnet recommendations (see Automatically Generating and Applying Subnet Recommendations on page 918), you must evaluate and apply the recommendations manually.

Applying recommendations is especially important if you have created and applied a new RNA detection policy that uses subnet detection. Because RNA gathers only secondary information (hops and MAC address data) about hosts in subnets that are set to auto-detect, until you explicitly assign an RNA detection engine to monitor each subnet and reapply the policy, RNA will not gather detailed information about the hosts in that network, including operating system and service identity data, flow data, and so on.

TIP! Because choosing which subnets to monitor with which detection engines is an iterative process (see How Do I Manage my RNA Deployment with Subnet Detection? on page 897), you may want to revisit subnet recommendations from time to time to adjust your detection engine assignments.

As described in the procedure in this section, you can accept or reject individual recommendations on the Detection Engine Auto-Detection Recommendations page by selecting the check boxes next to each recommendation you want to accept, and by leaving blank the check boxes next to the recommendations you want to reject. After you select the recommendations you want to implement, update your RNA detection policy by clicking **Accept**, then reapply the detection policy.

Manually rejecting recommendations ensures that the Defense Center will not automatically apply the rejected recommendations, even if you if you configured the Defense Center to automatically generate and apply daily subnet recommendations.

For example, consider a scenario where you are monitoring the 10.0.0.0/8 network with a single detection engine, you manually generate recommendations, and RNA recommends that you split monitoring of the

Keysight, EX-1004
Page 916

network amongst multiple detection engines. Rejecting the recommendation has the following consequences:

- RNA will not apply the rejected recommendation the next time it automatically generates and applies recommendations.
- RNA will not automatically apply any future recommendations to split monitoring of the 10.0.0.0/8 network, even if the new recommendations are not the same as the original split recommendation. That is, even if the new recommendations include different detection engines or more or fewer subnets, RNA will not automatically split up the 10.0.0.0/8 network.

  If you want to split monitoring of the network after you manually reject the recommendation to split, you must either edit your RNA detection policy or manually accept a new recommendation to split.

Note, however, that you cannot dismiss recommendations entirely. Even if you reject some of RNA's recommendations, RNA may present the same recommendations the next time it generates them even if it will not automatically apply them.

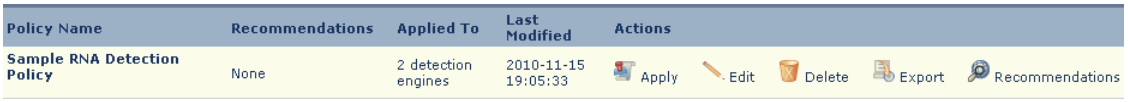
IMPORTANT! Rejected recommendations are those that you did not select before clicking **Accept**. Although clicking **Cancel** rather than **Accept** exits the evaluation process, the Defense Center does not consider any recommendations to be either accepted or rejected. In other words, if you configured the Defense Center to automatically apply daily subnet recommendations, recommendations you thought you rejected may be included in the RNA detection policy the next time the Defense Center automatically applies it, depending on RNA's view of your network at the time of the auto-apply.

To accept or reject subnet recommendations:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears.



   The recommendations status indicates whether RNA has subnet recommendations for the policy or if RNA is in the process of generating new recommendations.

2. Click **Recommendations** next to an RNA detection policy that has new recommendations.

   The Detection Engine Auto-Detection Recommendations page appears, displaying RNA's subnet recommendations for the policy. For information on how to interpret the recommendations, see Understanding RNA-Generated Subnet Recommendations on page 911.

Keysight, EX-1004
Page 917

3. Under **Accept**, select the check boxes next to the recommendations you want to accept. To reject all of RNA's recommendations, leave the check boxes cleared.
4. For each recommendation that prompts you to select a detection engine for a subnet, select a detection engine from the drop-down list.

   Make sure to also select the **Accept** check box next to the recommendations.
5. Click **Accept**.

   The RNA detection policy is saved with the changes you made and the Detection Policy page appears again. Note that if you want to reject all of RNA's recommendations, you still must click **Accept**.
6. Click **Apply** to reapply the policy to the detection engines that use it.

   See Applying an RNA Detection Policy on page 919 for more information.

Automatically Generating and Applying Subnet Recommendations

Platform: DC

Component: RNA

As RNA continuously monitors your network traffic, it may be able to refine any subnet recommendations it has made for your RNA detection policies, especially if your network configuration has been altered through routing or host changes. Choosing which subnets to monitor with which detection engines is an iterative process, as you revisit the subnet recommendations from time to time and then adjust your detection engine assignments.

As a time-saving and performance-maximizing measure, you can use the system policy to configure RNA to automatically generate subnet recommendations for your currently applied RNA detection policies on a daily basis.

Optionally, you can configure the Defense Center to automatically update those policies and apply the updated policies to your RNA detection engines. Note that during an automatic policy update and apply, if RNA determines that multiple detection engines are equally appropriate to monitor a particular subnet, RNA chooses the detection engine whose name is first alphabetically. For more information, see How Do I Manage my RNA Deployment with Subnet Detection? on page 897.

Note that you cannot apply an RNA detection policy to the combined detection engines on a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). Auto-apply does not occur for clustered pairs with mismatched versions.

To configure RNA subnet detection settings:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

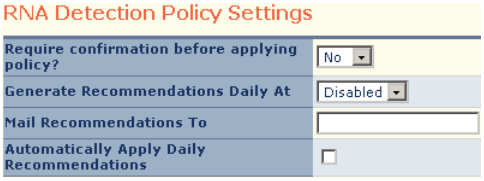
Keysight, EX-1004
Page 918

2. You have two options:
   - To modify the RNA subnet detection settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the RNA subnet detection 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 1733, and click **Save**.

   In either case, the Access List page appears.
3. Click **RNA Settings**.

   The RNA Settings page appears. You manage subnet detection settings in the RNA Detection Policy Settings section.



4. From the **Generate Recommendations Daily At** drop-down list, select the time when you want RNA to automatically generate daily subnet recommendations for all applied RNA detection policies.

   To disable daily generation of subnet recommendations, select **Disabled**.
5. Optionally, in the **Mail Notifications To** field, enter the email address where you want to receive notifications of new subnet recommendations.

   TIP! To receive email notifications, you must configure a valid mail relay host; see Configuring a Mail Relay Host and Notification Address on page 1751.
6. Select the **Automatically Apply Daily Recommendations** check box to automatically update and apply your RNA detection policies after RNA generates subnet recommendations.

   Note that this option has no effect unless you enable daily recommendations.
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 1736.

Applying an RNA Detection Policy

Platform: DC

Component: RNA

Applying an RNA detection policy allows your RNA detection engines to begin monitoring your network, and, optionally, to begin analyzing data exported by

Keysight, EX-1004
Page 919

NetFlow-enabled devices. It also allows RNA to generate recommendations for subnets that you set to auto-detect.

When you apply an RNA detection policy:

- RNA deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks
- the 3D Sensors where the affected detection engines reside discard any RNA data that has not yet been sent to the Defense Center

Note that you cannot apply an RNA detection policy to the combined detection engines on a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails).

IMPORTANT! If your Defense Center is not configured to automatically apply subnet recommendations, you should revisit your detection policy after you apply it for the first time so that you can manually evaluate the recommendations. Explicitly assigning RNA detection engines to monitor your network ensures that RNA collects detailed information about the hosts in your network, including operating system and service identity data, flow data, and so on. RNA gathers only secondary information (hops and MAC address data) about hosts in subnets that are set to auto-detect.

To apply an RNA detection policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears.

2. Click **Apply** next to the RNA detection policy that you want to apply.

   The Apply Policy page appears, listing the RNA detection engines in your deployment. You can list the detection engines by detection engine group, sensor, current policy, interface set type, or detection engine type.

3. Under Available Detection Engines, select the detection engines where you want to apply the policy, and click **Apply**.

   If you enabled the **Require confirmation before applying policy?** option in the system policy, you must confirm that you want to apply the detection policy. For more information, see Configuring RNA Settings on page 1756.

   The Defense Center also warns you if the detection engines have different policies applied to them than the one you are attempting to apply. You can determine which RNA detection policy is currently applied by viewing the Available Detection Engines page (**Operations > Configuration > Detection Engines**).

   After you dismiss any warnings, the policy is applied and the Detection Policy page appears.

Keysight, EX-1004
Page 920

Editing an RNA Detection Policy

Platform: DC

Component: RNA

You can edit an RNA detection policy, but you must be logged into the web interface where you created the policy.

To edit an existing RNA detection policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears.

2. Click **Edit** next to the RNA detection policy that you want to modify.

   The Create Policy page appears.

3. Modify the policy to suit your needs.

   See Creating RNA Detection Policies on page 900 and Managing Subnet Detection on page 911 for more information.

4. Click **Update Policy**.

   If any subnets are set to auto-detect but you have not configured the Defense Center to automatically apply subnet recommendations, confirm the message that appears. You must revisit your detection policy after you apply it so that you can manually evaluate the recommendations.

   The RNA detection policy is saved and the Detection Policy page appears again.

IMPORTANT! Your changes to the policy do not take effect until you reapply the policy to the detection engines that use it. See Applying an RNA Detection Policy on page 919 for more information.

Deleting an RNA Detection Policy

Platform: DC

Component: RNA

You can delete an RNA detection policy that you no longer want to use.

IMPORTANT! You can delete a policy that is currently in use, so make sure you first apply a different policy to any detection engines that use it.

Keysight, EX-1004
Page 921

To delete an existing RNA detection policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears.

2. Click **Delete** next to the RNA detection policy that you want to remove and confirm that you want to delete the policy.

   The policy is deleted.

Keysight, EX-1004
Page 922

Chapter 22
Using the Network Map

3D Sensors with the RNA component passively collect traffic traveling over the network, decode data, and then compare it to established operating system and service fingerprints. As RNA collects this information, it builds a *network map*, which is a detailed representation of your network. If your organization has the appropriate scripting capabilities, you can augment the data collected by RNA by adding operating system, service, client application, protocol, or host attribute information from a third party application through the host input feature or by editing it through the Defense Center web interface. You can actively scan hosts in the network map using Nmap and add the scan results to your network map.

The network map allows you to use the Defense Center to view your network topology in terms of the hosts and network devices (bridges, routers, NAT devices, and load balancers) running on your network as well as their associated host attributes, services, and vulnerabilities. In other words, you can select different views of the network map depending on the kind of analysis you are performing.

You can also use the custom topology feature to help you organize and identify subnets in the hosts and network devices views of the network map. For

Keysight, EX-1004
Page 923

example, if each department within your organization uses a different subnet, you can label those subnets using the custom topology feature.

TIP! The network map is a useful tool for a quick, overall view of your network. For more detailed information on RNA events, as well as hosts and their attributes, the services and client applications running on your network, vulnerability data, and sessions sustained by monitored hosts (flow data), see Working with RNA Events on page 994 and Working with Flow Data and Traffic Profiles on page 1071.

For more information, see the following sections:

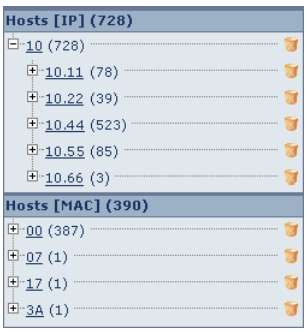
- Understanding the Network Map on page 924
- Working with the Hosts Network Map on page 926
- Working with the Network Devices Network Map on page 928
- Working with the Services Network Map on page 930
- Working with the Vulnerabilities Network Map on page 932
- Working with the Host Attributes Network Map on page 935
- Working with Custom Network Topologies on page 936

Understanding the Network Map

Platform: DC

Component: RNA

Each view of the network map has the same format: a hierarchical tree with expandable categories and sub-categories. When you click one of the top-level categories, it expands to show you the 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 the 3D Sensors with RNA that it manages (including those that process NetFlow data). If multiple 3D Sensors

Keysight, EX-1004
Page 924

generate events for the same host or service, the Defense Center combines the information from each sensor into a composite representation of that host or service.

IMPORTANT! Although you can configure the RNA detection policy to add hosts and services to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts and services 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 What Information Does RNA Provide? on page 880.

From any network map, you can view any host's *host profile*, which provides a complete view of all the information collected by RNA for a single host. The host profile contains general host information, such as the host name and operating system, as well as more specific information including RNA-detected protocols, services, and client applications that are running on the host. For more information on host profiles, see Using Host Profiles on page 944.

You can delete an item from the network map if you are no longer interested in investigating it. You can delete hosts and services from the network map; you can also delete or deactivate vulnerabilities.

Note that if RNA detects activity associated with the host after you delete it from the network map, it re-adds the host to the network map. Similarly, deleted services are re-added to the services network map if RNA detects a change in the service (for example, if an Apache web server is upgraded to a new version). Vulnerabilities are reactivated on specific hosts if RNA detects a change in the host's vulnerable operating system or service.

You can also use the network map to deactivate vulnerabilities network-wide, which means that you deem that none of the hosts that RNA has detected on your network are vulnerable to that particular attack or exploit.

TIP! If you want to permanently exclude a host or subnet from the network map, modify the RNA detection policy. See What is an RNA Detection Policy? on page 889 for more information.

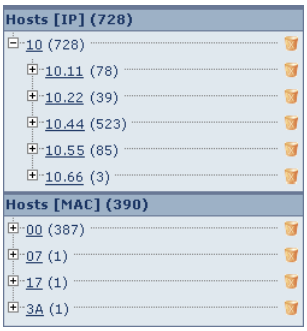
Keysight, EX-1004
Page 925

Working with the Hosts Network Map

Platform: DC

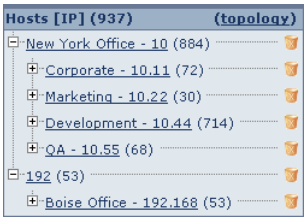
Component: RNA

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.



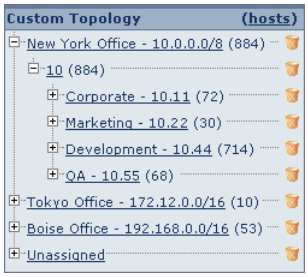
IMPORTANT! Although you can configure the RNA detection policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts is limited. For example, there is no operating system data available for hosts added to the network map using NetFlow data, unless you provide it using the host input feature. In addition, note that without a RNA host license applied to the Defense Center, RNA will not build the network map or store information about the hosts on your monitored network, even if you are using only NetFlow-enabled devices to monitor your network. For more information, see What Information Does RNA Provide? on page 880 and Managing Feature Licenses on page 1787.

Note that if you create a custom topology for your network, the labels you assign to your subnets appear in the hosts network map.



Keysight, EX-1004
Page 926

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 936.



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.

Note that if RNA detects activity associated with the host after you delete it from the network map, 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 RNA detection policy. See What is an RNA Detection Policy? on page 889 for more information.

IMPORTANT! Sourcefire **strongly** recommends that you do **not** delete network devices from the network map, because RNA 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: Any RNA/Admin

1. Select **Analysis & Reporting > RNA > Network Map > Hosts**.

   The Hosts page appears, displaying a list of host IP addresses and, for those hosts without IP addresses, 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 944.

Keysight, EX-1004
Page 927

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. Note that if RNA rediscovers the host, it is re-added 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.

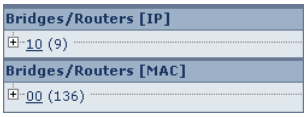
   For information on configuring custom topologies, see Working with Custom Network Topologies on page 936.

Working with the Network Devices Network Map

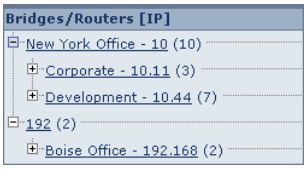
Platform: DC

Component: RNA

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 that do not have an IP address and are identified only by a MAC address.



Note that if you create a custom topology for your network, the labels you assign to your subnets appear in the network devices network map.



Keysight, EX-1004
Page 928

The methods RNA 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 an IP address. If it communicates using STP, it may only have a MAC address.

You cannot delete network devices from the network map because RNA uses their locations to determine network topology (including generating network hops and TTL values for monitored hosts).

To view the network devices network map:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > RNA > Network Map > Network Devices**.

   The Network Devices page appears, displaying a list of network device IP addresses and, for those network devices without IP addresses, 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 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 944.

Keysight, EX-1004
Page 929

Working with the Services Network Map

Platform: DC

Component: RNA

Use the services network map to view the services on your network, organized in a hierarchical tree by service name, vendor, version, and finally by the hosts running each service.



For a full list of the services that the Sourcefire 3D System detects, see Detected Services on page 2027. Keep in mind, however, that the detector list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed, as well as on any individual add-on detectors you may have imported. 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 RNA detectors.

IMPORTANT! Although you can configure the RNA detection policy to add services to the network map based on data exported by NetFlow-enabled devices, the available information about these services is limited. For more information, see What Information Does RNA Provide? on page 880.

From the services network map, you can view the host profile of each host that runs a specific service as well as delete any service category, any service running on all hosts, or any service running on a specific host. For example, you can delete a service from the network map if you know it is disabled on the host and you want to make sure it is not used for impact flag qualification.

Note that deleting a service from the network map does not remove it from your network. A deleted service reappears in the network map if RNA detects a

Keysight, EX-1004
Page 930

change in the service (for example, if an Apache web server is upgraded to a new version) or if you restart RNA.

Depending on what you delete, the behavior differs:

- If you delete a service category, the service category is removed from the network map. All services that reside beneath the category are removed from any host profile that contains the services.

  For example, if you delete **http**, all services identified as `http` are removed from all host profiles and **http** no longer appears in the Services view of the network map.

- If you delete a specific service, vendor, or version, the affected service 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 services 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 service list and the service 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) service is deleted from the host profile of IP address 172.16.1.50.

To view the services network map:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > RNA > Network Map > Services**.

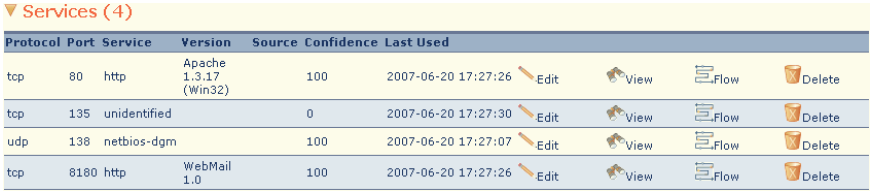
   The Services page appears.

2. Drill down to the specific service 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.

3. Click a specific IP address under the service you selected.

   The host profile of the host running the service appears with the Services section expanded. For more information about the Services section of the host profile, see Working with Services in the Host Profile on page 958.



Keysight, EX-1004
Page 931

4. Optionally, to delete any service category, any service running on all hosts, or any service 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 service is deleted. Note that if RNA rediscovers the service, it is re-added to the network map.

Working with the Vulnerabilities Network Map

Platform: DC

Component: RNA

Use the vulnerabilities network map to view the vulnerabilities that RNA has detected on your network, organized by RNA ID, 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

Keysight, EX-1004
Page 932

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 984 for more information.

If there is an identity conflict for a service or operating system on a host, RNA treats the host as potentially affected by vulnerabilities for both potential identities. Resolving the identity conflict to one current identity reduces the vulnerabilities to only those associated with that identity. For more information, see Understanding Current Identities on page 1360 and Understanding Identity Conflicts on page 1362.

If a user with the Administrator role maps vulnerabilities for a specific service in the system policy, RNA lists vulnerabilities for services that do not send vendor or version information. If a service is not mapped, RNA lists the service in the host profile for the host, but does not list any of the vulnerabilities associated with that service in the host profile. When RNA detects vendor and version information for a service, vulnerabilities are listed for that service unless deactivated. For more information, see Mapping Vulnerabilities for Services on page 1773.

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. In other words, for each host that is affected by each vulnerability in the range, the count is incremented by one. If a host is affected by more than one vulnerability, it is counted multiple times. Therefore, it is possible that the count can 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 RNA has determined are *potentially* affected by a vulnerability or vulnerabilities.

Note that any host that was not affected when a vulnerability was deactivated but that becomes affected due to an operating system or service change shows the vulnerability as valid in its host profile. In addition, the Defense Center no longer uses deactivated vulnerabilities for intrusion impact correlation.

To view the vulnerabilities network map:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > RNA > Network Map > Vulnerabilities**.

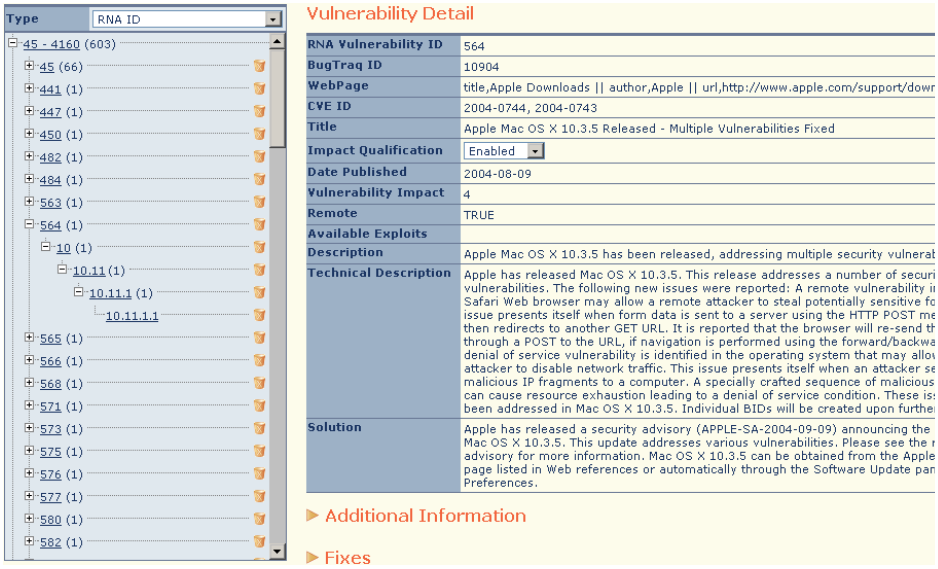
   The Vulnerabilities page appears.

2. From the **Type** drop-down list, select the class of vulnerability you want to view. By default, vulnerabilities are displayed by RNA ID.

Keysight, EX-1004
Page 933

3. Drill down to the specific vulnerability you want to investigate.

   The vulnerability details appear. For details on the information provided, see Viewing Vulnerability Details on page 980.



   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 984 for more information on reactivating vulnerabilities.

Keysight, EX-1004
Page 934

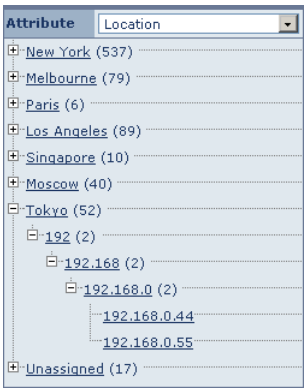
Working with the Host Attributes Network Map

Platform: DC

Component: RNA

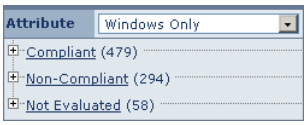
Use the host attributes network map to view the hosts on your network organized by their host attributes. When you select the host attribute you want to use to organize your hosts, the Defense Center lists the possible values for that attribute in the network map and groups hosts based on their assigned 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 987.

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

Keysight, EX-1004
Page 935

For more information on compliance white lists, see Using RNA as a Compliance Tool on page 1156.

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: Any RNA/Admin

1. Select **Analysis & Reporting > RNA > Network Map > Host Attributes**.
   The Host Attributes page 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.
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.

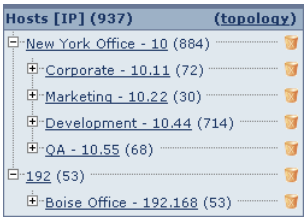
Working with Custom Network Topologies

Platform: DC

Component: RNA

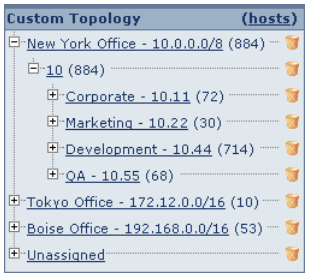
Use the custom topology feature to help you organize and identify subnets in your hosts and network devices network maps.

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.



Keysight, EX-1004
Page 936

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 926 and Working with the Network Devices Network Map on page 928.

When you create a custom topology, you must identify the networks in your organization. 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 Sourcefire 3D System has detected
- by importing networks from an RNA detection policy, which adds the networks that you configured the Sourcefire 3D System to monitor in an RNA detection policy
- by adding networks to your topology manually, if the other two methods create an inaccurate or incomplete representation of your deployment

After you identify your networks, you can name them so that when you view the network map, the organization imparted by the custom topology is meaningful to you. For more information, see the following sections:

- Creating Custom Topologies on page 937
- Managing Custom Topologies on page 942

Creating Custom Topologies

Platform: DC

Component: RNA

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
- by importing networks from an RNA detection policy
- by adding networks to your topology manually

You must save and activate the topology to start using it with the network map.

Keysight, EX-1004
Page 937

To create a custom topology:

Access: Admin

1. Select **Policy & Response > RNA > Network Map > Custom Topology**.

   The Custom Topology page appears.

2. Click **Create Topology**.

   The Edit Topology appears.

3. Provide basic topology information, such as the topology name and description.

   See Providing Basic Topology Information on page 939.

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 939.
   - To add networks to your topology by importing them from an RNA detection policy, follow the procedure in see Importing Networks from an RNA Detection Policy on page 940.
   - To add networks to your topology manually, follow the procedure in Manually Adding Networks to Your Custom Topology on page 940.

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**. Note that 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 942.

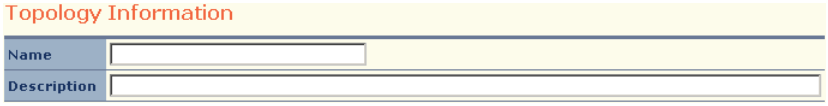
Keysight, EX-1004
Page 938

Providing Basic Topology Information

Platform: DC

Component: RNA

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 939
   - Importing Networks from an RNA Detection Policy on page 940
   - Manually Adding Networks to Your Custom Topology on page 940

Importing a Discovered Topology

Platform: DC

Component: RNA

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.



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 939
   - Importing Networks from an RNA Detection Policy on page 940
   - Manually Adding Networks to Your Custom Topology on page 940

Keysight, EX-1004
Page 939

Importing Networks from an RNA Detection Policy

Platform: DC

Component: RNA

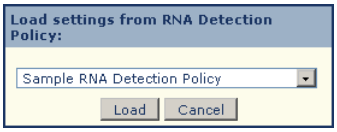
One way you can add networks to your custom topology is to import the networks that you configured the Sourcefire 3D System to monitor in an RNA detection policy. For information on RNA detection policies, see What is an RNA Detection Policy? on page 889.

To import networks from an RNA Detection 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 RNA detection policy you want to use and click **Load**.
3. The monitored networks in your RNA detection policy populate the page.

   For example, if you configured your RNA detection 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 RNA detection 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 939
   - Manually Adding Networks to Your Custom Topology on page 940

Manually Adding Networks to Your Custom Topology

Platform: DC

Component: RNA

If importing networks from your RNA detection policy and importing discovered topology is insufficient creates an inaccurate or incomplete representation of your network deployment, you can add networks to your custom topology manually.

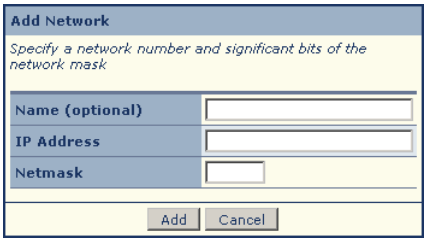
Keysight, EX-1004
Page 940

To add a network to a custom topology manually:

Access: Admin

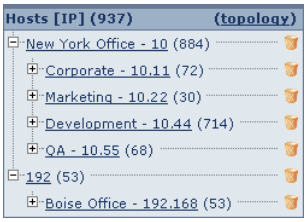
1. On the Edit Topology page, click **Add Network**.

   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.



3. For more information, see Working with the Hosts Network Map on page 926 and Working with the Network Devices Network Map on page 928.

4. 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 59.

5. Click Add.

   The network is added to your topology.

6. To add additional networks to your topology, repeat steps 1 through 5.

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.

---

7. Optionally, follow the procedures in the following sections, depending on how you want to build your custom topology:
   - Importing a Discovered Topology on page 939
   - Importing Networks from an RNA Detection Policy on page 940

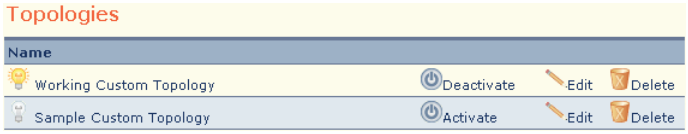
Keysight, EX-1004
Page 941

Managing Custom Topologies

Platform: DC

Component: RNA

Use the Custom Topology page to manage custom topologies. You can create, modify, and delete topologies.



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.

For more information on managing custom topologies, see:

- Activating and Deactivating Custom Topologies on page 942
- Modifying a Custom Topology on page 942
- Deleting a Custom Topology on page 943

Activating and Deactivating Custom Topologies

Platform: DC

Component: RNA

Use the following procedure to either activate or deactivate a custom topology. Note that only one custom topology can be active at any time. If you have created multiple topologies, activating one automatically deactivates the currently active topology.

To activate or deactivate a custom topology:

Access: Admin

1. Select **Policy & Response > RNA > Network Map > 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.

Modifying a Custom Topology

Platform: DC

Component: RNA

Use the following procedure to modify a custom topology.

Keysight, EX-1004
Page 942

To modify a custom topology:

Access: Admin

1. Select **Policy & Response > RNA > Network Map > Custom Topology**.

   The Custom Topology page appears.

2. Click **Edit** next to the topology.

   The Edit Topology page appears. See Creating Custom Topologies on page 937 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.

Deleting a Custom Topology

Platform: DC

Component: RNA

Use the following procedure to delete a custom topology. If you delete the active topology, your changes take effect immediately, that is, your network map no longer displays your custom topology.

To delete a custom topology:

Access: Admin

1. Select **Policy & Response > RNA > Network Map > 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.

Keysight, EX-1004
Page 943

Chapter 23
Using Host Profiles

A *host profile* provides a complete view of all the information for a single host that was collected by a 3D Sensor with RNA (including sensors that process NetFlow data) and sent to the Defense Center that manages it. 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 compliance policies and alerts based on host criticality.

Host profiles also provide you with detailed information about the services, protocols, and client applications running on a particular host, including whether they are in compliance with a compliance white list. You can view *flow data events* for services, remove services from the services list, and view service details. Flow data events are events generated at the end of a connection that contain information about the transaction. You can also view application details and flow data events for client applications and delete client applications or host protocols from the host profile.

If your Sourcefire 3D System deployment includes IPS, you can use data from the RNA network map to tailor processing of traffic by IPS to fit the type of operating system on the host and the services running on the host. For more information, see Using Adaptive Profiles on page 697.

Keysight, EX-1004
Page 944

You can see user history information for a host if you are using Real-Time User Awareness (RUA) to track user logins and at least one user has logged into the host within the last twenty-four hours.

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. After you address a vulnerability on a host, you can move that vulnerability to the Not Vulnerable list in the host profile. For example, if a Windows host has vulnerabilities that can be remediated by applying a service pack, you can then mark those vulnerabilities as Not Vulnerable. You can also apply fixes to the operating system for a host, causing all vulnerabilities addressed by the fix to be automatically marked invalid.

You can work with RNA's own vulnerability information for the hosts on your network, as well as information on vulnerabilities that were detected by third-party scanners and that you imported onto the Defense Center using the host input feature.

Optionally, you can perform an Nmap scan from the host profile, to augment the service and operating system information in your host profile. The Nmap scanner actively probes the host to obtain information about the operating system and services running on the host. The results of the scan are added to the list of operating system and service identities for the host. Note that the confidence for an operating system or service data from Nmap is always 100%.

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 host license limit
- the host resides in a network segment that is not monitored by your RNA detection policy
- you do not have RNA as part of your deployment

Note that the information displayed in a host profile can vary according to the type of host and the information available about the host. For example, if RNA detected 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 is available about an IP host.

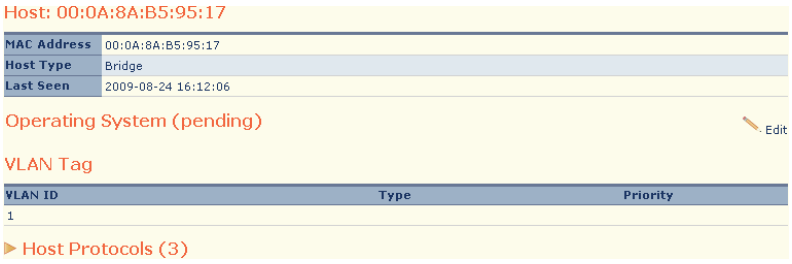
Keysight, EX-1004
Page 945

As another example, although you can configure the RNA detection policy to add hosts and services to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts and services is limited. For example, there is no operating system data available for these hosts, unless you provide it using a scanner or the host input feature. For more information, see What Information Does RNA Provide? on page 880.

The following graphic shows an example of a host profile.



The following graphic shows an example of a host profile for a MAC host.



For more information about each section of the host profile, see the following:

- Viewing Host Profiles on page 948 explains how to access a host profile.
- Working with Basic Host Information in the Host Profile on page 948 describes the information provided in the Host section of a host profile.

Keysight, EX-1004
Page 946

- Working with Operating Systems in the Host Profile on page 951 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 Services in the Host Profile on page 958 describes the information provided in the Services, Service Detail, and Service Banner sections of a host profile.
- Working with Client Applications in the Host Profile on page 969 describes the information provided in the Client Applications section of a host profile.
- Working with VLAN Tags in the Host Profile on page 972 describes the information provided in the VLAN Tag section of a host profile.
- Working with User History in the Host Profile on page 973 describes the information provided in the User History section of a host profile.
- Working with Host Attributes in the Host Profile on page 973 describes the information provided in the Attributes section of a host profile.
- Working with the Predefined Host Attributes on page 985 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 987, which provides information about creating and using user-defined host attributes.
- Working with Host Protocols in the Host Profile on page 974 describes the information provided in the Host Protocols section of a host profile.
- Working with White List Violations in the Host Profile on page 976 describes the information provided in the White List Violations section of a host profile.
- Working with Vulnerabilities in the Host Profile on page 978, which describes the information provided in the Vulnerabilities and Vulnerability Detail sections of a host profile.

Keysight, EX-1004
Page 947

Viewing Host Profiles

Platform: DC

Component: RNA

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 RNA 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: Any RNA/Admin

- 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: Any RNA/Admin

- On any network map, drill down to the IP address of the host whose profile you want to view.

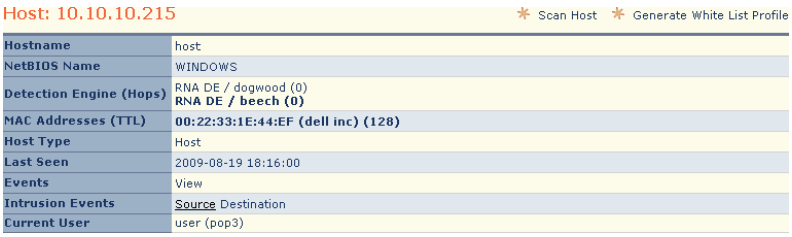
  The host profile appears. See Working with the Hosts Network Map on page 926 for an example of how to access a host profile from a network map.

Working with Basic Host Information in the Host Profile

Platform: DC

Component: RNA

Each host profile provides basic information about a detected host.



Keysight, EX-1004
Page 948

The Basic Host Profile Information table describes each field.

Basic Host Profile Information

Field	Description
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.
Detection Engine (Hops)	Either: • the reporting detection engine for the network where the host resides, as defined in the RNA detection policy, or • the detection engine that processed the NetFlow data that added the host to the network map The sensor where the detection engine resides and the number of network hops between the sensor that detected the host and the host itself follows the detection engine name, in parentheses. If multiple detection engines can see the host, the reporting detection engine is displayed in bold. If this field is blank, either: • the host was added to the network map by a detection engine that is not explicitly monitoring the network where the host resides, as defined in the RNA detection policy, or • the host was added using the host input feature and has not also been detected by RNA
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 RNA through ARP and DHCP traffic. If multiple sensors detected the host, the Defense Center displays all MAC addresses and TTL values associated with the host, regardless of which sensor 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.

Keysight, EX-1004
Page 949

Basic Host Profile Information (Continued)

Field	Description
Host Type	The type of device that RNA detected: bridge, router, NAT device, load balancer, or host. The methods RNA 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.
Last Seen	The date and time the host was last detected.
Current User	The user most recently logged into this host.

Viewing Events Associated with a Host Profile

Platform: DC

Component: RNA

You can view all the RNA network discovery or host input events associated with the host within the current time range.

To view associated RNA events:

Access: Any RNA/Admin

▶ In the **Events** field, click **View**.

The table view of RNA events appears, showing the events associated with the host.

Keysight, EX-1004
Page 950

If your Defense Center also manages 3D Sensors with IPS, to view intrusion events where the host is the source IP address:

Access: Any RNA/Admin

- Requires: DC + IPS + RNA In the **Intrusion Events** field, click **Source**.

  The first page of your preferred workflow for intrusion events appears, showing the events where the host is the source IP address. If you do not have a preferred workflow for intrusion events, you must select one.

If your Defense Center also manages 3D Sensors with IPS, to view intrusion events where the host is the destination IP address:

Access: Any RNA/Admin

- Requires: DC + IPS + RNA In the **Intrusion Events** field, click **Destination**.

  The first page of your preferred workflow for intrusion events appears, showing the events where the host is the destination IP address. If you do not have a preferred workflow for intrusion events, you must select one.

Working with Operating Systems in the Host Profile

Platform: DC

Component: RNA

RNA passively detects the identity of the operating system running on a host by analyzing the network stack in traffic generated by the host. RNA also collates operating system identity from other sources, such as the Nmap scanner or application data imported through the host input feature. RNA determines which identity to use by ranking the identity sources by priority. User input data has the highest priority, followed by application or scanner sources, then RNA.

Sometimes RNA 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, but RNA does collate information from the sources to use the most detailed definition possible.

Keysight, EX-1004
Page 951

For more information on the operating system information displayed in the host profile, see the Operating System Details table.

Operating System Details

Field Name	Description
Vendor	The operating system vendor.
Product	The operating system that RNA determines most likely to be running on the host, based on the identity data collected from all sources. If the operating system is `Pending`, RNA has not yet identified an operating system and no other identity data is available for the operating system. If the operating system is `unknown`, RNA cannot identify the OS and no other identity data is available for the operating system. If the host's operating system is not included in the list of operating systems RNA is able to detect (see Detected Operating Systems and Services on page 2019), 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 1363 • run an Nmap scan against the host, as described in Scanning a Host from the Host Profile on page 992 • 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 951

Keysight, EX-1004
Page 952

Operating System Details (Continued)

Field Name	Description
Version	The operating system version.
Source Type	One of the following values: • User: *user_name* • Application: *app_name* • Scanner: *scanner_type* (Nmap or scanner added through system policy) • RNA, RNA Port Match, or RNA Pattern Match, for services detected by 3D Sensors with RNA RNA may reconcile data from multiple sources to determine the identity of a service; see Understanding Current Identities on page 1360.
OS Confidence	One of: • the percentage of confidence that RNA has in the identity of the operating system running on the host, for hosts detected by 3D Sensors with RNA • 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner • unknown, for hosts for which RNA cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data If multiple sensors detected the host, the Defense Center displays the highest confidence reported among the sensors.

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 RNA 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 RNA 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, RNA uses both identities for vulnerabilities and impact correlation.

Although you can configure the RNA detection 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

Keysight, EX-1004
Page 953

identity. For more information, see What Information Does RNA Provide? on page 880.

You can set a custom display string for the host's operating system identity. That display string is then used in that host profile.

IMPORTANT! Note that changing the operating system information for a host can change its compliance with a compliance white list.

Viewing Operating System Identities

Platform: DC

Component: RNA

You can view the specific operating system identities discovered or added for a host. RNA 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 the Operating System Details table on page 952.

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: RNA/Admin

- 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 **Delete** 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 operating system identities added by RNA.



Keysight, EX-1004
Page 954

Editing an Operating System

Platform: DC

Component: RNA

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 RNA detects a conflicting operating system identity for the host after you edit the operating system, an operating system conflict occurs.

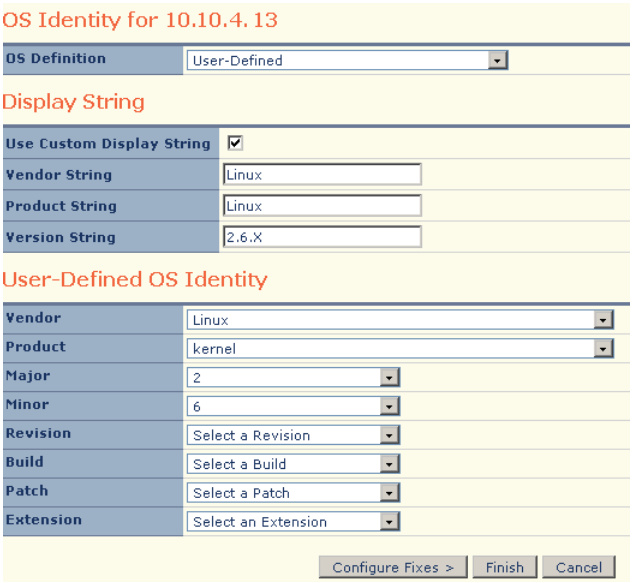
Both operating systems are then considered current until you resolve the conflict. For more information, see Resolving Operating System Identity Conflicts on page 956.

To change the operating system identity:

Access: RNA/Admin

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.



Keysight, EX-1004
Page 955

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 and skip to step 5.
   - Select a variation on the current operating system identity from the **OS Definition** drop-down list and skip to step 5.
   - Select **User-Defined** from the **OS Definition** drop-down list and 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.
5. Optionally, to configure fixes for the specified operating system identity, click **Configure Fixes**.
6. Add the patches you want to apply for that operating system to the fixes list.
7. Click **Finish** to complete the operating system identity configuration.

Resolving Operating System Identity Conflicts

Platform: DC

Component: RNA

An operating system identity conflict occurs when a new identity detected by RNA conflicts with the current identity, if that identity was provided from 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 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.

To make one of the conflicting identities current:

Access: RNA/Admin

▶ 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.

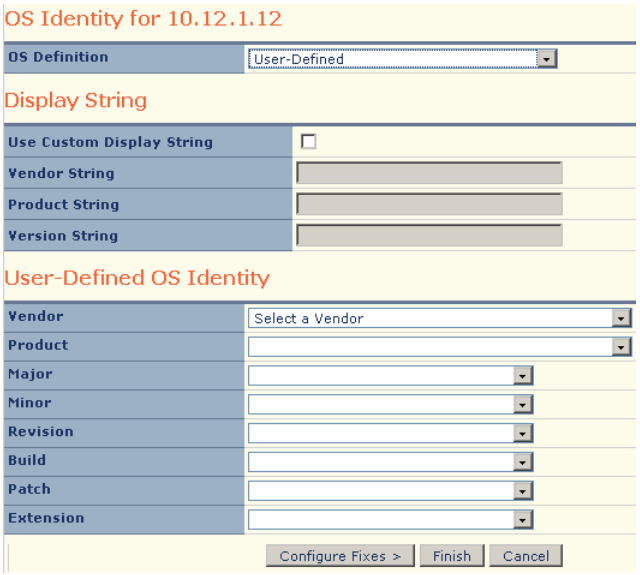
Keysight, EX-1004
Page 956

To resolve an operating system identity conflict:

Access: RNA/Admin

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 and skip to step 5.
   - Select a variation on one of the conflicting operating system identities from the **OS Definition** drop-down list and skip to step 5.
   - Select **User-Defined** from the **OS Definition** drop-down list and 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 fixes for the specified operating system identity, click **Configure Fixes**.
6. Add the patches you want to apply for that operating system to the fixes list.

Keysight, EX-1004
Page 957

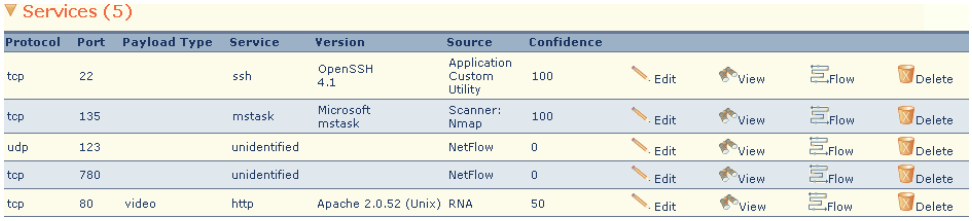
7. Click **Finish** to complete the operating system identity configuration and return to the host profile.

Working with Services in the Host Profile

Platform: DC

Component: RNA

If RNA detects services running on a host on your monitored network or if services are added through the host input feature or through a scanner or other active source, the Defense Center lists them in the Services section of the host profile.


Services (5)

Protocol	Port	Payload Type	Service	Version	Source	Confidence				
tcp	22		ssh	OpenSSH 4.1	Application Custom Utility	100	Edit	View	Flow	Delete
tcp	135		mstask	Microsoft mstask	Scanner: Nmap	100	Edit	View	Flow	Delete
udp	123		unidentified		NetFlow	0	Edit	View	Flow	Delete
tcp	780		unidentified		NetFlow	0	Edit	View	Flow	Delete
tcp	80	video	http	Apache 2.0.52 (Unix)	RNA	50	Edit	View	Flow	Delete

The Defense Center lists up to 100 services per host. Once that limit is reached, new service information from any source, whether active or passive, is discarded until you delete a service from the host or a service times out.

If you scan a host using Nmap, Nmap adds the results of previously undetected services running on open TCP ports to the Services 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 service information detected on the host by the Nmap scan. See Working with Scan Results in a Host Profile on page 992 and Setting up Nmap Scans on page 1428 for more information. In addition, note that if the host is deleted from the network map, the Nmap scan results for that service for the host are discarded.

IMPORTANT! Although you can configure the RNA detection policy to add services to the network map based on data exported by NetFlow-enabled devices, the available information about these services is limited. For more information, see What Information Does RNA Provide? on page 880.

Keysight, EX-1004
Page 958

The process for working with services in the host profile differs depending on how you accessed the profile.

- If you accessed the host profile by drilling down through the Services network map, the details for that service appears with the service name highlighted in bold. If you want to view the details for any other service on the host, click **View** next to that service name.
- If you accessed the host profile in any other way, expand the Services section and click **View** next to the service whose details you want to see.

You can also perform the following actions:

- To analyze the flow events associated with a particular service on the host, click **Flow** next to the service.

  The first page of your preferred workflow for flow events appears, showing flow events constrained by the port and protocol of the service, as well as the IP address of the host. If you do not have a preferred workflow for flow events, you must select one. For more information about flow data, see Working with Flow Data and Traffic Profiles on page 1071.
- To delete a service from the host profile, click **Delete** next to the service.

  The service is deleted from the host profile, but will appear again if RNA detects traffic from the service again. Note that deleting a service from a host can bring the host into compliance with a white list.
- To resolve a service identity conflict, click **Resolve** next to the service.

  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 service identity, click **Edit** next to the service.

  You can choose the current identity, choose a variation on that identity, or set a new user-defined identity.

The columns in the Services list are described in the following table.

Service Information

Column	Description
Protocol	The name of the protocol the service uses.
Port	The port where the service runs.

Keysight, EX-1004
Page 959

Service Information (Continued)

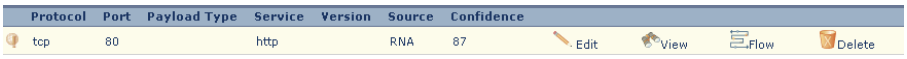
Column	Description
Payload Type	If the service is http, the general type of content (such as audio, video, or webmail) detected by RNA. For a list of detected payload types, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.
Service	One of: • the name of the service • `pending`, if RNA cannot positively or negatively identify the service for one of several reasons • `unknown`, if RNA cannot identify the service based on known service fingerprints For detailed information on how the Defense Center determines and displays service identities, see the Sourcefire 3D System Identification of Services table on page 1387.
Version	The service version as identified by RNA, Nmap or another active source, or that you specified using the host input feature. If the Sourcefire 3D System cannot identify the service version, this field is blank.

Keysight, EX-1004
Page 960

Service Information (Continued)

Column	Description
Source	One of the following values: • User: *user_name* • Application: *app_name* • Scanner: *scanner_type* • RNA, RNA Port Match, or RNA Pattern Match, for services detected by 3D Sensors with RNA • NetFlow, for services added to the network map based on NetFlow data RNA may reconcile data from multiple sources to determine the identity of a service; see Understanding Current Identities on page 1360.
Confidence	One of: • the percentage of confidence that RNA has in the identity of the service, for hosts detected by 3D Sensors with RNA • 100%, for services identified by an active source, such as the host input feature or Nmap scanner • unknown, for services for which RNA cannot determine an identity, and for services added to the network map based on NetFlow data

Note that if the host is running a service that violates a compliance white list in an activated compliance policy, the Defense Center marks the noncompliant service with the white list violation icon ( ).



See the following sections for more information:

- Service Detail on page 962
- Editing Service Identities on page 965
- Resolving Service Identity Conflicts on page 967

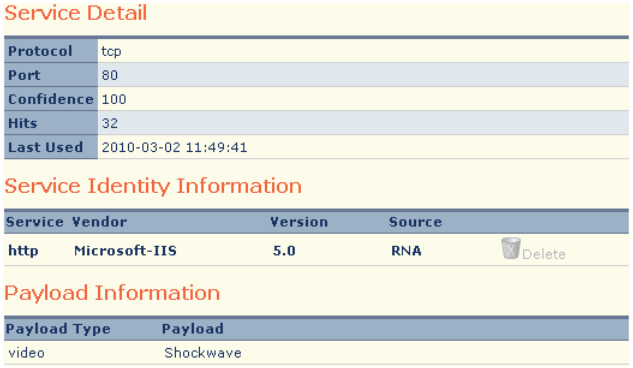
Keysight, EX-1004
Page 961

Service Detail

Platform: DC

Component: RNA

The service detail contains basic information about the detected service, as well as identity (vendor and version) and any available payload (HTTP content) information.



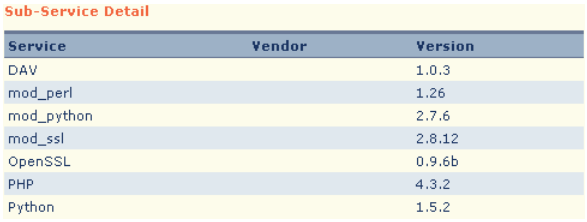
The Defense Center lists up to 16 passively detected (RNA or NetFlow) identities per service. A service can have multiple passive identities if RNA can identify a service but also detects multiple vendors or versions of that service. For example, a load balancer between your 3D Sensor and your web server farm can cause RNA to identify multiple passive identities for the HTTP service your web servers are not running the same version of the server software. Note that the Defense Center does not limit the number of service identities from active sources such as user input, scanners, or other applications.

The Defense Center displays the current identity in bold. RNA uses the current identity for a service for the multiple purposes, including to assign vulnerabilities to a host, impact assessment, when evaluating compliance rules written against host profile qualifications and compliance white lists, and so on.

TIP! For information on changing the service identity and resolving service identity conflicts from the service detail, see Editing Service Identities on page 965 and Resolving Service Identity Conflicts on page 967.

Keysight, EX-1004
Page 962

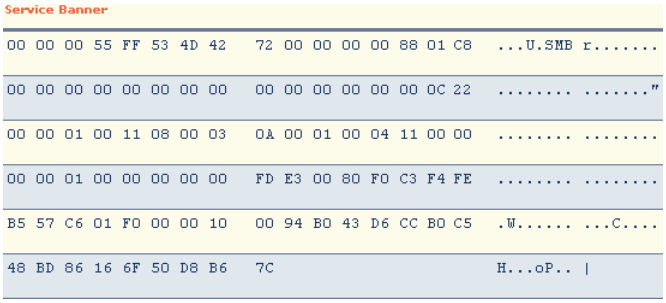
The service detail may also display sub-service information known about the selected service is also displayed. The following graphic displays sub-service information for a detected http service (Apache).



Sub-Service Detail

Service	Vendor	Version
DAV		1.0.3
mod_perl		1.26
mod_python		2.7.6
mod_ssl		2.8.12
OpenSSL		0.9.6b
PHP		4.3.2
Python		1.5.2

Finally, the service detail may display the service banner, which appears below the service details when you view a service from the host profile.



Service Banner

```
00 00 00 55 FF 53 4D 42   72 00 00 00 00 88 01 C8   ...U.SMB r.......
00 00 00 00 00 00 00 00   00 00 00 00 00 00 0C 22   ........ .......”
00 00 01 00 11 08 00 03   0A 00 01 00 04 11 00 00   ........ ........
00 00 01 00 00 00 00 00   FD E3 00 80 F0 C3 F4 FE   ........ ........
B5 57 C6 01 F0 00 00 10   00 94 B0 43 D6 CC B0 C5   .W...... ...C....
48 BD 86 16 6F 50 D8 B6   7C                       H...oP.. |
```

Service banners provide additional information about a service that may help you identify a service. RNA cannot identify or detect a misidentified service when an attacker purposely alters the service banner string. The service banner displays the first 256 bytes of the first packet detected for the service. It is collected only once, the first time the service is detected by RNA. 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 service banners, you must enable the **Capture Banners** check box in the RNA detection policy. This option is disabled by default.

The following table further explains the information provided in the service detail.

Service Detail Information

Detail	Description
Protocol	The name of the protocol the service uses.
Port	The port where the service runs.

Keysight, EX-1004
Page 963

Service Detail Information (Continued)

Detail	Description
Confidence	One of: • the percentage of confidence that RNA has in the identity of the service, for hosts detected by 3D Sensors with RNA • 100%, for services identified by an active source, such as the host input feature or Nmap scanner • `unknown`, for services for which RNA cannot determine an identity, and for services added to the network map based on NetFlow data
Hits	The number of times the service was detected by RNA or Nmap. Note that the number of hits is 0 for services imported through host input, unless RNA detects traffic for that service.
Last Used	The time and date the service was last detected. Note that the last used time for host input data reflects the initial data import time, unless RNA detects new traffic for that service. 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.
Service	The name of the service, if known.
Vendor	The service vendor. This field does not appear if the vendor is unknown.
Version	The service 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) • RNA, for services detected by 3D Sensors with RNA • NetFlow, for services added to the network map based on NetFlow data RNA may reconcile data from multiple sources to determine the identity of a service; see Understanding Current Identities on page 1360.

Keysight, EX-1004
Page 964

Service Detail Information (Continued)

Detail	Description
Payload Type	If the service is http, the general type of content (such as audio, video, or webmail) detected by RNA. For a list of detected payload types, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.
Payload	If the service is http, the specific type of content detected by RNA. For example, if the payload type is video, the payload itself could be WMV, QuickTime, MP4, and so on. For a list of detected payloads, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.

To view the service detail for a host:

Access: RNA/Admin

- Click **View** next to a service in the Services list section of a host profile.

  The Service Detail pop-up window appears.

Editing Service Identities

Platform: DC

Component: RNA

You can manually update the identity settings for a service 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 service identities.

Note that deleting an identity does not delete the service, even if that is the only identity. Deleting an identity does remove the identity from the Service Detail pop-up window and, if applicable, updates the current identity for the service in the host profile.

You cannot modify or delete service identities added by RNA.

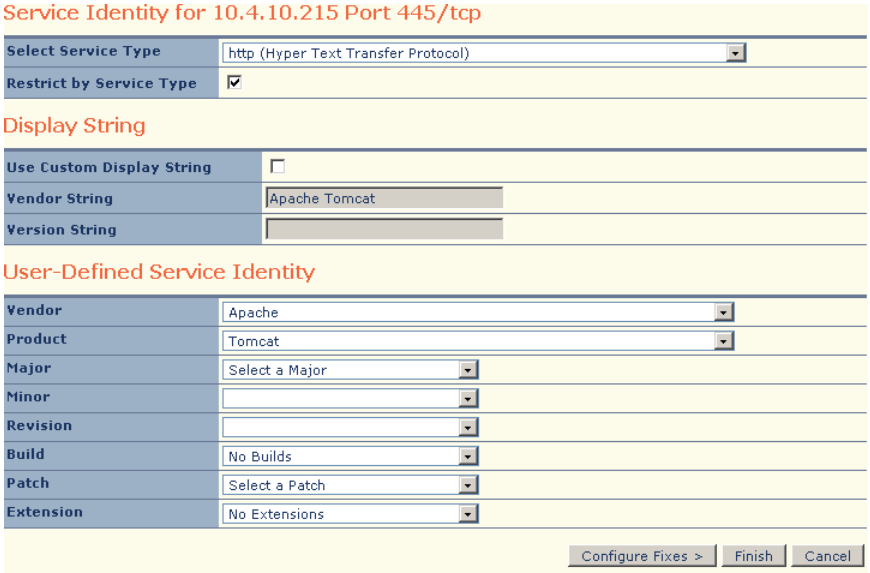
Keysight, EX-1004
Page 965

To modify the service identity:

Access: RNA/Admin

1. In the Services section in a host profile, click **View** to open the Service Detail pop-up window.
2. You have two options:
   - To delete a service identity, click **Delete** next to the service identity you want to remove.
   - To modify a service identity, click **Edit** next to the service in the Services list.

     The Service Identity pop-up window appears.



3. You have two options:
   - Select the current definition from the **Select Service Type** drop-down list.
   - Select the type of service from the **Select Service Type** drop-down list.
4. Optionally, to only list vendors and products for that service type, select the **Restrict by Service Type** check box.
5. Optionally, to customize the name and version of the service, 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 from the following lists (if applicable):
   - **Vendor**
   - **Product**
   - **Major Version**

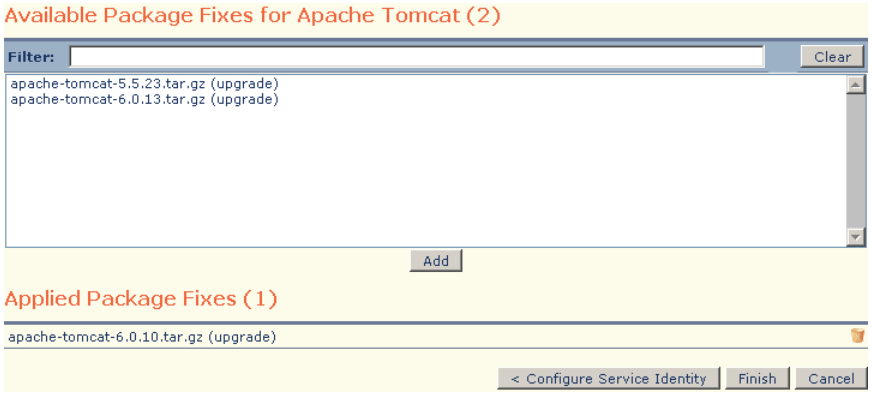
Keysight, EX-1004
Page 966

- **Minor Version**
- **Revision Version**
- **Build**
- **Patch**
- **Extension**

For example, if you want the service 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 service 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 service to the fixes list.
9. Click **Finish** to complete the service identity configuration.

Resolving Service Identity Conflicts

Platform: DC

Component: RNA

A service identity conflict occurs when an active source, such as an application or scanner, adds identity data for a service to a host, then RNA detects traffic for that port that indicates a conflicting service identity.

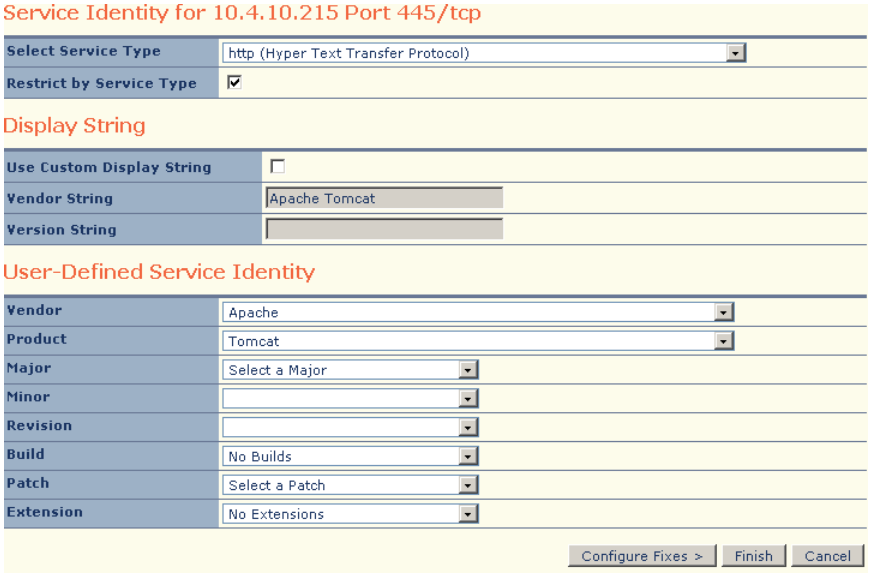
Keysight, EX-1004
Page 967

To resolve a service identity conflict:

Access: RNA/Admin

1. Click **Resolve** next to the service in the Services list.

   The Service Identity pop-up window appears.



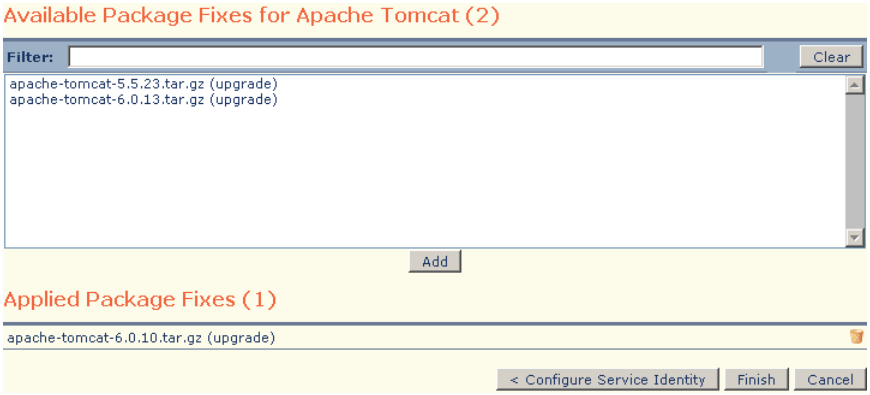
2. Select the type of service from the **Select Service Type** drop-down list.
3. Optionally, to only list vendors and products for that service type, select the **Restrict by Service Type** check box.
4. Optionally, to customize the name and version of the service, select the **Use Custom Display String** 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 from the following lists (if applicable):
   - **Vendor**
   - **Product**
   - **Major Version**
   - **Minor Version**
   - **Revision Version**
   - **Build**
   - **Patch**
   - **Extension**

   For example, if you want the service to map to Red Hat Linux 9, select **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Keysight, EX-1004
Page 968

6. If you want to indicate that fixes for the service 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 service to the fixes list.
8. Click **Finish** to complete the service identity configuration and return to the host profile.

Working with Client Applications in the Host Profile

Platform: DC

Component: RNA

You can see the client applications running on a host in the host profile. If you want to remove a client application from a host profile, you can delete that client application.

For more information on managing client applications in the host profile, see:

- Viewing Client Applications in the Host Profile on page 970
- Deleting Client Applications from the Host Profile on page 972

IMPORTANT! The host profile of a host added to the network map based on NetFlow data does not include any information on the client applications running on the host. For more information, see What Information Does RNA Provide? on page 880.

Keysight, EX-1004
Page 969

Viewing Client Applications in the Host Profile

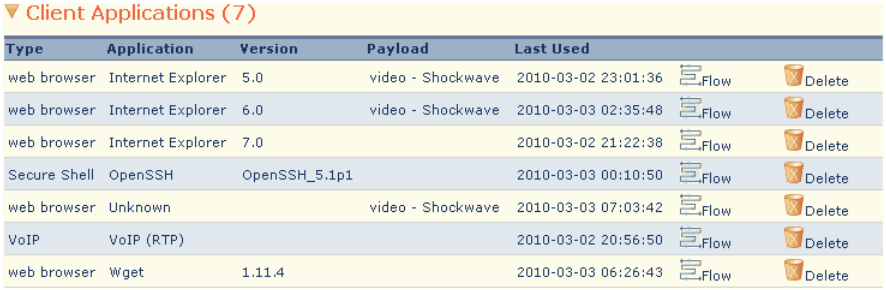
Platform: DC

Component: RNA

RNA can detect a variety of client applications running on the hosts on your network. For a list of the applications that RNA recognizes, see Detected Client Applications on page 2036.

IMPORTANT! Note that you must select the **Client Application Detection** check box in your RNA detection policy for RNA to detect client applications on the hosts in your monitored network. This option is enabled by default.

The host profile displays the product and version of the detected client applications on a host, any available payload (HTTP content) information, and the time that the application was last detected in use.



The following table describes the client application information that appears in a host profile.

Client Application Information

Column	Description
Type	Displays the category of application (HTTP browser, SMTP agent, and so on).
Application	Displays the vendor and product name of the client application.
Version	Displays the version of the client application.

Keysight, EX-1004
Page 970

Client Application Information (Continued)

Column	Description
Payload	For web browsers, the content detected by RNA in the http traffic. Payload information is comprised of a payload type, which represents the general content type (for example, audio or video), as well as a payload, which represents the specific type of content (for example, WMV or QuickTime). For a list of detected payloads and payload types, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.
Last Used	Shows the last time the host used the client application or the last time the client application was detected. If the client application was added through the host input API or command line import, the last used time indicates the time the client application data was originally imported for the host.

Note that if the host is running a client application that violates a compliance white list in an activated compliance policy, the Defense Center marks the noncompliant application with the white list violation icon ( ).



To analyze the flow events associated with a particular client application on the host, click **Flow** next to the application. The first page of your preferred workflow for flow events appears, showing flow events constrained by the type, product, and version of the application, as well as the IP address of the host. If you do not have a preferred workflow for flow events, you must select one. For more information about flow data, see Working with Flow Data and Traffic Profiles on page 1071.

Keysight, EX-1004
Page 971

Deleting Client Applications from the Host Profile

Platform: DC

Component: RNA

You can delete a client application from a host profile to remove client applications that you know are not running on the host. Note that deleting a client application from a host can bring the host into compliance with a white list.

IMPORTANT! If RNA detects the client application again, it re-adds it to the network map and the host profile.

To delete a client application from a host profile:

Access: RNA/Admin

- In the Client Applications section of the host profile, click the **Delete** icon next to the client application you want to delete.

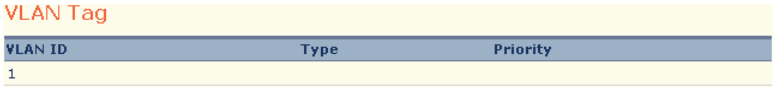
  The client application is deleted for that host.

Working with VLAN Tags in the Host Profile

Platform: DC

Component: RNA

The VLAN Tag section of the host profile appears if the host is a member of a Virtual LAN (VLAN).



Physical network equipment often uses VLANs to create logical network segments from different network blocks. RNA 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, RNA processes and the Defense Center display the innermost VLAN tag. RNA collects and the Defense Center displays VLAN tag information only for MAC addresses that it identifies through ARP and DHCP traffic.

Keysight, EX-1004
Page 972

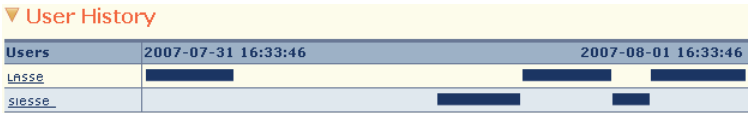
VLAN tag information can be useful, for example, if you have a VLAN composed entirely of printers and RNA detects a Microsoft Windows 2000 operating system in that VLAN. VLAN information also helps RNA generate more accurate network maps.

Working with User History in the Host Profile

Platform: DC

Component: RNA

The user history portion of the host profile provides a graphic representation of the last twenty-four hours of user activity. A typical user would log 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 the user identities is provided with bar graphs to approximate the login and logout times.



Working with Host Attributes in the Host Profile

Platform: DC

Component: RNA

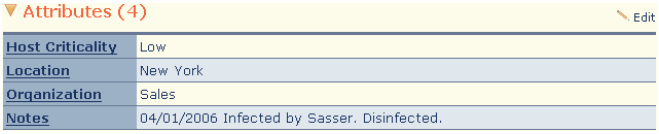
You can use *host attributes* to classify hosts in ways that are important to your network environment. Host attribute values can be positive integers, strings, or URLs. You can also create a list of string values and assign them automatically based on host IP addresses. For information about creating and managing user-defined host attributes, see Working with User-Defined Host Attributes on page 987.

The Sourcefire 3D System includes two predefined host attributes: Host Criticality and Notes. See Working with the Predefined Host Attributes on page 985 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**

Keysight, EX-1004
Page 973

manually change the value of a white list host attribute. For more information on white lists, see Using RNA as a Compliance Tool on page 1156.



Assigning Host Attribute Values

Platform: DC

Component: RNA

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: RNA/Admin

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

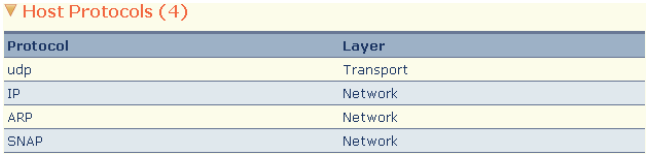
Platform: DC

Component: RNA

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.

Keysight, EX-1004
Page 974

Each host profile contains information about the protocols detected in the network traffic associated with the host.



The host profile displays protocol and network layer information as described in the Protocol Information table.

Protocol Information

Column	Description
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 compliance policy, the Defense Center marks the noncompliant protocol with the white list violation icon ( ).



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 can bring a host into compliance with a compliance white list.

IMPORTANT! If RNA 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: RNA/Admin

- 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.

Keysight, EX-1004
Page 975

Working with White List Violations in the Host Profile

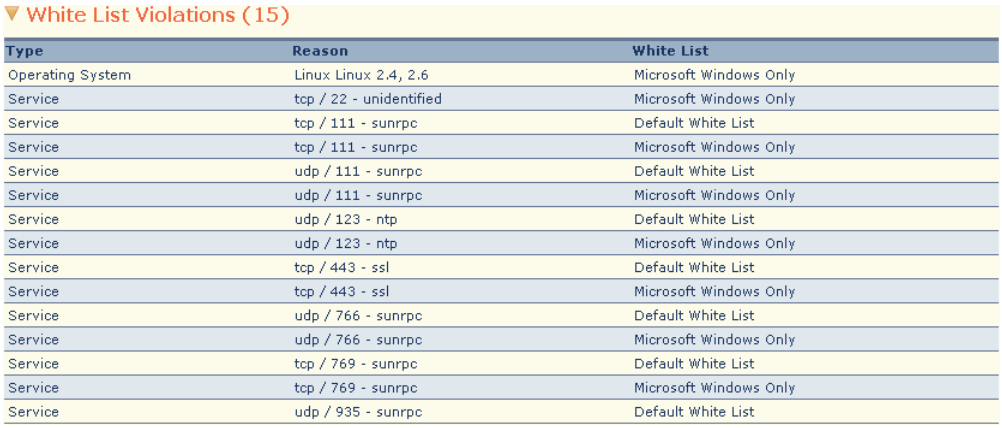
Platform: DC

Component: RNA

A *compliance white list* (or *white list*) is a set of criteria that allows you to specify the operating systems, client applications, services, and protocols that are allowed to run on a specific subnet.

If you add a white list to an active compliance policy, when RNA detects that a host is violating the white list, the Defense Center logs a white list event—which is a special kind of compliance 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 (15)

Type	Reason	White List
Operating System	Linux Linux 2.4, 2.6	Microsoft Windows Only
Service	tcp / 22 - unidentified	Microsoft Windows Only
Service	tcp / 111 - sunrpc	Default White List
Service	tcp / 111 - sunrpc	Microsoft Windows Only
Service	udp / 111 - sunrpc	Default White List
Service	udp / 111 - sunrpc	Microsoft Windows Only
Service	udp / 123 - ntp	Default White List
Service	udp / 123 - ntp	Microsoft Windows Only
Service	tcp / 443 - ssl	Default White List
Service	tcp / 443 - ssl	Microsoft Windows Only
Service	udp / 766 - sunrpc	Default White List
Service	udp / 766 - sunrpc	Microsoft Windows Only
Service	tcp / 769 - sunrpc	Default White List
Service	tcp / 769 - sunrpc	Microsoft Windows Only
Service	udp / 935 - sunrpc	Default White List

The host profile displays white list violation information as described in the White List Violation Information table.

White List Violation Information

Column	Description
Type	The type of the violation, that is, whether the violation occurred as a result of a non-compliant operating system, service, client application, or protocol.

Keysight, EX-1004
Page 976

White List Violation Information (Continued)

Column	Description
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, services, protocols, and client applications, the Defense Center marks noncompliant 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.



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

Platform: DC

Component: RNA

Shared host profiles for compliance white lists specify which operating systems, client applications, services, 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 whose IP address is known to create a shared host profile that your compliance white lists can use. However, note that you cannot create a shared host profile that is based on an individual host's host profile if RNA 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 from any event view.

   For more information, see Viewing Host Profiles on page 948.

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.

Keysight, EX-1004
Page 977

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 1189.

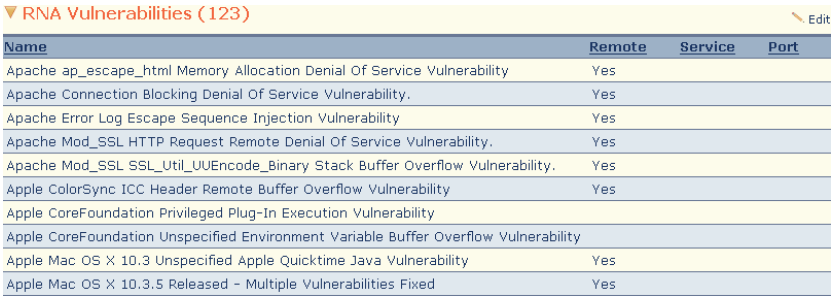
Working with Vulnerabilities in the Host Profile

Platform: DC

Component: RNA

The RNA and third-party Vulnerabilities sections of the host profile list the vulnerabilities that affect that host.

The RNA Vulnerabilities section lists vulnerabilities based on the operating system and services that RNA detected on the host.


RNA Vulnerabilities (123)

Edit

Name	Remote	Service	Port
Apache ap_escape_html Memory Allocation Denial Of Service Vulnerability	Yes		
Apache Connection Blocking Denial Of Service Vulnerability.	Yes		
Apache Error Log Escape Sequence Injection Vulnerability	Yes		
Apache Mod_SSL HTTP Request Remote Denial Of Service Vulnerability.	Yes		
Apache Mod_SSL SSL_Util_UUEncode_Binary Stack Buffer Overflow Vulnerability.	Yes		
Apple ColorSync ICC Header Remote Buffer Overflow Vulnerability	Yes		
Apple CoreFoundation Privileged Plug-In Execution Vulnerability			
Apple CoreFoundation Unspecified Environment Variable Buffer Overflow Vulnerability			
Apple Mac OS X 10.3 Unspecified Apple Quicktime Java Vulnerability	Yes		
Apple Mac OS X 10.3.5 Released - Multiple Vulnerabilities Fixed	Yes		

If there is an identity conflict for either the identity of the host's operating system or one of the services on the host, RNA lists vulnerabilities for both identities until the conflict is resolved.

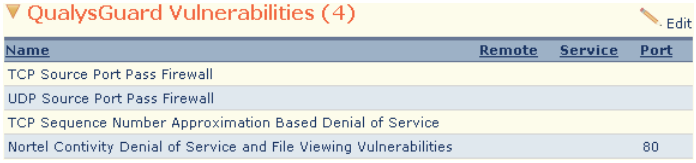
Note that because there is no operating system information available for hosts added to the network map based on NetFlow data, the Defense Center cannot determine which vulnerabilities might affect those hosts, unless you use the host input feature to manually set the hosts' operating system identity.

Also, many services do not include vendor and version information. RNA cannot automatically map vulnerabilities to a host for vendorless or versionless service traffic received or sent by the host. However, using the system policy, you can manage whether vulnerabilities are associated with services that do not have vendor or version information. For more information, see Mapping Vulnerabilities for Services on page 1773.

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

Keysight, EX-1004
Page 978

you import vulnerabilities from a QualysGuard Scanner, host profiles on your Defense Center include a QualysGuard Vulnerabilities section.



For information on importing third-party vulnerabilities, see the *Sourcefire 3D System Host Input API Guide*.

The Vulnerability Information table describes the columns in the Vulnerabilities sections of the host profile.

Vulnerability Information

Column	Description
Name	The name of the vulnerability.
Remote	Indicates whether the vulnerability can be remotely exploited.
Service	Lists the service name if the vulnerability is associated with a specific existing service.
Port	Lists the port number if the vulnerability is associated with a service 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 980 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 982 for more information.

Keysight, EX-1004
Page 979

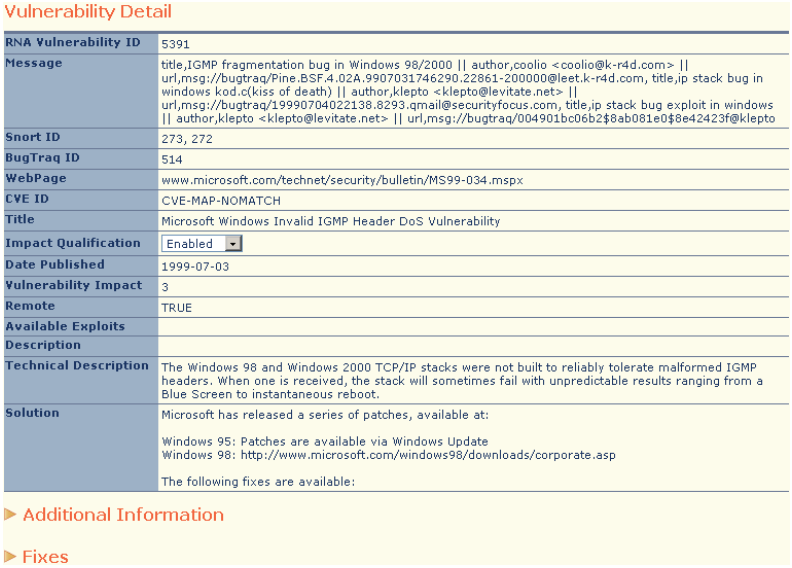
- download patches to mitigate the vulnerabilities discovered on the hosts on your network. See Downloading Patches for Vulnerabilities on page 983 for more information.
- mark hosts as not vulnerable to individual vulnerabilities if you know that the hosts have been patched. See Setting Vulnerabilities for Individual Hosts on page 984 for more information.

Viewing Vulnerability Details

Platform: DC

Component: RNA

Vulnerability details include a technical description of the vulnerability and known solutions.



Vulnerability Detail

RNA Vulnerability ID	5391
Message	title,IGMP fragmentation bug in Windows 98/2000 || author,coolio <coolio@k-r4d.com> || url,msg://bugtraq/Pine.BSF.4.02A.9907031746290.22861-200000@leet.k-r4d.com, title,ip stack bug in windows kod.c(kiss of death) || author,klepto <klepto@levitate.net> || url,msg://bugtraq/19990704022138.8293.qmail@securityfocus.com, title,ip stack bug exploit in windows || author,klepto <klepto@levitate.net> || url,msg://bugtraq/004901bc06b2$8ab081e0$8e42423f@klepto
Snort ID	273, 272
BugTraq ID	514
WebPage	www.microsoft.com/technet/security/bulletin/MS99-034.mspx
CVE ID	CVE-MAP-NOMATCH
Title	Microsoft Windows Invalid IGMP Header DoS Vulnerability
Impact Qualification	Enabled
Date Published	1999-07-03
Vulnerability Impact	3
Remote	TRUE
Available Exploits	
Description	
Technical Description	The Windows 98 and Windows 2000 TCP/IP stacks were not built to reliably tolerate malformed IGMP headers. When one is received, the stack will sometimes fail with unpredictable results ranging from a Blue Screen to instantaneous reboot.
Solution	Microsoft has released a series of patches, available at: Windows 95: Patches are available via Windows Update Windows 98: http://www.microsoft.com/windows98/downloads/corporate.asp The following fixes are available:

▶ Additional Information

▶ Fixes

Keysight, EX-1004
Page 980

The Vulnerability Detail Information table describes the fields on the Vulnerability Detail page.

Vulnerability Detail Information

Row	Description
RNA Vulnerability ID	The RNA Vulnerability identification number that RNA 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 982 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.

Keysight, EX-1004
Page 981

Vulnerability Detail Information (Continued)

Row	Description
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

Platform: DC

Component: RNA

If RNA 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 an identity conflict exists for the identity of the host's operating system or one of the services on the host, RNA lists vulnerabilities for both conflicting identities until the conflict is resolved. For more information, see Understanding Identity Conflicts on page 1362 and Resolving Operating System Identity Conflicts on page 956.

Keysight, EX-1004
Page 982

Note also that RNA 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 Managing RNA Rule State Recommendations on page 468.

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 932 and Deactivating Sourcefire Vulnerabilities on page 1061.

To change the use of a vulnerability across the system:

Access: RNA/Admin

1. Access the host profile of a host affected by the vulnerability you want to deactivate.
2. Expand the RNA 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 980.
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

Platform: DC

Component: RNA

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: RNA/Admin

1. Access the host profile of a host for which you want to download a patch.
2. Expand the RNA Vulnerabilities section.

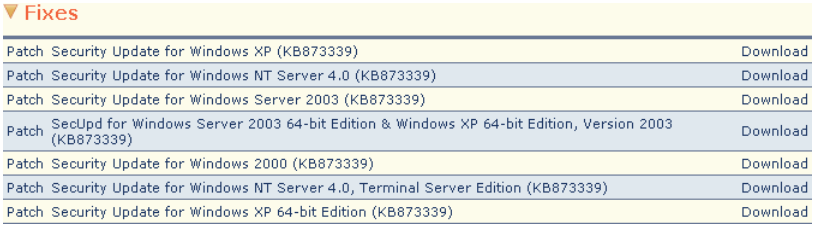
Keysight, EX-1004
Page 983

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.



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

Platform: DC

Component: RNA

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: RNA/Admin

1. Open a host profile.

Keysight, EX-1004
Page 984

2. Next to **RNA 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 980.

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**.

Working with the Predefined Host Attributes

Platform: DC

Component: RNA

There are two predefined host attributes that you can assign to each host: host criticality and host-specific notes. The following sections describe these host attributes and explain how to use them:

- Assigning a Host Criticality Value to a Host on page 986
- Adding Notes to a Host on page 986

Keysight, EX-1004
Page 985

Assigning a Host Criticality Value to a Host

Platform: DC

Component: RNA

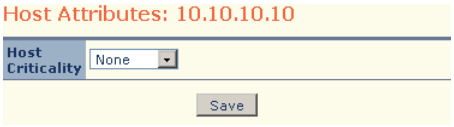
Use the host criticality attribute to designate the business criticality of a given host and to tailor compliance policies and alerts based on host criticality. For example, if you consider your organization's mail servers more critical to your business than a typical user workstation, you can assign 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 compliance policy that launches different alerts based on the criticality of an affected host.

To set host criticality in the host profile:

Access: RNA/Admin

1. Open the host profile for the host for which you want to set a business criticality.
2. Under Attributes, click **Host Criticality**.

   The Host Attributes pop-up window appears.



3. Form 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.

Adding Notes to a Host

Platform: DC

Component: RNA

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.

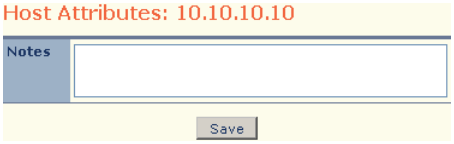
Keysight, EX-1004
Page 986

To add a note to a host profile:

Access: RNA/Admin

1. Open the host profile for the host for which you want to set a business criticality.
2. Under Attributes, click **Notes**.

   The Notes pop-up window appears.



3. Enter up to 255 alphanumeric characters, special characters, and spaces in the text box.
4. Click **Save**.

   Your notes are saved and associated with the host.

Working with User-Defined Host Attributes

Platform: DC

Component: RNA

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 compliance policies and searches. You can also view the attributes on the host attribute table view of events and generate reports based on them.

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 compliance 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.

Keysight, EX-1004
Page 987

Host attributes can be one of the following types:

Host Attribute Types

Type	Description
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 and then manually assign one of these numbers to a host.
List	Allows you to create a list of string values and then manually assign one of the values to a host. You can also automatically assign values to hosts based on the host's IP address.
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), 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 RNA as a Compliance Tool on page 1156.

See the following sections for more information:

- Restrictions on User-Defined Host Attributes on page 988
- Creating User-Defined Host Attributes on page 989
- Editing a User-Defined Host Attribute on page 991
- Deleting a User-Defined Host Attribute on page 991

Restrictions on User-Defined Host Attributes

Platform: DC

Component: RNA

Host attributes are defined globally rather than per policy. After you create a host attribute, it is available regardless of the policy applied.

Keysight, EX-1004
Page 988

Creating User-Defined Host Attributes

Platform: DC

Component: RNA

The following procedure explains how to create a user-defined host attribute.

To create a new host attribute:

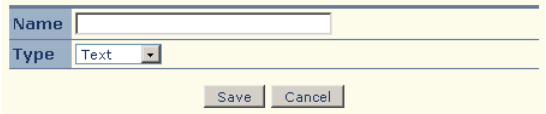
Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Host Attributes**.

   The Host Attributes page appears.

2. Click **Create Attribute**.

   The Create Attribute page appears.



3. In the **Name** field, type a name for the host attribute using alphanumeric characters and spaces.
4. Select the type of attribute that you want to create from the **Type** drop-down list, as described in the Host Attribute Types table on page 988.
   - If you are creating a **Text** or **URL** host attribute, continue with step 5.
   - If you are creating an **Integer** host attribute, see Creating Integer Host Attributes on page 989.
   - If you are creating a **List** host attribute, see Creating List Host Attributes on page 990.
5. Click **Save**.

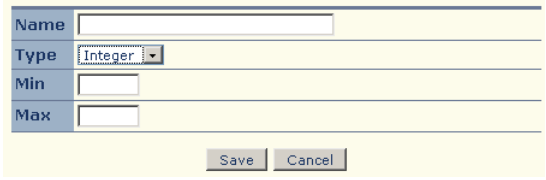
   The new user-defined host attribute is saved.

Creating Integer Host Attributes

Platform: DC

Component: RNA

When you define an integer-based host attribute, you must specify the minimum and maximum values of the range of numbers that the attribute accepts.



Keysight, EX-1004
Page 989

To create an integer-based host attribute:

Access: P&R Admin/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.

Creating List Host Attributes

Platform: DC

Component: RNA

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.

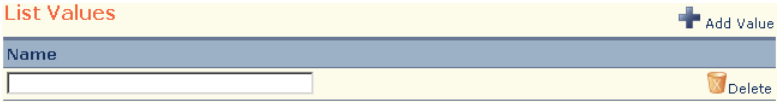
When you create a value for the host attribute, you can also auto-assign it to a range of IP addresses so that when a new host is discovered, it is automatically assigned a value for the host attribute.

To create a list-based host attribute:

Access: P&R Admin/Admin

1. To add a value to the list, click **Add Value**.

   The List Values section expands.



2. In the **Name** field, enter the first value you want to add using alphanumeric characters, symbols, and spaces.
3. Optionally, to auto-assign your hosts to the value you just added, click **Add Networks**.

   The Auto-Assign Networks section expands.



4. Select the value you added from the **Value** drop-down list.
5. 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 59.

Keysight, EX-1004
Page 990

6. 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 Assigning a Host Criticality Value to a Host on page 986.

Editing a User-Defined Host Attribute

Platform: DC

Component: RNA

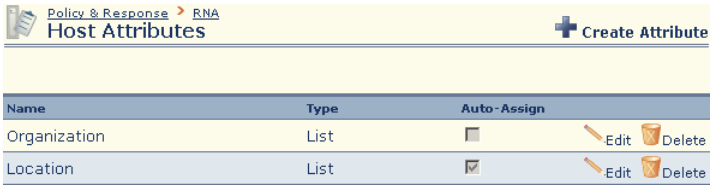
When you modify an existing user-defined host attribute, you can change the definition of a value, but you cannot change the attribute type (text, list, integer, URL). In addition, you **cannot** modify compliance white list host attributes.

To edit an existing user-defined host attribute:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Host Attributes**.

   The User-Defined Host Attribute page appears, displaying a list of the user-defined Host Attributes.



2. Click **Edit** next to the host attribute you want to edit.

   The Host Attribute page appears with the selected attribute's settings.

3. Modify any of the settings that you want and click **Save**.

   See Creating User-Defined Host Attributes on page 989 for more information about the attribute types that you can edit and the values those attributes can contain.

Deleting a User-Defined Host Attribute

Platform: DC

Component: RNA

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.

Keysight, EX-1004
Page 991

To delete a host attribute:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Host Attributes**.
   The User-Defined Host Attribute page appears.
2. Click **Delete** next to the host attribute you want to delete.
   The selected host attribute is removed from the system.

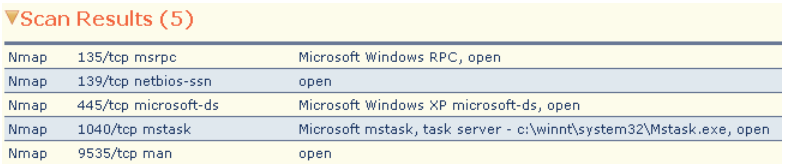
Working with Scan Results in a Host Profile

Platform: DC

Component: RNA

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.

Nmap adds information about the host operating system and any services running on open unfiltered ports directly into the Operating System and Services 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 (5)

Nmap	135/tcp msrpc	Microsoft Windows RPC, open
Nmap	139/tcp netbios-ssn	open
Nmap	445/tcp microsoft-ds	Microsoft Windows XP microsoft-ds, open
Nmap	1040/tcp mstask	Microsoft mstask, task server - c:\winnt\system32\Mstask.exe, open
Nmap	9535/tcp man	open

Each result indicates the source of the information, the number and type of the scanned port, the name of the service running on the port, and any additional information detected by Nmap, such as the state of the port or the vendor name for the service. If you scan for UDP ports, services 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

Platform: DC

Component: RNA

You can perform a Nmap scan against a host from the host profile. Once the scan completes, service and operating system information for that host are updated in

Keysight, EX-1004
Page 992

the host profile. Any additional scan results are added to the Scan Results section of the host profile.

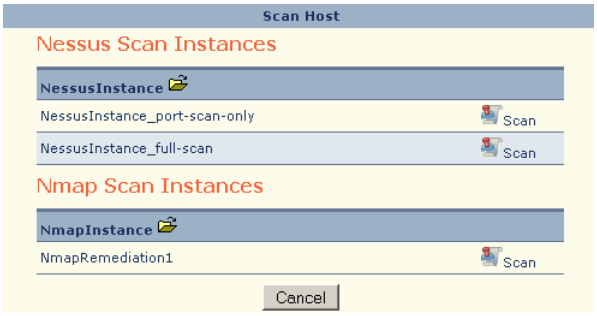
WARNING! Nmap-supplied service and operating system data remains static until you run another Nmap scan or override it with higher priority host input. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and service data up to date. For more information, see Automating Nmap Scans on page 1722.

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.

Keysight, EX-1004
Page 993

Chapter 24
Working with RNA Events

RNA events are triggered by the changes that your 3D Sensors with RNA detect in the network segments they monitor. Settings in two policies specify the kinds of information RNA collects and stores.

- The *system policy* controls the kinds and amount of RNA data stored in the database, and therefore determines the data that other parts of the Sourcefire 3D System can use.
- The *RNA detection policy* specifies the kinds of data RNA collects, as well as the monitored network segments.

These policies also control other aspects of Sourcefire 3D System behavior. For more information, see Managing System Policies on page 1732 and What is an RNA Detection Policy? on page 889.

RNA events alert you to the activity on your network and provide you with the information you need to respond appropriately. As a simple example, you may have conference rooms or a series of spare work spaces where visiting employees attach to your network. You would expect to see New Host events generated by RNA 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.

RNA 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. You can also set up RNA to detect the services, network protocols, client applications, and potential vulnerabilities running on your network hosts. RNA can provide information on vulnerabilities that were detected by third-party scanners and that you imported onto the Defense Center using the host input feature. In addition, you can track RNA-specific features such as changes to host criticality,

Keysight, EX-1004
Page 994

user-defined host attributes, and user-initiated interactions with host vulnerabilities.

RNA provides a set of predefined workflows that you can use to analyze the events that are generated for your network. You can use these predefined workflows, or you can create custom workflows that display only the information that matches your specific needs.

To collect and store RNA data for analysis, make sure that:

- you apply a host license to your Defense Center. For more information, see Understanding Licenses on page 1780.
- your RNA detection policy is configured to monitor networks whose traffic your managed 3D Sensors with RNA and NetFlow-enabled devices can see. For more information, see What is an RNA Detection Policy? on page 889.

IMPORTANT! You must use a Defense Center with an RNA host license to view RNA events. You cannot view RNA events on a Master Defense Center.

For more information, see:

- Viewing RNA Event Statistics on page 996
- Understanding RNA Event Workflows on page 1002
- Working with Network Discovery and Host Input Events on page 1004
- Working with Hosts on page 1019
- Working with Host Attributes on page 1033
- Working with Services on page 1040
- Working with Client Applications on page 1050
- Working with Sourcefire Vulnerabilities on page 1057
- Working with Third-Party Vulnerabilities on page 1064

IMPORTANT! When a connection between a monitored host and any other host is terminated, RNA generates a flow event. For more information on flow events, see Understanding Flow Data on page 1072.

Keysight, EX-1004
Page 995

Viewing RNA Event Statistics

Platform: DC

Component: RNA

The RNA Statistics page displays a summary of the hosts, events, protocols, services, and operating systems detected by RNA:

- The statistics summary provides general statistics about the total events, services, hosts, network devices, and information about your host limit usage; see Statistics Summary on page 997.
- The event breakdown provides statistics about the types of events occurring on the system; see Event Breakdown on page 998.
- The protocol breakdown provides statistics about the protocols that detected hosts are using. See Protocol Breakdown on page 1000.
- The service breakdown provides statistics about the services running on the network; see Service Breakdown on page 1000.
- 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 1001.

The page lists statistics for the last hour and the total accumulated statistics. You can select statistics for a particular detection engine, a particular sensor, or all detection engines on all sensors. You can also view events that match the entries on the page by clicking the event, service, operating system, or operating system vendor listed within the summary.

To view RNA Statistics:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Event Summary > RNA Statistics**.

   The RNA Statistics page appears.

2. From the **Select Detection Engine** list, select the detection engine whose statistics you want to view. Select **All** to view statistics for all RNA detection engines on any sensors managed by the Defense Center.

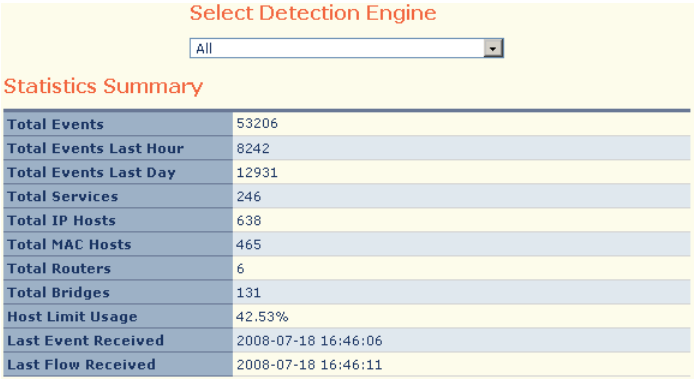
Keysight, EX-1004
Page 996

Statistics Summary

Platform: DC

Component: RNA

The statistics summary provides general statistics about the total events, services, hosts, network devices, and information about your host limit usage.



The Statistics Summary table describes the rows of the Statistics Summary section of the RNA Statistics page.

Statistics Summary

Row	Description
Total Events	Total number of RNA events stored on the Defense Center.
Total Events Last Hour	Total number of RNA events generated in the last hour.
Total Events Last Day	Total number of RNA events generated in the last day.
Total Services	Total number of services running on detected hosts.
Total IP Hosts	Total number of detected hosts RNA has identified by unique IP address.

Keysight, EX-1004
Page 997

Statistics Summary (Continued)

Row	Description
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 RNA statistics for all detection engines or for a specific detection engine. This occurs because RNA detection engines monitor hosts based on their IP addresses, and this statistic totals all hosts that cannot be identified by their IP addresses.
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 the RNA host license. Note that the host limit usage only appears if you are viewing statistics for all detection engines. **IMPORTANT!** If the host limit is reached and a host is deleted, the host will not reappear on the network map until you restart RNA.
Last Event Received	The date and time that the most recent RNA event occurred.
Last Flow Received	The date and time that the most recent RNA flow was completed.

Event Breakdown

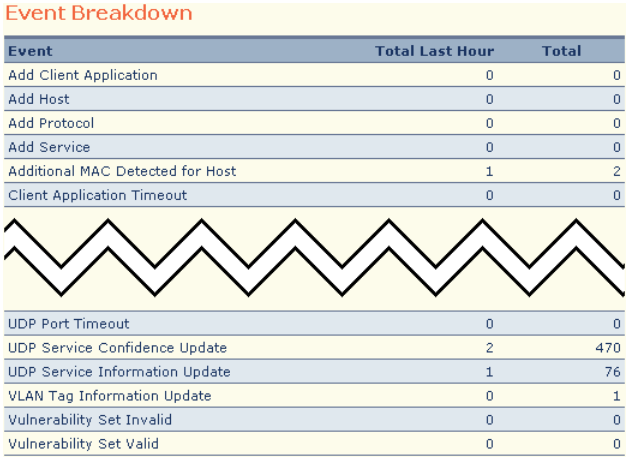
Platform: DC

Component: RNA

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

Keysight, EX-1004
Page 998

event type, see the RNA Network Discovery Event Types table on page 1006 and the RNA Host Input Event Types table on page 1011.


Event Breakdown

Event	Total Last Hour	Total
Add Client Application	0	0
Add Host	0	0
Add Protocol	0	0
Add Service	0	0
Additional MAC Detected for Host	1	2
Client Application Timeout	0	0
UDP Port Timeout	0	0
UDP Service Confidence Update	2	470
UDP Service Information Update	1	76
VLAN Tag Information Update	0	1
Vulnerability Set Invalid	0	0
Vulnerability Set Valid	0	0

You can also use the Events Breakdown section to view details on network discovery and host input events.

To view network discovery and host input events by type:

Access: Any RNA/Admin

- Click the type of event you want to view.

  The first page of the default RNA events workflow appears, constrained by the event type you picked. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

  For information on working with RNA events, see Working with Network Discovery and Host Input Events on page 1004.

Keysight, EX-1004
Page 999

Protocol Breakdown

Platform: DC

Component: RNA

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	456
ARP	Network	448
udp	Transport	374
tcp	Transport	230
icmp	Transport	115
SNAP	Network	5
Loopback	Network	2
egp	Transport	1
IPX	Network	1
NetBIOS	Network	1
Novell NetWare	Network	1
OldIPX	Network	1

Service Breakdown

Platform: DC

Component: RNA

The Service Breakdown section lists the services that are currently in use by detected hosts. It lists the service name, the total number of hosts running the

Keysight, EX-1004
Page 1000

service in the past hour, and the total number of hosts that have been detected running the service at any point.

Service Breakdown

Service	Total Last Hour	Total
unidentified	0	106
ssh	0	92
ssl	0	83
netbios-dgm	0	55
sunrpc	0	9
netbios-ssn	0	5
ntp	0	4
http	0	4
domain	0	4
unknown	0	3
netbios-ns	0	3
mysql	0	2
smtp	0	2
telnet	0	1
snmp	0	1
bootps	0	1
rsync	0	1

You can also use the Service Breakdown section to view details on the services detected by RNA.

To view service events that match a listed service name:

Access: Any RNA/Admin

- Click the name of the service you want to view.

  The first page of the default services workflow appears, constrained by the service you picked. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

  For information on working with services, see Working with Services on page 1040.

OS Breakdown

Platform: DC

Component: RNA

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 RNA's known fingerprints.

Keysight, EX-1004
Page 1001

A value of `pending` means that RNA has not yet gathered enough information to identify the operating system or its version.

OS Breakdown

OS Name	OS Vendor	Total
pending	pending	521
Linux	Linux Online, Inc.	427
Windows	Microsoft	212
Mac OS X	Apple	72
unknown	unknown	21
Linux	Gentoo	19
FreeBSD	FreeBSD Project, The	13
NetBSD	NetBSD Foundation, Inc.	11
Linux	RedHat	5
Debian Linux	Debian	4
PXE	PXE	4
Sun Solaris	Sun Microsystems, Inc.	3
SunOS	Sun Microsystems, Inc.	3
Catalyst Switch	Cisco Systems, Inc.	2
LaserJet	HP	2
OpenBSD	OpenBSD	2
Blade 100	Etherboot/Sun	1
OpenBSD OpenBSD	OpenBSD	1

You can also use the OS Breakdown section to view details on the operating systems detected by RNA.

To view hosts by operating system or vendor:

Access: Any RNA/Admin

- 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

  For information on working with hosts, see Working with Hosts on page 1019.

Understanding RNA Event Workflows

Platform: DC

Component: RNA

Defense Center provides a set of workflows that you can use to analyze the RNA events that are generated for your network. The RNA workflows are, along with the network map, a key source of information about your network assets. These workflows contain tables that are populated with data generated by RNA.

Keysight, EX-1004
Page 1002

Access RNA workflows from the **Analysis & Reporting > RNA** menu. The Defense Center provides predefined workflows for network discovery events (also called simply RNA events), as well as for detected hosts and their host attributes, services, client applications, and vulnerabilities. You can also create custom workflows. For more information on workflows, see Understanding and Using Workflows on page 1573.

TIP! Select **Analysis & Reporting > Custom Tables** to access workflows based on custom tables.

When you are using an RNA workflow, you can perform many common actions, whatever the type of event. These common functions are described in the Common RNA Event Actions table.

Common RNA 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 1495.
sort data	click the column title. Click the column title again to reverse the sort order.
constrain the columns that appear	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, use the Expand arrow ( ) to expand the search constraints, then click the column name under Disabled Columns.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1613.
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 1593.

Keysight, EX-1004
Page 1003

Common RNA Event Actions (Continued)

To...	You can...
delete items from the system, including: • network discovery and host input events from RNA events workflows • hosts and network devices from host workflows • host attributes from host attributes workflows • services from services workflows • client applications from client application workflows • third-party vulnerabilities from third-party vulnerability workflows	use one of the following methods: • 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 RNA is restarted, when they may be detected again. **TIP!** See Purging the RNA and RUA Databases on page 2014 for information on deleting all RNA events from the database. Note that you **cannot** delete Sourcefire RNA (as opposed to third-party) vulnerabilities; you can, however, mark them reviewed. For more information, see Working with Sourcefire Vulnerabilities on page 1057.
navigate to other event views to view associated events	find more information in Navigating between Workflows on page 1614.
temporarily use a different workflow	click **Workflows**. For more information, see Selecting Workflows on page 1590.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.

Working with Network Discovery and Host Input Events

Platform: DC

Component: RNA

A 3D Sensor with RNA generates network discovery events that communicate the details of changes in your monitored network segments. *New* events are

Keysight, EX-1004
Page 1004

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, RNA generates new events for each host and any TCP or UDP services discovered running on each host. Optionally, you can configure RNA to use data exported by NetFlow-enabled devices to generate these new host and service events.

In addition, RNA generates new events for each network and transport protocol running on each discovered host. Optionally, you can configure RNA to generate new events when it detects specific applications running on a host (by default, this option is enabled in RNA detection policies).

After the initial network mapping is complete, RNA continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered host, or service changes.

When an RNA event is generated, it is logged to the database. You can use the Defense Center web interface to view, search, and delete RNA events.You can also use RNA events in compliance rules. Based on the type of RNA event generated as well as other criteria that you specify, you can build compliance rules that, when used in a compliance policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria.

You can add data to the RNA network map using the host input feature. You can add, modify, or delete operating system information, which causes RNA to stop updating that information for that host. You can also manually add, modify, or delete services, client applications, and host attributes or modify vulnerability information. When you do this, RNA generates host input events.

See the following sections for more information:

- Understanding RNA Network Discovery Event Types on page 1005
- Understanding RNA Host Input Event Types on page 1010
- Viewing RNA Network Discovery and Host Input Events on page 1012
- Understanding the RNA Events Table on page 1014
- Searching for RNA Network Discovery Events on page 1016

Understanding RNA Network Discovery Event Types

Platform: DC

Component: RNA

There are many types of RNA network discovery events. For example, RNA generates and logs a New Host event when it detects a new host on your monitored network segment. When you view a table of RNA events, the event type is listed in the **Event** column. For more information, see Viewing RNA Network Discovery and Host Input Events on page 1012.

Contrast RNA network discovery events, which are generated when RNA detects a change in your monitored network (such as detecting traffic from a previously

Keysight, EX-1004
Page 1005

undetected host), with RNA 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 RNA Host Input Event Types on page 1010.

You can configure the types of network discovery events that RNA logs by modifying your system policy. By default, RNA logs all types of network discovery events. For more information, see Configuring Database Event Limits on page 1747.

If you understand the information the different types of network discovery events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in compliance policies. In addition, knowing the names of the event types can help you craft more effective event searches. The following table describes all the different types of network discovery events.

RNA Network Discovery Event Types

Event Name	Description
Additional MAC Detected for Host	This event is generated when RNA detects a new MAC address for a previously discovered host. This event is often generated when RNA 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 RNA 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 Application Timeout	This event is generated when RNA drops a client application from the database due to inactivity.
Client Application Update	This event is generated when RNA 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 RNA 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 RNA detects a change in the number of network hops between a host and the sensor running the detection engine that detects the host. This may happen if the detection engine 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 detection engine detects an ARP transmission from the host, indicating that the host is on a local segment.

Keysight, EX-1004
Page 1006

RNA Network Discovery Event Types (Continued)

Event Name	Description
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. **IMPORTANT!** If the host limit is reached and a host is deleted, the host does not re-appear on the network map until you restart RNA on all registered sensors with RNA. For more information on restarting RNA, see Purging the RNA and RUA Databases on page 2014.
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. New hosts are dropped if you enable the **Drop New Hosts When Host Limit Reached** option in the RNA Settings section of the system policy. See Configuring RNA Settings on page 1756 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 RNA Settings section of the system policy. See Configuring RNA Settings on page 1756 for information about configuring the host timeout value. If you change the networks you want to monitor in your RNA detection policy, you may want to manually delete any old hosts from the network map so that they do not count against your host license. For more information, see Working with the Hosts Network Map on page 926.
Host Type Changed to Network Device	This event is generated when RNA detects that a detected host is actually a network device.
Identity Conflict	This event is generated when RNA detects a new service or operating system identity that conflicts with a current active identity for that service 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 1333. For more information, see Understanding Identity Conflicts on page 1362 and Configuring RNA Settings on page 1756. For information on manually resolving conflicts, see Resolving Operating System Identity Conflicts on page 956 and Resolving Service Identity Conflicts on page 967.

Keysight, EX-1004
Page 1007

RNA Network Discovery Event Types (Continued)

Event Name	Description
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 1333. For more information, see Configuring RNA Settings on page 1756.
MAC Information Change	This event is generated when RNA detects a change in the information associated with a specific MAC address or TTL value. This event often occurs when RNA 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 RNA 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 RNA detects the actual MAC address of the host.
NETBIOS Name Change	This event is generated when RNA detects a change to a host's NetBIOS name. This event will only be generated for hosts using the NetBIOS protocol.
New Client Application	This event is generated when RNA detects a new client application. **IMPORTANT!** To collect and store client application data for analysis, make sure that you enable client application detection in your RNA detection policy. For more information, see What is an RNA Detection Policy? on page 889.
New Host	This event is generated when RNA detects a new host running on the network. If you enabled the **Generate Hosts from NetFlow Data** option in your RNA detection policy, this event is also generated when an RNA detection engine processes NetFlow data that involves a new host.
New Network Protocol	This event is generated when RNA detects that a host is communicating with a new network protocol (IP, ARP, and so on).
New OS	This event is generated when RNA either detects a new operating system for a host, or a change in a host's operating system.

Keysight, EX-1004
Page 1008

RNA Network Discovery Event Types (Continued)

Event Name	Description
New TCP Service	This event is generated when RNA detects a new TCP service 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 service; that information is transmitted in the TCP Service Information Update event. If you enabled the **Generate Services from NetFlow Data** option in your RNA detection policy, this event is also generated when an RNA detection engine processes NetFlow data involving a service on your monitored networks that does not already exist in the network map.
New Transport Protocol	This event is generated when RNA detects that a host is communicating with a new transport protocol, such as TCP or UDP.
New UDP Service	This event is generated when RNA detects a new UDP service running on a host. If you enabled the **Generate Services from NetFlow Data** option in your RNA detection policy, this event is also generated when an RNA detection engine processes NetFlow data involving a service on your monitored networks that does not already exist in the network map.
TCP Port Closed	This event is generated when RNA detects that a TCP port has closed on a host.
TCP Port Timeout	This event is generated when RNA has not detected activity from a TCP port within the interval defined in RNA Settings section of the system policy. See Configuring RNA Settings on page 1756 for information about configuring the service timeout value.
TCP Service Confidence Update	This event is generated when RNA updates its confidence level in the correct identification of a TCP service it detected running on a host.
TCP Service Information Update	This event is generated when RNA detects a change in a discovered TCP service running on a host. This event may be generated if a TCP service is upgraded.
UDP Port Closed	This event is generated when RNA detects that a UDP port has closed on a host.
UDP Port Timeout	This event is generated when RNA has not detected activity from a UDP port within the interval defined in RNA Settings section of the system policy. See Configuring RNA Settings on page 1756 for information about configuring the service timeout value.

Keysight, EX-1004
Page 1009

RNA Network Discovery Event Types (Continued)

Event Name	Description
UDP Service Confidence Update	This event is generated when RNA updates its confidence level in the correct identification of a UDP service it detected running on a host.
UDP Service Information Update	This event is generated when RNA detects a change in a discovered UDP service running on a host. This event may be generated if a UDP service is upgraded.
VLAN Tag Information Update	This event is generated when RNA 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 972.

Understanding RNA Host Input Event Types

Platform: DC

Component: RNA

There are many types of RNA host input events. For example, RNA generates and logs an Add Host event when a user adds a host using the host import feature. When you view a table of RNA events, the event type is listed in the **Event** column. For more information, see Viewing RNA Network Discovery and Host Input Events on page 1012.

Contrast RNA host input events, which are generated when a user takes a specific action (such as manually adding a host), with RNA network discovery events, which are generated when RNA 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 RNA Network Discovery Event Types on page 1005.

You can configure the types of host input events that RNA logs by modifying your system policy. By default, RNA logs all types of host input events. For more information, see Configuring Database Event Limits on page 1747.

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 compliance policies. In addition, knowing the names of the event types can help you craft more effective event searches. The

Keysight, EX-1004
Page 1010

RNA Host Input Event Types table describes all the different types of host input events.

RNA Host Input Event Types

Event Name	Description
Add Client Application	This event is generated when a user adds a client application.
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 RNA adds the results of an Nmap scan to a host.
Add Service	This event is generated when a user adds a service.
Delete Client Application	This event is generated when a user deletes a client application 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 Service	This event is generated when a user deletes a service or group of services 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.

Keysight, EX-1004
Page 1011

RNA Host Input Event Types (Continued)

Event Name	Description
Set Service Definition	This event is generated when a user sets the vendor and version definitions for a service.
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 RNA Network Discovery and Host Input Events

Platform: DC

Component: RNA

You can use the Defense Center to view a table of RNA events. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access RNA events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of RNA 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 1618.

Keysight, EX-1004
Page 1012

The RNA Event Actions table below describes some of the specific actions you can perform on an RNA events workflow page. You can also perform the tasks described in the Common RNA Event Actions table on page 1003.

RNA Event Actions

To...	You can...
modify the time and date range for displayed events	find more information in Setting Event Time Constraints on page 1597. 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.
learn more about the contents of the columns in the table	find more information in Understanding the RNA Events Table on page 1014.
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
search for RNA events	click **Search**. For more information, see Searching for RNA Network Discovery Events on page 1016.

Keysight, EX-1004
Page 1013

To view RNA events:

Access: Any RNA/Admin

- Select **Analysis & Reporting > RNA > Events**.

  The first page of the default RNA events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

TIP! If you are using a custom workflow that does not include the table view of RNA events, from the **Workflows** menu on the toolbar, select **RNA Events**.

Understanding the RNA Events Table

Platform: DC

Component: RNA

RNA generates network discovery events (also called RNA 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, RNA generates new events for each host and any TCP or UDP services it discovers on each host. In addition, RNA generates new events for each network and transport protocol running on each discovered host. Optionally, you can configure RNA to generate new events when it detects specific applications running on a host. After the initial network mapping is complete, RNA continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered host, service, or client application changes.

The fields in the RNA events table are described in the RNA Event Fields table.

RNA Event Fields

Field	Description
Time	The time that RNA generated the event.
Event	The event type. See Understanding RNA Network Discovery Event Types on page 1005 for a description of each available event.
IP Address	The IP address of the host involved in the event

Keysight, EX-1004
Page 1014

RNA Event Fields (Continued)

Field	Description
User	The last user to log into the host involved in the event before the event was generated.
MAC Address TTL	The MAC address of the NIC used by the network traffic that triggered the network 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 network discovery event.
Port	The port used by the traffic that triggered the event, if applicable.
Description	The text description of the event.
Confidence	One of: • the percentage of confidence that RNA has in the identity of a service, for services detected by 3D Sensors with RNA • 100%, for services identified by an active source, such as the host input feature or Nmap scanner • unknown, for services for which RNA cannot determine an identity, and for services running on hosts added to the network map based on NetFlow data • blank, for non-service related RNA events
Detection Engine	The name of the detection engine that generated the event. For new host and new service events based on NetFlow data, this is the detection engine that processed the NetFlow data.
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.

Keysight, EX-1004
Page 1015

Searching for RNA Network Discovery Events

Platform: DC

Component: RNA

You can search for specific RNA events 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. The default searches are:

- NetSky.S Worm Search, which searches for TCP activity on port 6789, a port commonly used by the NetSky.S worm.
- New Events, which searches for any kind of new event.
- Sasser Worm Search, which searches for TCP activity on ports 9996 and 5554, ports commonly used by the Sasser worm.
- Subseven Trojan Search, which searches for TCP activity on a number of ports used by the Subseven Trojan and a Subseven description.
- Timeout Events, which searches for host timeout events.
- Update Events, which searches for any kind of update event.

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 RNA Event Search Criteria table.

RNA Event Search Criteria

Field	Search Criteria Rules
Event	Enter a complete or partial event name. For example, you could enter `New` to search for all new events, `Change` to search for all change events, or `Update` to search for update events. You could also enter `Host` to view host events, or `Host IP address changed` to view all IP address change events. The event names you can use are listed in Understanding RNA Network Discovery Event Types on page 1005.
MAC Address	Specify the MAC address of a host whose events you want to view. Enter MAC addresses in the standard format, for example, `0A:BB:CD:AF:5F:07`. You can use an asterisk (*) as a wildcard character.

Keysight, EX-1004
Page 1016

RNA Event Search Criteria (Continued)

Field	Search Criteria Rules
MAC Vendor	Specify all or part of the NIC hardware vendor associated with the MAC address of the hosts whose events you want to view. For example, entering `Apple` returns all events where the NIC hardware vendor contains the word Apple, including Apple, Inc. **TIP!** To search for virtual MAC vendors, that is, for events that involve virtual machines, type `virtual_mac_vendor`. **TIP!** 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. For example, searching for `VMWare, Inc.` without quotes returns all events with vendor names that include `VMWare` as well as all events with vendor names that include `Inc.` such as Apple, Inc. and Dell, Inc.
IP Address	Type a specific IP address or use CIDR notation to specify a range of IP addresses for the hosts whose events you want to view. See Specifying IP Addresses in Searches on page 1557 for a full description of the syntax allowed for IP addresses.
User	Specify all or part of the user identity (user name) for the user involved in the events you want to view. Specify multiple user identities in a comma-separated list. This field is case insensitive. **TIP!** Use `blank` or `unknown` to search for events that do not have associated user identities; use `!blank` or `!unknown` to search for only the events that do have associated user identities.
Port	Specify the port used by the traffic that triggered the event. Note that you cannot: • enter a port/protocol combination as you can when searching for other kinds of events • use spaces when specifying port numbers or ranges. For more information, see Specifying Ports in Searches on page 1558.
Description	Specify the description of the events you want to view. For example, enter `HTTP` to view any web server-related events, `IIS` to view IIS-related events, or `PHP` to view PHP-related events. This search field is case insensitive. You can use an asterisk (*) as a wildcard character in this field.

Keysight, EX-1004
Page 1017

RNA Event Search Criteria (Continued)

Field	Search Criteria Rules
Confidence	Specify the level of confidence that RNA has in its assessment of a service for the service-related events you want to view. You can precede the confidence with greater than (>), greater than or equal to (>=), less than (<), less than or equal to (<=), or equal to (=) operators. For example, use `<=50` to view service-related events with a confidence of 50% or less. **TIP!** Type `unknown` to search for service-related events where RNA cannot determine a service identity, which includes services added to the network map based on NetFlow data.
Detection Engine	Specify the name of the detection engine that generated the events you want to view. For more information, see Specifying Detection Engine/Appliance Names on page 1559.

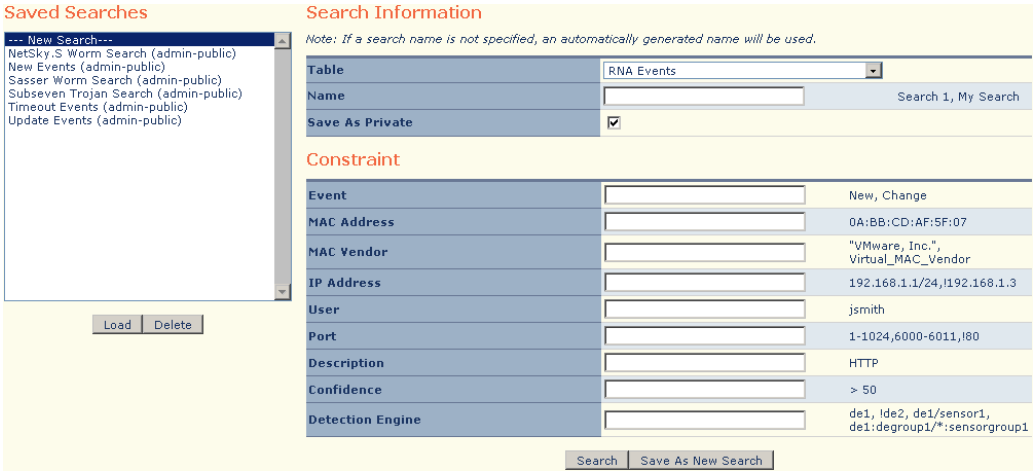
For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

To search for RNA events:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > RNA Events**.

   The RNA Events search page appears.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

Keysight, EX-1004
Page 1018

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 RNA Event 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 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 RNA events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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

Platform: DC

Component: RNA

RNA generates an event when it detects a host. RNA then collects information about the detected hosts and uses that information to build host profiles. You can use the Defense Center web interface to view, search, and delete RNA 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

Keysight, EX-1004
Page 1019

then use these criticality values, white lists, and traffic profiles within compliance rules and policies.

IMPORTANT! Although you can configure the RNA detection 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 What Information Does RNA Provide? on page 880.

For more information, see the following sections.

- Viewing Hosts on page 1020
- Understanding the Hosts Table on page 1022
- Creating a Traffic Profile for Selected Hosts on page 1025
- Creating a Compliance White List Based on Selected Hosts on page 1026
- Locating Whois Information for a Host on page 1027
- Searching for Hosts on page 1027
- Setting Host Attributes for Specific Hosts on page 1036

Viewing Hosts

Platform: DC

Component: RNA

You can use the Defense Center to view a table of hosts that RNA has detected. Then, you can manipulate the view depending on the information you are looking for.

The page you see when you access RNA hosts differs depending on the workflow you use. There are two predefined workflows:

- The Operating System Summary workflow provides a series of pages that constrains the detected hosts based on its operating system name, version, and host criticality. It also contains a table view of hosts that lists all detected hosts with the most recently active at the top of the list. Each row in the table contains a single detected host.
- The RNA Hosts workflow includes table views of hosts (with and without MAC addresses) that list all detected hosts with the most recently active at the top of the list. Each row in the table contains a single detected host.

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 1618.

Keysight, EX-1004
Page 1020

The RNA Host Actions table below describes some of the specific actions you can perform on an RNA hosts workflow page. You can also perform the tasks described in the Common RNA Event Actions table on page 1003.

RNA 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 1022.
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
assign a host attribute to selected hosts	find more information in Setting Host Attributes for Specific Hosts on page 1036.
create traffic profiles for selected hosts	find more information in Creating a Traffic Profile for Selected Hosts on page 1025.
create a compliance white list based on selected hosts	find more information in Creating a Compliance White List Based on Selected Hosts on page 1026.
locate whois information for an external host	find more information in Locating Whois Information for a Host on page 1027.
search for hosts	click **Search**. For more information, see Searching for Hosts on page 1027.

Keysight, EX-1004
Page 1021

To view hosts:

Access: Any RNA/Admin

- Select **Analysis & Reporting > RNA > Hosts**.

  The first page of the default hosts workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of hosts, from the **Workflows** menu on the toolbar, select **RNA Hosts**.

Understanding the Hosts Table

Platform: DC

Component: RNA

When RNA discovers a host, it collects data about that host. That data can include the host's IP address, 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 RNA collects about detected hosts, see Using Host Profiles on page 944.

IMPORTANT! Although you can configure the RNA detection 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 What Information Does RNA Provide? on page 880.

Keysight, EX-1004
Page 1022

The fields in the hosts table are described in the Host Fields table.

Host Fields

Field	Description
Last Seen	The date and time the host was last detected by RNA. The last seen value is updated at least as often as the update interval you configured in the RNA detection policy, as well as when RNA generates a new host event for the host's IP address. 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 address of the host.
Current User	The user identity (username) of the currently logged in user on the host.
Host Criticality	The user-specified criticality value assigned to the host. See the description of the Host Criticality column in the Host Attribute Search Criteria table on page 1038 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 972.
Hops	The number of network hops from the sensor running the detection engine that detected the host to the host.
Host Type	The type of host (host, bridge, router, NAT device, or load balancer). The methods RNA 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.

Keysight, EX-1004
Page 1023

Host Fields (Continued)

Field	Description
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. In this field, a value of `unknown` means that the operating system does not match any of RNA's known fingerprints. A value of `pending` means that RNA 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. In this field, a value of `unknown` means that the operating system does not match any of RNA's known fingerprints. A value of `pending` means that RNA 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. In this field, a value of `unknown` means that the operating system does not match any of RNA's known fingerprints. A value of `pending` means that RNA 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. In this field, a value of `unknown` means that the operating system does not match any of RNA's known fingerprints. A value of `pending` means that RNA 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 system policy) • RNA, for services detected by 3D Sensors with RNA RNA may reconcile data from multiple sources to determine the identity of an operating system; see Understanding Current Identities on page 1360.
Confidence	One of: • the percentage of confidence that RNA has in the identity of the operating system running on the host, for hosts detected by 3D Sensors with RNA • 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner • `unknown`, for hosts for which RNA cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data

Keysight, EX-1004
Page 1024

Host Fields (Continued)

Field	Description
Notes	The user-defined content of the Notes host attribute.
Detection Engine	Either: • the reporting detection engine for the network where the host resides, as defined in the RNA detection policy • the detection engine that processed the NetFlow 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 detection engine that is not explicitly monitoring the network where the host resides, as defined in the RNA detection policy, or • the host was added using the host input feature and has not also been detected by RNA
MAC Address	The host's detected MAC address of the NIC. The MAC Address field appears in the Table View of Hosts with MAC, which you can find in the RNA Hosts workflow. You can also add the MAC Address field to: • custom tables that include fields from the RNA Hosts table • drill-down pages in custom workflows based on the RNA 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 with MAC, which you can find in the RNA Hosts workflow. You can also add the MAC Vendor field to: • custom tables that include fields from the RNA Hosts table • drill-down pages in custom workflows based on the RNA Hosts table
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

Platform: DC

Component: RNA

A traffic profile is a profile of the traffic on your network, based on flow data collected by RNA over a timespan that you specify. Once you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

Keysight, EX-1004
Page 1025

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 flows 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. 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 1118.

Creating a Compliance White List Based on Selected Hosts

Platform: DC

Component: RNA

Compliance white lists allow you to specify which operating systems, services, client applications, and 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. 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 1167.

Keysight, EX-1004
Page 1026

Locating Whois Information for a Host

Platform: DC/MDC/Standalone Sensor

Component: None

If your Defense Center is connected to the Internet, you can use the Whois feature to look up ARIN information about an external host.

To use the Whois feature:

Access: Any

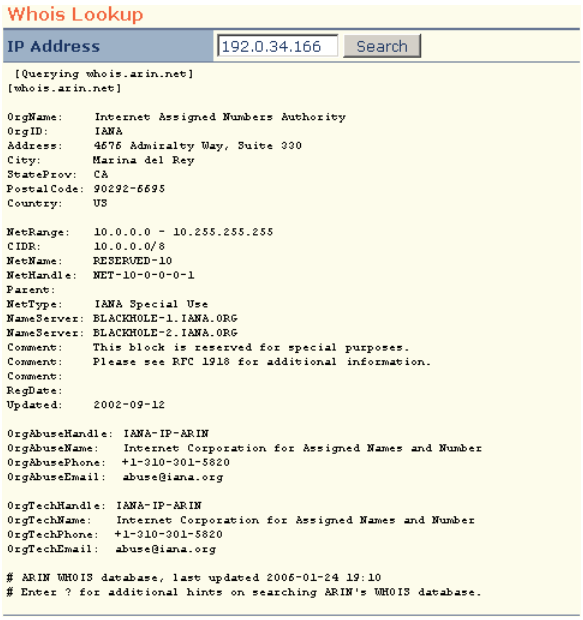
1. Select **Operations > Tools > Whois**.

   The Whois page appears.



2. Type the IP address you want to look up and click **Search**.

   The results of the lookup appear.



Searching for Hosts

Platform: DC

Component: RNA

You can search for specific hosts by using one of the predefined searches or by using your own search criteria. The predefined searches serve as examples and

Keysight, EX-1004
Page 1027

can provide quick access to important information about your network. The default searches are:

- Local Systems search, which searches for local systems less than two network hops away from the sensor running the detection engine that detected the host.
- Remote Systems search, which searches for remote systems farther than one network hop away from the sensor running the detection engine that detected the host.
- Unidentified Systems search, which searches for systems whose operating system has not yet been identified by RNA.

  IMPORTANT! Unidentified hosts are not the same as unknown hosts. *Unidentified* hosts are hosts about which RNA has not yet gathered enough information to identify their operating systems. *Unknown* hosts are hosts whose traffic has been analyzed by RNA, but whose operating systems do not match any of RNA's known fingerprints.

- Unknown Systems search, which searches for hosts whose operating systems do not match any operating system fingerprints in the RNA database.

IMPORTANT! When searching for hosts, you should keep in mind that although you can configure the RNA detection 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 What Information Does RNA Provide? on page 880.

Keysight, EX-1004
Page 1028

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 Host Search Criteria table.

Host Search Criteria

Field	Search Criteria Rules
Host Type	Type the kind of device you want to search for: `bridge`, `router`, `NAT device`, `load balancer`, or `host`. **TIP!** To search for all network devices, type `!host`.
IP Address	Type a specific IP address or use CIDR notation to specify a range of IP addresses for the hosts you want to view. See Specifying IP Addresses in Searches on page 1557 for a full description of the syntax allowed for IP addresses.
Current User	Specify all or part of the user identity (user name) for the user logged in to the hosts you want to view. Specify multiple user identities in a comma-separated list. This field is case insensitive. **TIP!** Use `blank` or `unknown` to search for hosts that do not have associated user identities; use `!blank` or `!unknown` to search for only the hosts that do have associated user identities.
MAC Address	Specify one of: • the MAC address of the host you want to view. Enter MAC addresses in the standard format, for example, `0A:BB:CD:AF:5F:07`. You can use an asterisk (*) as a wildcard character. • `blank` or `none` to search for hosts added by the host import feature that do not have an associated MAC address
MAC Vendor	Specify all or part of the NIC hardware vendor associated with the MAC address of the hosts whose events you want to view. For example, entering `Apple` returns all events where the NIC hardware vendor contains the word Apple, including Apple, Inc. **TIP!** To search for virtual MAC vendors, that is, for events that involve virtual machines, type `virtual_mac_vendor`. **TIP!** 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. For example, searching for `VMWare, Inc.` without quotes returns all events with vendor names that include `VMWare` as well as all events with vendor names that include `Inc.` such as Apple, Inc. and Dell, Inc.
NetBIOS Name	Specify the NetBIOS name of the host you want to view. You can use an asterisk (*) as a wildcard character in this field.

Keysight, EX-1004
Page 1029

Host Search Criteria (Continued)

Field	Search Criteria Rules
Host Criticality	Specify the host criticality of the hosts you want to view by typing `None`, `Low`, `Medium`, or `High`. For more information on host criticality, see Setting Host Attributes for Specific Hosts on page 1036.
VLAN ID	Specify the VLAN ID of the host or hosts you want to view. For example, if you have a VLAN with an ID of 5 and want to view all hosts in that VLAN, type `5` in the VLAN ID field. For more information about VLAN tags, see Working with VLAN Tags in the Host Profile on page 972.
Hops	Specify the number of network hops from the hosts you want to view to the sensor that detected them.
Last Seen	Specify the date and time the host was last detected by RNA or scanned by Nmap, or when operating system data was originally updated using the host input feature. See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time. Note that the last seen value is updated at least as often as the update interval you configured in the RNA detection policy, as well as when any instance of RNA managed by the Defense Center generates a new host event for the host's IP address. For information on setting the update interval, see What is an RNA Detection Policy? on page 889.
OS Vendor	Specify the vendor of the operating system that is running on the hosts you want to view. You can use an asterisk (*) as a wildcard character in this field. **TIP!** Type `unknown` to search for hosts where the operating system vendor is unknown. Type `pending` to search for hosts where the operating system vendor has not yet been identified.
OS Name	Specify the name of the operating system that is running on the hosts you want to view. You can use an asterisk (*) as a wildcard character in this field. **TIP!** Type `unknown` to search for hosts where the operating system name is unknown. Type `pending` to search for hosts where the operating system name has not yet been identified.
OS Version	Specify the version of the operating system that is running on the hosts you want to view. You can use an asterisk (*) as a wildcard character in this field.
Confidence	Specify the level of confidence that RNA has in its assessment of the host operating system for the hosts you want to view. You can precede the confidence with greater than (`>`), greater than or equal to (`>=`), less than (`<`), less than or equal to (`<=`), or equal to (`=`) operators. For example, use `<=50` to view hosts with a confidence of 50% or less. **TIP!** Type `unknown` to search for hosts where RNA cannot determine the operating system identity, which includes hosts added to the network map based on NetFlow data.

Keysight, EX-1004
Page 1030

Host Search Criteria (Continued)

Field	Search Criteria Rules
Notes	Specify text in the host's profile notes. You can use an asterisk (*) as a wildcard character in this field. For more information, see Adding Notes to a Host on page 986.
Source Type	Specify 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 system policy) • RNA, for hosts detected by 3D Sensors with RNA RNA may reconcile data from multiple sources to determine the identity of an operating system; see Understanding Current Identities on page 1360.
Detection Engine	Specify the name of the detection engine that detected the hosts you want to view, or, for flows exported by NetFlow-enabled devices, the detection engine that processed the NetFlow data that added the host to the network map. For more information, see Specifying Detection Engine/Appliance Names on page 1559.
OS Conflict	You have two options: • Type `conflict` to display all hosts with an operating system identity conflict. • Type `!conflict` to display all hosts that do **not** have an operating system identity 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 956.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

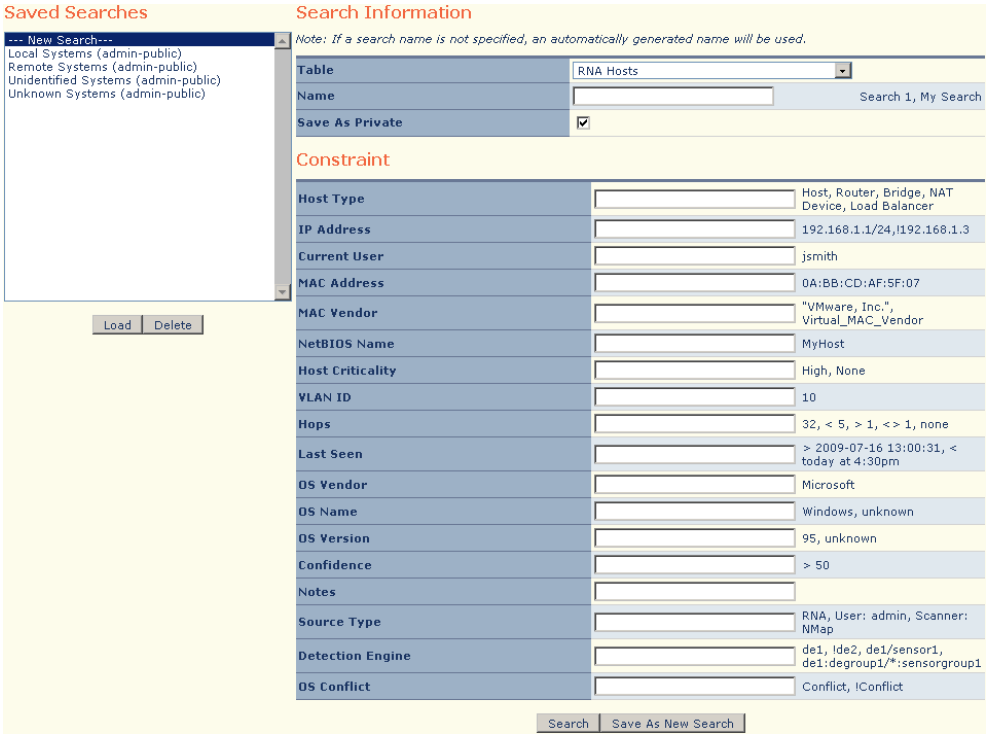
Keysight, EX-1004
Page 1031

To search for hosts:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Hosts**.

   The Hosts search page appears.



> **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.

Keysight, EX-1004
Page 1032

3. 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.

   TIP! 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

4. 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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

Platform: DC

Component: RNA

RNA 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 compliance rule. You can also set attribute values in response to a compliance rule.

For more information, see:

- Viewing Host Attributes on page 1034
- Understanding the Host Attributes Table on page 1035
- Setting Host Attributes for Specific Hosts on page 1036
- Searching for Host Attributes on page 1038
- Configuring Set Attribute Remediations on page 1340

Keysight, EX-1004
Page 1033

Viewing Host Attributes

Platform: DC

Component: RNA

You can use the Defense Center to view a table of hosts detected by RNA, 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 RNA host attributes differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of RNA 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 1618.

The RNA 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 RNA Event Actions table on page 1003.

RNA 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 1035.
drill down to the next page in the workflow	use one of the following methods: • 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 that you created in a custom workflow. 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.

Keysight, EX-1004
Page 1034

RNA Host Attribute Actions (Continued)

To...	You can...
assign a host attribute to selected hosts	find more information in Setting Host Attributes for Specific Hosts on page 1036
search for hosts with specific host attributes	click **Search**. For more information, see Searching for Host Attributes on page 1038.

To view host attributes:

Access: Any RNA/Admin

- Select **Analysis & Reporting > RNA > Host Attributes**.

  The first page of the default host attributes workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

> **TIP!** If you are using a custom workflow that does not include the table view of host attributes, from the **Workflows** menu on the toolbar, select **Attributes**.

Understanding the Host Attributes Table

Platform: DC

Component: RNA

RNA 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 compliance rule.

For more information on host attributes, see Working with the Predefined Host Attributes on page 985 and Working with User-Defined Host Attributes on page 987.

Keysight, EX-1004
Page 1035

The fields in the host attributes table are described in the Host Attribute Fields table.

Host Attribute Fields

Field	Description
IP Address	The IP address of a host.
Current User	The user identity (username) of the currently logged in user on the host.
Host Criticality	The user-assigned importance of a host to your enterprise. You can use the host criticality in compliance 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 Assigning a Host Criticality Value to a Host on page 986 and Setting Host Attributes for Specific Hosts on page 1036.
Notes	Information about the host that you want other analysts to view. For information on how to add note, see Adding Notes to a Host on page 986.
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 987.
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 Specific Hosts

Platform: DC

Component: RNA

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 compliance 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 compliance 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,

Keysight, EX-1004
Page 1036

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 989.

TIP! 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 Assigning a Host Criticality Value to a Host on page 986 or Configuring Set Attribute Remediations on page 1340.

To set host attributes for selected hosts:

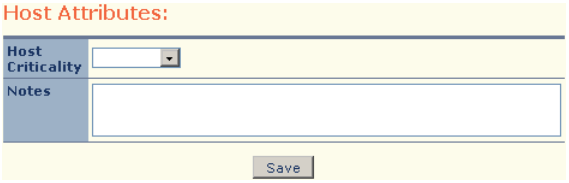
Access: RNA/Admin

1. Select the checkboxes 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**.

   A pop-up window appears.



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.

Keysight, EX-1004
Page 1037

Searching for Host Attributes

Platform: DC

Component: RNA

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 987.

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 Host Attribute Search Criteria table.

Host Attribute Search Criteria

Search Field	Description
IP Address	Type a specific IP address or use CIDR notation to specify a range of IP addresses for the hosts you want to view. See Specifying IP Addresses in Searches on page 1557 for a full description of the syntax allowed for IP addresses.
Host Criticality	Specify the host criticality of the hosts you want to view by typing `None`, `Low`, `Medium`, or `High`. For more information on host criticality, see Setting Host Attributes for Specific Hosts on page 1036.
Notes	Specify text in the host's profile notes. You can use an asterisk (*) as a wildcard character in this field. For more information, see Adding Notes to a Host on page 986.
Current User	Specify all or part of the user identity (user name) for the user logged into the hosts you want to view. Specify multiple user identities in a comma-separated list. This field is case insensitive. **TIP!** Use `blank` or `unknown` to search for hosts that do not have associated user identities; use `!blank` or `!unknown` to search for only the hosts that do have associated user identities.
any user-defined host attribute	Specify the appropriate value, which depends on the type of host attribute you select: • If the host attribute type is Integer, enter an integer 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 string from the drop-down list. • If the host attribute type is URL, enter a URL. For more information on host attributes, see Working with User-Defined Host Attributes on page 987.

Keysight, EX-1004
Page 1038

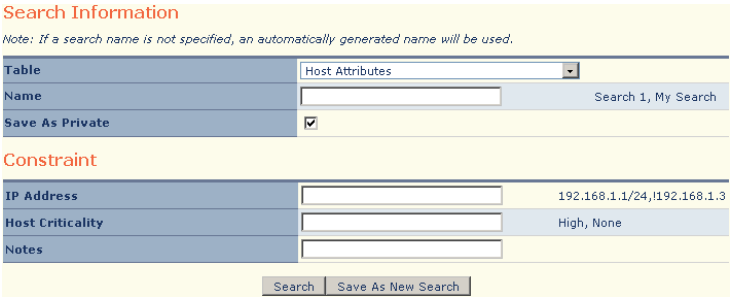
For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

To search for host attributes:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Host Attributes**.

   The Host Attributes search page appears.



   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 Host Attribute 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 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 host attributes workflow. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

Keysight, EX-1004
Page 1039

- 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 Services

Platform: DC

Component: RNA

RNA collects information about all services run by hosts on monitored network segments. The information that RNA collects includes the name of the service, the protocol used by the service, the IP address of the host running a service, and the port on which the service is running.

When RNA detects a running service, it generates a service event and logs it to the database. You can use the Defense Center web interface to view, search, and delete RNA service events.

You can also base compliance rules on service events. For example, you could trigger a compliance rule when RNA detects a chat service, such as ircd, running on one of your hosts.

For a full list of the services that the Sourcefire 3D System detects, see Detected Services on page 2027. Keep in mind, however, that the detector list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed, as well as on any individual add-on detectors you may have imported. 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 RNA detectors.

IMPORTANT! Although you can configure the RNA detection policy to add services to the network map based on data exported by NetFlow-enabled devices, the available information about these services is limited. For more information, see What Information Does RNA Provide? on page 880.

See the following sections for more information.

- Viewing Services on page 1041
- Understanding the Services Table on page 1043
- Searching for Services on page 1045
- Editing Service Identities on page 965

Keysight, EX-1004
Page 1040

Viewing Services

Platform: DC

Component: RNA

You can use the Defense Center to view a table of detected services. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access RNA services differs depending on the workflow you use. There are three predefined workflows:

- The Network Services by Count workflow provides a series of pages that constrain the detected services based on how many hosts are running particular services, and further constrain by the vendor and version of those services.
- The Network Service by Hit workflow provides a series of pages that constrain the detected services based on how often they are accessed, and further constrain by the vendor and version of those services.
- The RNA Services workflow includes a table view of RNA services. The table view contains a row for every service discovered on each detected host, organized by IP address and port. Each row contains data from a single service discovery event.

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 1618.

Keysight, EX-1004
Page 1041

The RNA Service Actions table below describes some of the specific actions you can perform on an RNA services workflow page. You can also perform the tasks described in the Common RNA Event Actions table on page 1003.

RNA Service Actions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the Services Table on page 1043.
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
search for services	click **Search**. For more information, see Searching for Services on page 1045.
edit service identities	select the checkboxes next to the events for services you want to edit, then click **Set Service Identity**. For more information, see Editing Service Identities on page 965.

Keysight, EX-1004
Page 1042

To view services:

Access: Any RNA/Admin

▶ Select **Analysis & Reporting > RNA > Services**.

The first page of the default services workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of services, from the **Workflows** menu on the toolbar, select **RNA Services**.

Understanding the Services Table

Platform: DC

Component: RNA

RNA collects information about services run by hosts on monitored network segments. See Detected Services on page 2027 for a list of services that RNA detects.

IMPORTANT! Although you can configure the RNA detection policy to add services to the network map based on data exported by NetFlow-enabled devices, the available information about these services is limited. For more information, see What Information Does RNA Provide? on page 880.

The fields in the services table are described in the Service Fields table.

Service Fields

Field	Description
Last Used	The date and time the service was last used on the network or the date and time that the service 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 RNA detection policy, as well as when RNA detects a service information update. For information on setting the update interval, see What is an RNA Detection Policy? on page 889.
IP Address	The IP address of the host running the service.
Current User	The user identity (user name) of the currently logged in user on the host where the service is running.
Port	The port where the service is running.

Keysight, EX-1004
Page 1043

Service Fields (Continued)

Field	Description
Protocol	The protocol used by the service.
Hits	The number of times the service was accessed. For services added using the host input feature, this value is always 0.
Service	One of: • the name of the service • `pending`, if RNA cannot positively or negatively identify the service for one of several reasons • `unknown`, if RNA cannot identify the service based on known service fingerprints For detailed information on how the Defense Center determines and displays service identities, see the Sourcefire 3D System Identification of Services table on page 1387.
Vendor	One of: • the service vendor as identified by RNA, Nmap or another active source, or that you specified using the host input feature • `pending`, if RNA does not yet have enough information to identify its vendor • `unknown`, if RNA cannot identify its vendor based on known service fingerprints, or if the service was added to the network map using NetFlow data
Version	One of: • the service version as identified by RNA, Nmap or another active source, or that you specified using the host input feature • `pending`, if RNA does not yet have enough information to identify its version • `unknown`, if RNA cannot identify its version based on known service fingerprints, or if the service was adding to the network map using NetFlow data
Payload	If the service is http, the content detected by RNA in the http traffic. Payload information is comprised of a payload type, which represents the general content type (for example, audio or video), as well as a payload, which represents the specific type of content (for example, WMV or QuickTime). For a list of detected payloads and payload types, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.

Keysight, EX-1004
Page 1044

Service Fields (Continued)

Field	Description
Source Type	One of the following values: • User: *user_name* • Application: *app_name* • Scanner: *scanner_type* (Nmap or scanner added through system policy) • RNA, RNA Port Match, or RNA Pattern Match, for services detected by 3D Sensors with RNA • NetFlow, for services added to the network map based on NetFlow data RNA may reconcile data from multiple sources to determine the identity of a service; see Understanding Current Identities on page 1360.
Confidence	One of: • the percentage of confidence that RNA has in the identity of the service, for hosts detected by 3D Sensors with RNA • 100%, for services identified by an active source, such as the host input feature or Nmap scanner • unknown, for services for which RNA cannot determine an identity, and for services added to the network map based on NetFlow data
Detection Engine	The name of the detection engine that either detected the service or processed the NetFlow data that added the service to the network map.
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 Services

Platform: DC

Component: RNA

You can search for specific services 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. The default searches are:

- Non-standard searches for many of the services detected by RNA (see Detected Services on page 2027 for a list of detected services), which searches for services running on non-standard ports.
- Possible MSSQL search, which searches for TCP traffic on port 1433, indicating possible Microsoft SQL Server activity.
- Possible MySQL search, which searches for TCP traffic on port 3306, indicating possible MySQL activity.

Keysight, EX-1004
Page 1045

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 Service Search Criteria table.

IMPORTANT! When searching for services, you should keep in mind that although you can configure the RNA detection policy to add services to the network map based on data exported by NetFlow-enabled devices, the available information about these service is limited. For more information, see What Information Does RNA Provide? on page 880.

Service Search Criteria

Field	Search Criteria Rules
IP Address	Type a specific IP address or use CIDR notation to specify a range of IP addresses for the services whose events you want to view. See Specifying IP Addresses in Searches on page 1557 for a full description of the syntax allowed for IP addresses.
Current User	Specify all or part of the user identity (user name) for the user logged into the host where a service is running. Specify multiple user identities in a comma-separated list. This field is case insensitive. **TIP!** Use `blank` or `unknown` to search for services that do not have associated user identities; use `!blank` or `!unknown` to search for only the services that do have associated user identities.
Port	Specify the port on which the service is running. For more information, see Specifying Ports in Searches on page 1558.
Protocol	Specify the protocol for which you want to view service events.
Hits	Specify the minimum number of times the service has been accessed, or for services added using the host input feature, type 0.
Last Used	Specify the date and time the service was last detected by RNA or the date and time the service was originally updated using the host input feature. See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time. Note that the Last Used value is updated at least as often as the update interval you configured in the RNA detection policy, as well as when RNA detects a service information update. For information on setting the update interval, see What is an RNA Detection Policy? on page 889.

Keysight, EX-1004
Page 1046

Service Search Criteria (Continued)

Field	Search Criteria Rules
Service	Specify the service you want to view. For a list of services detected by RNA, see Detected Services on page 2027. **TIP!** Type `pending` to search for services that RNA has not yet identified; type `unknown` to search for services that RNA cannot identify. For detailed information on how the Defense Center determines and displays service identities, see the Sourcefire 3D System Identification of Services table on page 1387.
Vendor	Specify the vendor of the services you want to view. You can use an asterisk (*) as a wildcard character in this field. **TIP!** Type `pending` to search for services where RNA has not yet identified the service vendor; type `unknown` to search for services where RNA cannot identify the vendors.
Version	Specify the version of the services you want to view. You can use an asterisk (*) as a wildcard character in this field. **TIP!** Type `pending` to search for services where RNA has not yet identified the service version; type `unknown` to search for services where RNA cannot identify the versions.
Payload	Specify the payload or payload type for the http services you want to view. Payload information is comprised of a payload type, which represents the general content type (for example, audio or video), as well as a payload, which represents the specific type of content (for example, WMV or QuickTime). For a list of payloads and payload types that you can specify, see Detected Payloads on page 2039.
Confidence	Specify the level of confidence that RNA has in its assessment of the service you want to view. You can precede the confidence with greater than (`>`), greater than or equal to (`>=`), less than (`<`), less than or equal to (`<=`), or equal to (`=`) operators. For example, use `<=50` to view services with a confidence of 50% or less. **TIP!** Type `unknown` to search for services for which RNA cannot determine an identity, which includes services added to the network map based on NetFlow data.

Keysight, EX-1004
Page 1047

Service Search Criteria (Continued)

Field	Search Criteria Rules
Source Type	One of the following values: • User: *user_name* • Application: *app_name* • Scanner: *scanner_type* (Nmap or scanner added through system policy) • RNA, RNA Port Match, or RNA Pattern Match, for services detected by 3D Sensors with RNA • NetFlow, for services added to the network map based on NetFlow data RNA may reconcile data from multiple sources to determine the identity of a service; see Understanding Current Identities on page 1360.
Detection Engine	Specify the name of the detection engine that detected the hosts you want to view, or, for flows exported by NetFlow-enabled devices, the detection engine that processed the NetFlow data that added the host to the network map. For more information, see Specifying Detection Engine/Appliance Names on page 1559.
Service Conflict	You have two options: • Type `conflict` to display all services with an identity conflict. • Type `!conflict` to display all services that do **not** have any identity conflicts. Note that the Service Conflict column does not appear in search results. To determine whether you are viewing services with or without identity conflicts, expand the search constraints on the workflow page. For more information on resolving service identity conflicts, see Resolving Service Identity Conflicts on page 967.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

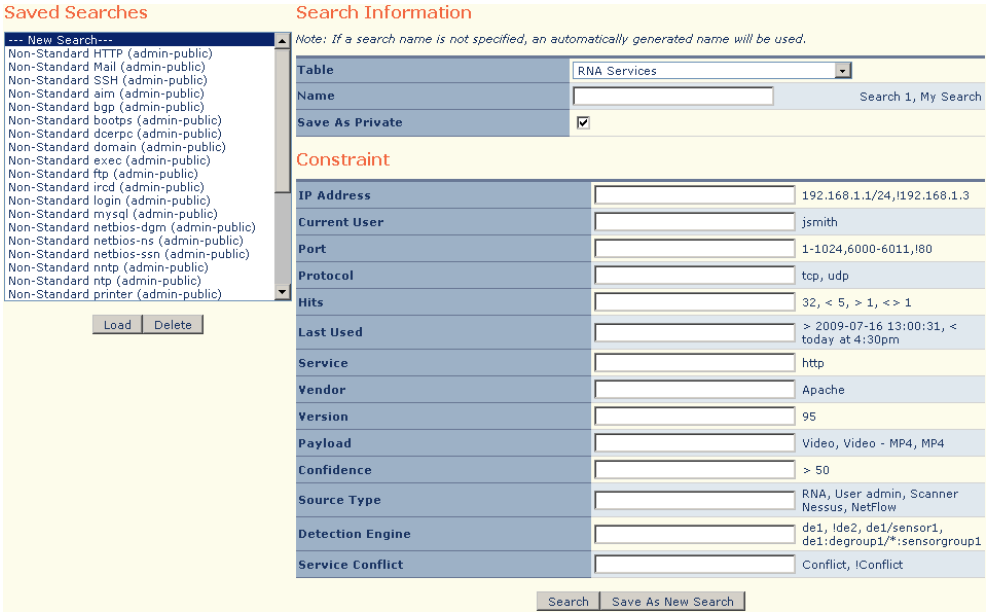
Keysight, EX-1004
Page 1048

To search for services:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Services**.

   The Services search page appears.



> **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 Service Search Criteria table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria.

> **TIP!** 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

Keysight, EX-1004
Page 1049

4. You have the following options:
   - Click **Search** to start the search.

     Your search results appear in the default services workflow. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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 Client Applications

Platform: DC

Component: RNA

When a monitored host connects to another host, RNA can, in many cases, determine what application was used. RNA detects the use of many email, instant messaging, peer to peer, web clients, as well as other types of applications.

For each detected client application, RNA 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 RNA client application events. You can also update client application data on a host or hosts using the host input feature.

If you know which client 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 compliance rule. You can also base compliance rules on the detection of client applications. For example, if you want your employees to use a specific mail client, you could trigger a compliance rule when RNA detects a different mail client running on one of your hosts.

For a full list of the client applications that the Sourcefire 3D System detects, see Detected Client Applications on page 2036. Keep in mind, however, that the detector list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed, as well as on any individual add-on detectors you may have imported. You should carefully read the release notes for each

Keysight, EX-1004
Page 1050

Sourcefire 3D System update as well as the advisories for each VDB update for information on updated RNA detectors.

IMPORTANT! To collect and store client application data for analysis, make sure that you enable client application detection in your RNA detection policy. For more information, see What is an RNA Detection Policy? on page 889.

See the following sections for more information:

- Viewing Client Applications on page 1051
- Understanding the Client Applications Table on page 1053
- Searching for Client Applications on page 1054

Viewing Client Applications

Platform: DC

Component: RNA

You can use the Defense Center to view a table of detected client applications. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access RNA client applications differs depending on the workflow you use. There are two predefined workflows:

- The Client Application Summaries workflow provides a series of pages that constrain the detected client applications based their name, version, and when they were last used. It also contains a table view of client applications, which contains a row for each client application discovered on each detected host. Each row includes a count of the number of times the client application was detected on the host.
- The RNA Client Applications workflow includes a table view of RNA client applications. The table view contains a row for each client application discovered on each detected host. Each row includes a count of the number of times the client application was detected on the host.

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 1618.

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 1618.

The RNA Client Application Actions table below describes some of the specific actions you can perform on an RNA client applications workflow page. You can

Keysight, EX-1004
Page 1051

also perform the tasks described in the Common RNA Event Actions table on page 1003.

RNA Client Application Actions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the Client Applications Table on page 1053.
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
search for client applications	click **Search**. For more information, see Searching for Client Applications on page 1054.

To view client applications:

Access: Any RNA/Admin

▶ Select **Analysis & Reporting > RNA > Client Applications**.

The first page of the default client applications workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of client applications, from the **Workflows** menu on the toolbar, select **RNA Client Applications**.

Keysight, EX-1004
Page 1052

Understanding the Client Applications Table

Platform: DC

Component: RNA

When a monitored host connects to another host, RNA can, in many cases, determine what application was used. RNA detects various web browsers, email clients, instant messengers, peer-to-peer applications, and so on; for a full list see Detected Client Applications on page 2036.

When RNA detects traffic for a known client application, RNA logs information about the application and the host running it. The fields in the client applications table are described in the Client Application Fields table.

Client Application Fields

Field	Description
Last Used	The time that the application was last used or the time that the client 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 RNA detection policy, as well as when RNA detects a client application information update. For information on setting the update interval, see What is an RNA Detection Policy? on page 889.
IP Address	The IP address of the host using the application.
Current User	The user identity (username) of the currently logged in user on the host with the client application is running.
Application Type	The type of application. RNA detects various web browsers, email clients, and instant-messaging clients.
Application	The name of the application.
Version	The version of the application.
Payload	For web browsers, the content detected by RNA in the http traffic. Payload information is comprised of a payload type, which represents the general content type (for example, audio or video), as well as a payload, which represents the specific type of content (for example, WMV or QuickTime). For a list of detected payloads and payload types, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.

Keysight, EX-1004
Page 1053

Client Application Fields (Continued)

Field	Description
Hits	The number of times RNA detected the application in use. For applications added using the host input feature, this value is always 0.
Detection Engine	The detection engine that generated the client application event.
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 Client Applications

Platform: DC

Component: RNA

You can search for hosts that are running specific client applications. 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 Client Application Search Criteria table.

Client Application Search Criteria

Search Field	Description
Last Used	Specify the date and time the application was last detected by RNA or the date and time when the client application was originally added using the host input feature. See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time. Note that the Last Used value is updated at least as often as the update interval you configured in the RNA detection policy, as well as when RNA detects a client application information update. For information on setting the update interval, see What is an RNA Detection Policy? on page 889.
IP Address	Type a specific IP address or use CIDR notation to specify a range of IP addresses for the hosts whose client applications you want to view. See Specifying IP Addresses in Searches on page 1557 for a full description of the syntax allowed for IP addresses.
Current User	Specify all or part of the user identity (user name) for the user logged into the host where a client application is running. Specify multiple user identities in a comma-separated list. This field is case insensitive. **TIP!** Use `blank` or `unknown` to search for client applications that do not have associated user identities; use `!blank` or `!unknown` to search for only the client applications that do have associated user identities.
Application Type	Specify the type of application that is running on the hosts you want to view.

Keysight, EX-1004
Page 1054

Client Application Search Criteria (Continued)

Search Field	Description
Application	Specify the name of the application that is running on the hosts you want to view. You can use an asterisk (*) as a wildcard character in this field.
Version	Specify the version of the application that is running on the hosts you want to view. You can use an asterisk (*) as a wildcard character in this field.
Hits	Specify the minimum number of times the application has been accessed, or for applications added using the host input feature, type 0.
Payload	Specify the payload or payload type for the web browser client applications you want to view. Payload information is comprised of a payload type, which represents the general HTTP content type (for example, audio or video), as well as a payload, which represents the specific type of content (for example, WMV or QuickTime). For a list of payloads and payload types that you can specify, see Detected Payloads on page 2039.
Detection Engine	Specify the name of the detection engine that detected the client application you want to view. For more information, see Specifying Detection Engine/ Appliance Names on page 1559.

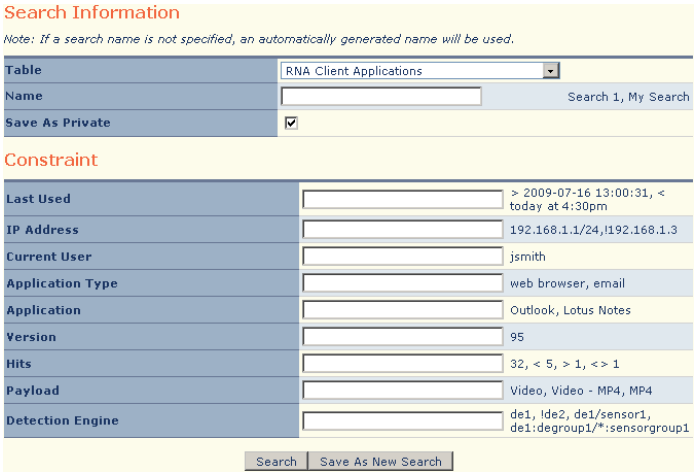
Keysight, EX-1004
Page 1055

To search for client applications:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Client Applications**.

   The Client Applications search page appears.



   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 Client Application Search Criteria table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria.

   TIP! 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

Keysight, EX-1004
Page 1056

4. You have the following options:
   - Click **Search** to start the search.

     Your search results appear in the default client applications workflow. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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

Platform: DC

Component: RNA

RNA includes its own vulnerability tracking database which is used, in conjunction with RNA's operating system fingerprinting capability, to identify the known vulnerabilities of each fingerprinted host in a monitored network.

Different operating systems have different sets of vulnerabilities. Each host with an identified operating system has a set 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, see:

- Viewing Sourcefire Vulnerabilities on page 1057
- Understanding the Sourcefire Vulnerabilities Table on page 1059
- Deactivating Sourcefire Vulnerabilities on page 1061
- Searching for Sourcefire Vulnerabilities on page 1062

Viewing Sourcefire Vulnerabilities

Platform: DC

Component: RNA

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 RNA vulnerabilities differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of RNA vulnerabilities. The table view contains a row for each vulnerability in the database, regardless of whether any of your detected hosts exhibit the

Keysight, EX-1004
Page 1057

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 1062.

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 1618.

The RNA Vulnerability Actions table below describes some of the specific actions you can perform on an RNA vulnerabilities workflow page. You can also perform the tasks described in the Common RNA Event Actions table on page 1003.

RNA 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 1059.
drill down to the next page in the workflow	use one of the following methods: • 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 that you created in a custom workflow. 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.

Keysight, EX-1004
Page 1058

RNA 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 1061.
search for vulnerabilities	click **Search**. For more information, see Searching for Sourcefire Vulnerabilities on page 1062.

To view vulnerabilities:

Access: Any RNA/Admin

- Select **Analysis & Reporting > RNA > Vulnerabilities**.

  The first page of the default vulnerabilities workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

  TIP! If you are using a custom workflow that does not include the table view of vulnerabilities, from the **Workflows** menu on the toolbar, select **Vulnerabilities**.

Understanding the Sourcefire Vulnerabilities Table

Platform: DC

Component: RNA

RNA includes its own vulnerability tracking database which is used in conjunction with RNA's operating system fingerprinting capability to identify the known vulnerabilities of each fingerprinted host in a monitored network.

Different operating systems have a different sets of vulnerabilities. Each host with an identified operating system has a set of associated vulnerabilities. You can deactivate vulnerabilities for a host when you patch the host or otherwise judge it immune. You can also use the Defense Center to track and review the vulnerabilities for each host.

The first page of the default vulnerabilities workflow displays every vulnerability in the database, regardless of whether it applies to any of your hosts. The second page of the default workflow provides you with a table view of vulnerabilities that do apply to detected hosts on your network.

Keysight, EX-1004
Page 1059

For more information on vulnerabilities, see Working with the Vulnerabilities Network Map on page 932 and Working with Vulnerabilities in the Host Profile on page 978.

The fields in the vulnerabilities table are described in the Vulnerabilities Table Fields table.

Vulnerabilities Table Fields

Field	Description
RNA ID	The RNA Vulnerability identification number that RNA uses to track vulnerabilities.
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.
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.
IP Address	The IP address of host affected by 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.

Keysight, EX-1004
Page 1060

Deactivating Sourcefire Vulnerabilities

Platform: DC

Component: RNA

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 RNA 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 932.
- Use the host's host profile.

  For more information, see Setting Vulnerabilities for Individual Hosts on page 984.
- Constrain the vulnerabilities workflow based on the IP address of the host or hosts for which you want to deactivate vulnerabilities.

TIP! To constrain the Vulnerabilities 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 1062.

To deactivate vulnerabilities:

Access: RNA/Admin

1. Select the check boxes next to vulnerabilities you want to deactivate, then click **Review**.

Keysight, EX-1004
Page 1061

Searching for Sourcefire Vulnerabilities

Platform: DC

Component: RNA

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. The search criteria you can use are described in the Vulnerabilities Search Criteria table.

Vulnerabilities Search Criteria

Search Field	Description
RNA ID	Specify one or more RNA vulnerability identification numbers to match. Specify multiple RNA IDs in a comma-separated list.
Bugtraq ID	Specify one or more Bugtraq identification numbers to match. Specify multiple Bugtraq IDs in a comma-separated list. Find Bugtraq ID numbers at http://www.securityfocus.com/bid.
Snort ID	Specify one or more Snort IDs (SIDs) to match. Specify multiple SIDs in a comma-separated list.
Title	Specify one or more words to search for in the vulnerability title. The search is case insensitive. You can use an asterisk (*) as a wildcard character in this field.
IP Address	Specify the IP address of the host or hosts whose vulnerabilities you want to view. See Specifying IP Addresses in Searches on page 1557 for information about entering IP addresses.
Date Published	Specify the date the vulnerability was published. See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time.
Vulnerability Impact	Specify a vulnerability impact, from 0 to 10, with 10 being the most severe. Vulnerability impact is the severity assigned to the vulnerability in the Bugtraq database.
Remote	Enter TRUE to search for vulnerabilities that are exploited over a network, or FALSE to exclude such vulnerabilities.
Available Exploits	Enter TRUE to search for vulnerabilities for which an exploit is known to exist, or FALSE to exclude such vulnerabilities.

Keysight, EX-1004
Page 1062

Vulnerabilities Search Criteria (Continued)

Search Field	Description
Description	Specify one or more words to search for in the general description of the vulnerability. The search is case insensitive. You can use an asterisk (*) as a wildcard character in this field.
Solution	Specify one or more words to search for in the solution of the vulnerability. You can use an asterisk (*) as a wildcard character in this field.
Technical Description	Specify one or more words to search for in the detailed technical description of the vulnerability. The search is case insensitive. You can use an asterisk (*) as a wildcard character in this field.

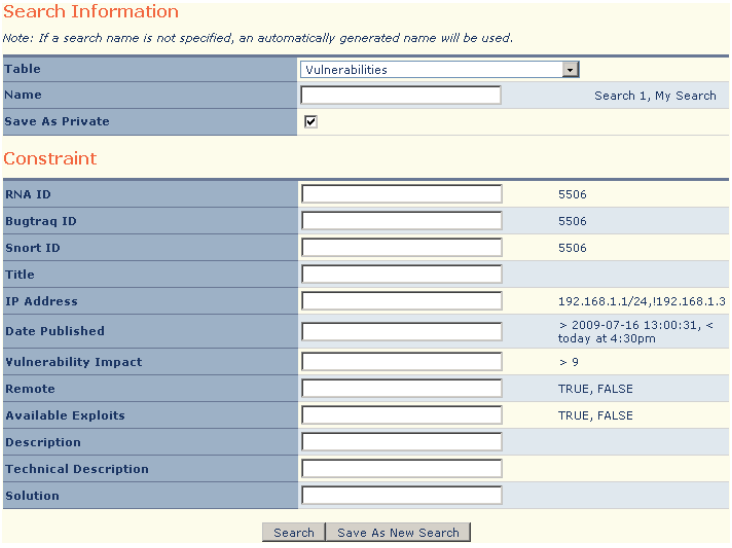
For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

To search for vulnerabilities:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Vulnerabilities**.

   The Vulnerabilities search page appears.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

Keysight, EX-1004
Page 1063

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 Vulnerabilities Search Criteria table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria.

   TIP! 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

4. 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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

Platform: DC

Component: RNA

RNA includes its own vulnerability tracking database which is used, in conjunction with RNA's operating system fingerprinting capability, to identify the known vulnerabilities of each fingerprinted host in a monitored 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 RNA's vulnerability data. For more information, see the *Sourcefire 3D System Host Input API Guide*.

Note that to include imported data in impact correlations, you must map third-party vulnerability information to the operating system and service definitions in the RNA database.

Keysight, EX-1004
Page 1064

For more information, see:

- Viewing Third-Party Vulnerabilities on page 1065
- Understanding the Third-Party Vulnerabilities Table on page 1067
- Searching for Third-Party Vulnerabilities on page 1068

Viewing Third-Party Vulnerabilities

Platform: DC

Component: RNA

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:

- The Vulnerabilities by Source workflow provides a quick view of how many third-party vulnerabilities you have detected per third-party vulnerability source, plus additional details about those vulnerabilities.
- The Vulnerabilites by IP Address workflow provides a quick view of how many third-party vulnerabilities you have detected per host on your monitored network.

Both predefined workflows terminate in a table view of third-party vulnerabilities and then 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 1618.

The RNA Third-Party Vulnerability Actions table below describes some of the specific actions you can perform on a third-party vulnerabilities workflow page.

Keysight, EX-1004
Page 1065

You can also perform the tasks described in the Common RNA Event Actions table on page 1003.

RNA 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 1067.
view the vulnerability details for a third-party vulnerability	click the vulnerability icon ( )that appears next to the vulnerability ID. For more information, see Viewing Vulnerability Details on page 980.
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
search for third-party vulnerabilities	click **Search**. For more information, see Searching for Third-Party Vulnerabilities on page 1068.

To view third-party vulnerabilities:

Access: Any RNA/Admin

▶ Select **Analysis & Reporting > RNA > Third-Party Vulnerabilities**.

The first page of the default third-party vulnerabilities workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of third-party vulnerabilities, from the **Workflows** menu on the toolbar, select **Vulnerabilities by Source** or **Vulnerabilities by IP Address**.

Keysight, EX-1004
Page 1066

Understanding the Third-Party Vulnerabilities Table

Platform: DC

Component: RNA

When you import third-party vulnerability information using the host input feature, RNA stores that information in its database. The fields in the third-party vulnerabilities table are described in the Third-Party Vulnerabilities Table Fields table.

Third-Party Vulnerabilities Table Fields

Field	Description
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 of host affected by the vulnerability.
Port	A port number, if the vulnerability is associated with a service 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/).
RNA ID	The RNA Vulnerability identification number that RNA 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 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.

Keysight, EX-1004
Page 1067

Searching for Third-Party Vulnerabilities

Platform: DC

Component: RNA

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. The search criteria you can use are described in the Third-Party Vulnerabilities Search Criteria table.

Third-Party Vulnerabilities Search Criteria

Search Field	Description
Vulnerability Source	Specify the source of the third-party vulnerabilities you want to view, for example, QualysGuard or NeXpose.
Vulnerability ID	Specify a vulnerability identification number associated with the Vulnerability Source to match. Specify multiple Vulnerability IDs in a comma-separated list.
IP Address	Specify the IP address of the host or hosts whose third-party vulnerabilities you want to view. See Specifying IP Addresses in Searches on page 1557 for information about entering IP addresses.
Port	Specify a port to view third-party vulnerabilities associated with a service running on that port. For more information, see Specifying Ports in Searches on page 1558.
Bugtraq ID	Specify one or more Bugtraq identification numbers to match. Specify multiple Bugtraq IDs in a comma-separated list. Find Bugtraq ID numbers at http://www.securityfocus.com/bid.
CVE ID	Specify one or more CVE identification numbers to match. Specify multiple CVE IDs in a comma-separated list. Find CVE ID numbers in MITRE's CVE database at (http://www.cve.mitre.org/.
RNA ID	Specify one or more RNA vulnerability identification numbers to match. Specify multiple RNA IDs in a comma-separated list.
Snort ID	Specify one or more Snort IDs (SIDs) to match. Specify multiple SIDs in a comma-separated list.
Title	Specify one or more words to search for in the third-party vulnerability title. The search is case insensitive. You can use an asterisk (*) as a wildcard character in this field.
Description	Specify one or more words to search for in the general description of the third-party vulnerability. The search is case insensitive. You can use an asterisk (*) as a wildcard character in this field.

Keysight, EX-1004
Page 1068

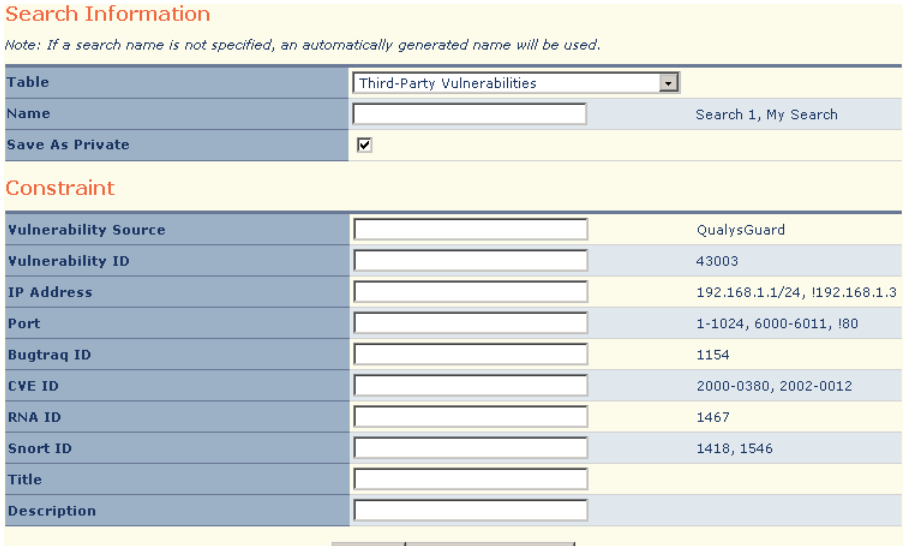
For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

To search for third-party vulnerabilities:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Third-Party Vulnerabilities**.

   The Third-Party Vulnerabilities search page appears.



> **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 Third-Party Vulnerabilities Search Criteria table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria.

> **TIP!** 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

Keysight, EX-1004
Page 1069

4. 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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.

Keysight, EX-1004
Page 1070

Chapter 25
Working with Flow Data and Traffic Profiles

If your Sourcefire 3D System deployment includes 3D Sensors with RNA, you can continuously monitor and collect data on the traffic generated by each host on your network. You can also supplement this data with IP traffic information collected by NetFlow-enabled devices.

RNA generates a flow event when a connection between a monitored host and any other host is terminated. A flow event includes detailed information about the collected traffic, including the timestamps of the first and last packet of the transaction, the source and destination IP addresses, the ports and protocol used by the flow, the number of packets and bytes sent and received, user identity information, and more.

You can view flow data using a set of predefined workflows, or you can create custom workflows; flow data workflows include both graphs and tables.

IMPORTANT! You must use a Defense Center with an RNA host license to view flow data. You cannot view flow data on a Master Defense Center.

The Flow Summary dashboard uses flow data to create customizable graphs of the activity on your monitored network that provide you with an at-a-glance understanding of the state of your network. See Working with Dashboards on page 92 for more information. Note that Restricted Event Analysts use the Flow Summary page instead of the Flow Summary Dashboard; see Viewing the Flow Summary Page on page 1116.

You can also use flow data to create and view a profile of your normal network traffic, called a *traffic profile*. After you establish parameters for normal traffic, you

Keysight, EX-1004
Page 1071

can detect abnormal network traffic by evaluating new traffic against that profile within a compliance policy.

Settings in two policies—the RNA detection policy and the system policy—specify how the Sourcefire 3D System collects and aggregates flow data, as well as how flow data is transmitted from the 3D Sensors to the Defense Center.

For more information, see:

- Understanding Flow Data on page 1072
- Manipulating Flow Data Graphs on page 1100
- Viewing the Flow Summary Page on page 1116
- Creating Traffic Profiles on page 1118
- Viewing Traffic Profiles on page 1138

Understanding Flow Data

Platform: DC

Component: RNA

RNA continuously monitors traffic generated by the hosts on your network. To supplement the data gathered by RNA, you can use records generated by NetFlow-enabled devices. For example, if you have NetFlow-enabled devices deployed on networks that your 3D Sensors with RNA cannot monitor, you can use the data exported by those devices to monitor those networks.

What Is Flow Data?

RNA generates flow events when connections between monitored hosts and any other hosts are terminated. Flow events include information about the collected traffic, including:

- the type of flow (whether the flow was detected by a 3D Sensor with RNA, or by a NetFlow-enabled device), and if the flow was detected by a NetFlow-enabled device, the IP address of the device
- the timestamps of the first and last packet of the transaction
- the source IP address, destination IP address, and the ports and protocol used by the flow
- any TCP flags detected in the flow (for flows exported by NetFlow-enabled devices)
- the user logged into the source and destination hosts (if RUA is enabled and the hosts are on the monitored network)
- the number of packets and bytes sent and received by the monitored host
- the client application and URL involved in the transaction, if applicable

Keysight, EX-1004
Page 1072

Note that some of the information in a flow event depends on the flow type, that is, on whether the flow was detected by a 3D Sensor with RNA or by a NetFlow-enabled device:

- For flows detected directly by the detection engines on 3D Sensors with RNA, RNA generates a single bidirectional flow event. However, because NetFlow-enabled devices export unidirectional flow data, RNA generates at least two flow events for each flow detected by NetFlow-enabled devices, depending on how you configured the devices.
- Flow events based on NetFlow data do not contain NetBIOS information or any information on flows' client applications (type, name, or version) or any URLs accessed during monitored sessions.
- Hosts added to the network map based on NetFlow records do not include operating system or NetBIOS data for the hosts involved in the flow, nor can RNA identify if the hosts are network devices (bridges, routers, NAT devices, or load balancers).
- Services added to the network map based on NetFlow records are identified by the port-service correlation in `/etc/sf/services` to extrapolate service names because flow records exported by NetFlow do not contain service information. In addition, RNA is unable to determine vendors and versions for services added to the network map based on NetFlow data.
- Flow data exported by NetFlow does not contain initiator or responder information. When RNA 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.
- Flow data exported by NetFlow-enabled devices contains information on the TCP flags set for the flow. Flows detected directly by the detection engines on 3D Sensors with RNA do not contain TCP flag information.

For detailed information, see How is NetFlow Data Different from RNA Flow Data? on page 1142.

How is Flow Data Collected?

The RNA detection policy specifies how the RNA collects flow data, as well as how flow data is transmitted from the 3D Sensors to the Defense Center.

When you configure your RNA detection policy, you must set the flow data mode, which controls how flow data is transmitted from 3D Sensors with RNA to the Defense Center. Because RNA uses 3D Sensors to process the data exported by your NetFlow-enabled devices before it transmits the data to the Defense Center, the flow data mode affects NetFlow data as well as flow data collected directly by 3D Sensors.

Keysight, EX-1004
Page 1073

There are three flow data modes:

- **Enabled** configures your 3D Sensors to transmit individual flows to the Defense Center.
- **Summary** (the default) configures your 3D Sensors to transmit flow summaries, which are sets of flow data aggregated over five-minute intervals.
- **Disabled** prevents 3D Sensors from transmitting flow data to the Defense Center.

IMPORTANT! If you configure your 3D Sensors to transmit individual flows, the Defense Center performs its own data aggregation so that you can view flow data both as individual flows and as flow summaries. For more information on flow summaries, see Understanding Flow Summary Data on page 1076.

Also, for every network you monitor with an RNA detection engine (flow filtering is not supported for NetFlow data), you must set a data collection mode. You can collect host data and flow data for the hosts in each subnet, you can collect only host data (the default), or you can collect neither. When RNA detects a flow between two hosts, the information that RNA records in the database depends on the data collection mode you set for the networks where the hosts reside.

If you set the same data collection mode for two networks, RNA collects data for hosts in those networks consistent with the collection mode. For example, if you set the data collection mode for Networks A and B to host only, and then RNA detects a flow between a host in Network A (Host A) and a host in Network B (Host B), RNA does not log a record of the flow, but it does record general host information about the two hosts, such as host type and operating system data.

On the other hand, the following table explains what kind of data RNA records if the data collection modes for two networks differ.

Flow Filtering Consequences

For Host A, if you collect...	And for Host B, if you collect...	Then, if there is a flow between the two hosts, RNA collects...
host and flow data	host data only	host and flow data for both hosts
host and flow data	no data	host and flow data for Host A no data for Host B
host data only	no data	host data for Host A no data for Host B

Keysight, EX-1004
Page 1074

Disabling or otherwise limiting flow data collection for certain networks can reduce the traffic between your 3D Sensors and your Defense Center as well as reduce the space required to store flow data on the Defense Center. You may want to only turn on flow collection for subnets that are most critical to your operations, while continuing to collect host data for your entire network.

On the other hand, for networks where you have disabled flow collection, you cannot take advantage of any feature that relies on flow data, such as the Flow Summary dashboard, traffic profiles, compliance rules based on flows or traffic profile changes, and flow trackers in compliance rules.

If you allow host data collection only, you can still use the network map, view host profiles, perform vulnerability and intrusion event impact assessment, as well as use compliance rules and white lists. Excluding a network excludes it from monitoring altogether.

For more information, see Introduction to Sourcefire RNA on page 878.

How Can You View and Use Flow Data?

You can view flow data using a set of predefined workflows (see Predefined Flow Data Workflows on page 1582), or you can create custom workflows. Flow data workflows include both graphs and tables. For example, the Flows over Time workflow contains a graph of the total number of detected flows over time, as well as table views of flow summaries and of flow events.

The Flow Summary dashboard (see Using Dashboards on page 61) uses flow data to create customizable graphs of the activity on your monitored network that provide you with an at-a-glance understanding of the state of your network.

In addition to viewing flow data, you can use it to create and view a profile of your normal network traffic, called a traffic profile. After you establish parameters for normal traffic, you can detect abnormal network traffic by evaluating new traffic against that profile within a compliance policy. For more information, see Creating Compliance Policies on page 1262.

For more information on flow data and traffic profiles, see:

- Understanding Flow Summary Data on page 1076
- Understanding Flow Data Graphs on page 1078
- Understanding Flow Data Tables on page 1086
- Viewing Flow Data on page 1093
- Searching for Flow Data on page 1094

Keysight, EX-1004
Page 1075

Understanding Flow Summary Data

Platform: DC

Component: RNA

RNA uses specific criteria to aggregate flows into flow summaries, whether on a 3D Sensor before flow data is transmitted to the Defense Center, or on the Defense Center after it receives individual flows from 3D Sensor.

TIP! You can control when and how flow data is aggregated using the flow data mode. For more information, see How is Flow Data Collected? on page 1073.

To be aggregated, multiple flows must:

- have the same source and destination IP addresses, and use the same port on the responder (destination) host

  Because NetFlow records do not contain information about which host in the flow is the initiator and which is the responder, when RNA 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.
- use the same protocol (TCP or UDP)
- use the same service

  For flows detected by 3D Sensors with RNA, this is the service as detected by RNA. For flows exported by NetFlow-enabled devices, this is the service identified by the port-service correlation in `/etc/sf/services`.
- either be detected by the same detection engine (for flows detected by 3D Sensors with RNA) or be exported by the same NetFlow-enabled device and processed by the same detection engine

The aggregated data in a flow summary includes the total number of packets and bytes send by the initiator and responder hosts, as well as the number of flows in the summary. Note that because NetFlow-enabled devices generate unidirectional flows, a summary's flow count is incremented by two for every flow based on NetFlow data.

For more information on flow summary data contrasted to individual flow data, see Understanding Flow Data Tables on page 1086.

Long-Running Flows

If a monitored session spans two or more five-minute intervals over which flow data is aggregated, the flow is considered a *long-running flow*. When calculating the number of flows in a flow summary, RNA increments the count only for the five-minute interval in which a long-running flow was initiated.

Also, when calculating the number of packets and bytes transmitted by the initiator and responder in a long-running flow, RNA does not report the number of packets and bytes that were actually transmitted during each five-minute interval.

Keysight, EX-1004
Page 1076

Instead, RNA assumes a constant rate of transmission and calculates estimated figures based on the total number of packets and bytes transmitted, the length of the flow, and what portion of the flow occurred during each five minute interval.

Combining Flow Summaries from External Responders

In the RNA detection policy settings as well as in the system policy, you must decide whether you want to combine flows for out-of-network responders.

This setting combines flow summaries that involve a host on your monitored network and one or more hosts not on your monitored network, according to the criteria listed above. The web interface displays `external` instead of an IP address for the aggregated external hosts in the flow summary.

When the aggregation occurs depends on where you enable the option:

- Enabling this option in the RNA detection policy (see Configuring RNA Detection Policy Settings on page 902) configures your 3D Sensors with RNA to combine flow summaries involving external responders **before** the data is transmitted to the Defense Center.

  Note that this option has no effect unless the RNA detection policy specifies that the flow data mode is **Summary** (the default), which prevents you from taking advantage of any feature that requires data from individual flows.
- Enabling this option in the system policy (see Understanding RNA Data Storage Settings on page 1759) configures your Defense Center to combine flow summaries involving external responders **after** they are transmitted by your 3D Sensors with RNA.

  Note that if you enable this option in the system policy and you attempt to drill down to the table view of flow data (that is, access data on individual flows) for a flow summary that involves an external responder, the table view contains no information.

Enabling either option can reduce the space required to store flow data and speed up the rendering of flow data graphs. In addition, enabling this option in the RNA detection policy can reduce the number of events sent to the Defense Center.

TIP! You should combine flow summaries for external responders in the RNA detection policy if you kept the flow data mode as **Summary**. Otherwise, enable the option in the system policy.

Special Considerations for Flow Summaries

The primary consequence of using a flow data mode of **Summary** is that you cannot access data on individual flows; the table view of flow data contains no information.

Keysight, EX-1004
Page 1077

However, where possible, the Sourcefire 3D System treats individual flows and flow summaries identically. Therefore, even if the flow data mode is **Summary**, you can still view flow data graphs and create traffic profiles, create compliance rules that trigger when a flow summary meets criteria that you specify, and add flow trackers to compliance rules. But because you cannot access data on individual flows, you must make sure that when you build compliance rules or flow trackers based on flow summaries, you use data that the flow summaries contain.

For compliance rules, a trigger criterion based on information that is not included in a flow summary automatically evaluate to true. As an example, flow summaries do not include information on the initiator port. Therefore, if you build a compliance rule that requires that the initiator must be using a specific port, the rule will trigger for every flow summary that meets the other criteria in the rule. If that is the only criterion in the rule, the rule will trigger for every flow summary that the Defense Center receives.

Similarly, if you use information that is not contained in a flow summary when you are building a flow tracker, the Defense Center tracks all flow summaries that meet the other criteria in the flow tracker.

Finally, because you use the same page in the Defense Center web interface to search both flow summary data and individual flows, you must pay careful attention to whether your search constraints apply to the table you are searching. For example, client application data is not available in flow summaries. If you search for flow events based on the detected client application in the flow and then view the results in a flow summary event view, the Defense Center displays flow data as if you had not constrained it at all. Search constraints that do not apply to the current event view are labeled as not applicable (N/A) and are marked with a strikethrough, as shown in the following graphic.



For your convenience, the flow data search page marks fields that constrain flow summary event views and flow graphs (flow graphs are built using flow summary data) with an asterisk (*). For information on how to search flow data, see Searching for Flow Data on page 1094.

Understanding Flow Data Graphs

Platform: DC

Component: RNA

One of the ways the Defense Center can present flow data is graphically. There are three different types of flow data 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.

Keysight, EX-1004
Page 1078

You can manipulate flow data 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, services, ports, and detection engines

Because traffic profiles are based on flow data (see Creating Traffic Profiles on page 1118), you can view traffic profiles as line graphs. You can manipulate these graphs in the same way as you would any other flow graph, with some restrictions. Note, however, that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

For more information on flow graphs, see:

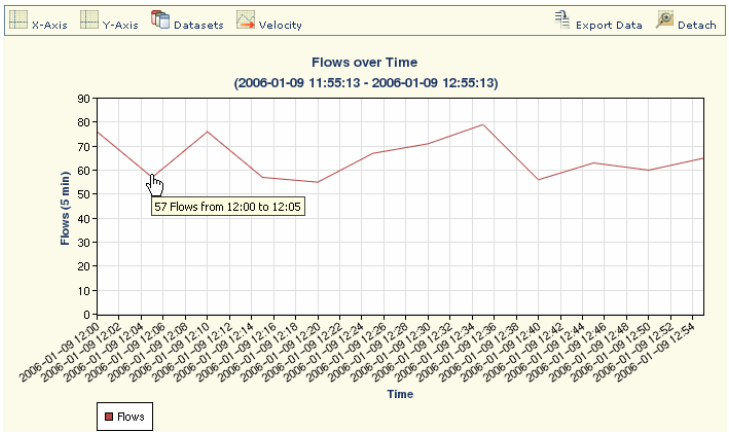
- Understanding Graph Types on page 1079
- Understanding Datasets on page 1082
- Understanding Flow Data Graph Functions on page 1084

Understanding Graph Types

Platform: DC

Component: RNA

There are three different flow data 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 flows detected on a monitored network segment over a one-hour time span. Traffic profiles are always displayed as line graphs.

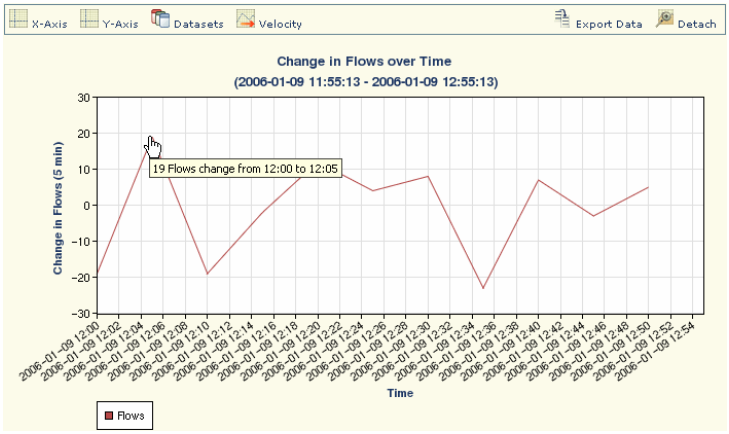


By default, line graphs appear in *standard view*. A standard line graph aggregates data over five minute intervals, plots the aggregated data points, and connects the points. For example, in the standard graph above, the first data point shows 76 flows at the starting point. The second data point shows 57 flows at five

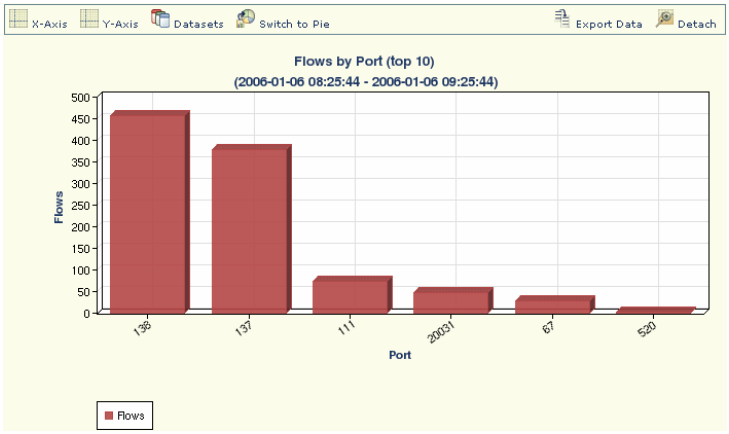
Keysight, EX-1004
Page 1079

minutes, which means that 57 flows were detected between minute zero and minute five.

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 flows to indicating the change in the number of flows over time. In the example below, the velocity graph shows that the number of flows increased by 19.

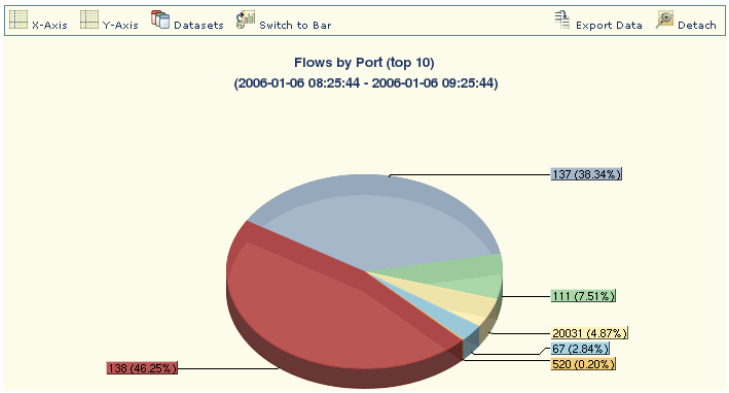


Bar graphs display data grouped into discrete categories. For example, a bar graph could show the number of flows detected on a monitored network segment for the 10 most active ports over a one-hour time span.



Keysight, EX-1004
Page 1080

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.



Access: Any RNA/P&R Admin/Admin

Follow the directions in the Changing Graph Types table to switch between standard and velocity line graphs, and to switch between bar graphs and pie charts. Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

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. For more information, see Understanding Datasets on page 1082.
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**.

Keysight, EX-1004
Page 1081

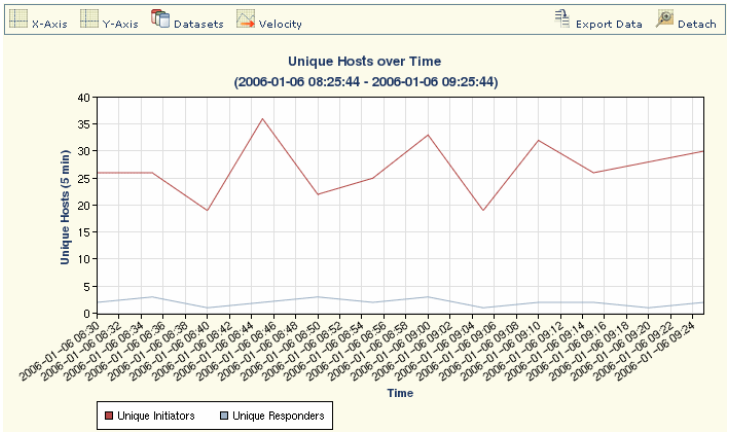
Understanding Datasets

Platform: DC

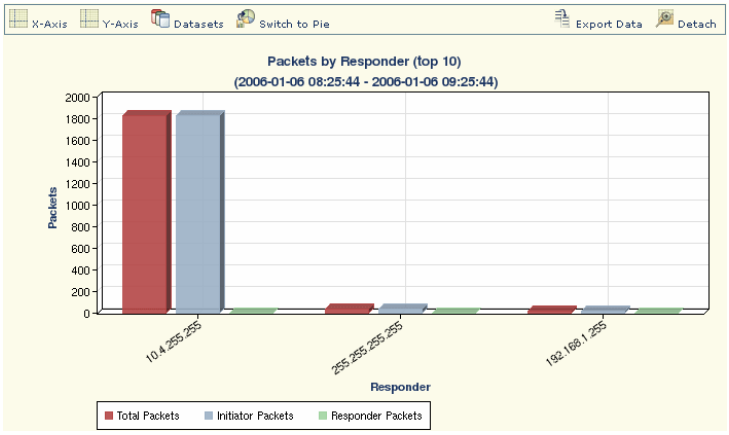
Component: RNA

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 responders detected on a monitored network segment over a one-hour interval.

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 segment over a one hour interval.

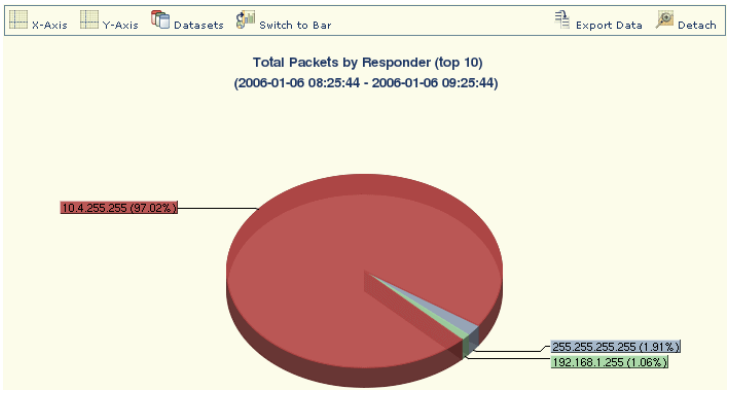


On bar graphs, multiple datasets appear as a set of colored bars for each x-axis data point. For example, the following bar graph displays the total packets transmitted on a monitored network segment, packets transmitted by initiators, and packets transmitted by responders.

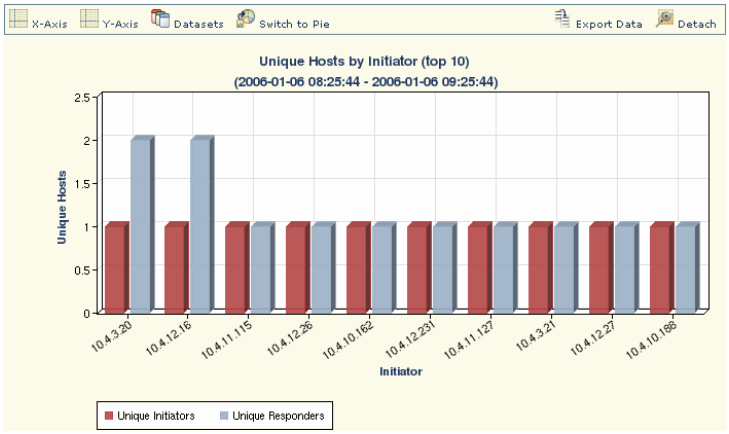


Keysight, EX-1004
Page 1082

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. For example, if you switch the preceding bar graph to a pie chart, the chart displays only the total number of packets by responder.

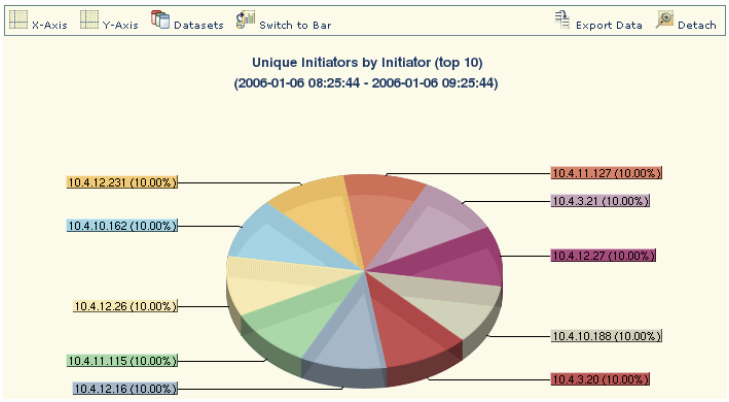


Contrast this with the case where you view a bar graph that displays the number of unique initiating hosts and the number of unique responding hosts on a monitored network segment.



Keysight, EX-1004
Page 1083

In this case, if you switch to a pie chart, the chart displays only the unique initiating hosts.



For information on how to display multiple datasets on a graph, see Selecting Datasets on page 1113.

Understanding Flow Data Graph Functions

Platform: DC

Component: RNA

The Flow Data Graph Functions table describes the different actions you can perform on a flow data workflow page that contains a graph.

Access: Any RNA/P&R Admin/Admin

Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

Flow Data Graph Functions

To...	You can...
modify the time and date range for the flow data graph	find more information in Setting Event Time Constraints on page 1597.
view a host's profile	on a graph displaying flow 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**.
learn more about the different kinds of graphs	see Understanding Graph Types on page 1079.

Keysight, EX-1004
Page 1084

Flow Data Graph Functions (Continued)

To...	You can...
change between bar graphs and pie chart, and between standard line graphs and velocity graphs	find more information in Changing the Graph Type on page 1109.
change the kinds of data displayed on the graph	find more information in Selecting Data to Graph on page 1109 and Selecting Datasets on page 1113.
get more information about the aggregated flow data used to construct individual data points	see Viewing Aggregated Flow Data on page 1101.
manipulate and constrain a flow data graph	find more information in Manipulating a Flow Data Graph on a Workflow Page on page 1103.
navigate between pages in the current workflow	find more information in Using Workflow Pages on page 1593.
constrain flow data by drilling down through the workflow	find more information in Drilling Down through a Flow Data Workflow on page 1105.
detach a flow graph to a pop-up window so that you can further manipulate it	find more information in Detaching Flow Data Graphs on page 1114.
export the aggregated data used in a flow data graph as a comma-separated values (CSV) file	find more information in Exporting Flow Data on page 1115.
navigate to other event views to view associated events	find more information in Navigating between Workflows on page 1614.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.

Keysight, EX-1004
Page 1085

Flow Data Graph Functions (Continued)

To...	You can...
generate a report based on the data in the flow data view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for flow data	click **Search**. For more information, see Searching for Flow Data on page 1094.

Understanding Flow Data Tables

Platform: DC

Component: RNA

As with all other event types, the Defense Center can present flow data in tables.

Note that there are several differences between the data collected and exported by NetFlow-enabled devices and the flow data detected directly by the detection engines on 3D Sensors with RNA. It is important to keep these differences in mind when performing any kind of analysis that includes flow data, or that uses information derived from flow data. For more information, see How is NetFlow Data Different from RNA Flow Data? on page 1142.

Access: Any RNA/Admin

The Flow Data Table Functions table describes some of the actions you can perform on a flow data workflow page that contains a table.

Flow Data Table Functions

To...	You can...
learn more about the contents of the columns in the table	depending on whether you are viewing individual flows or flow summaries, see either: • Flow Data Fields on page 1089 • Flow Summary Fields on page 1092
modify the time and date range for the flow data graph	click the time range link. For more information, see Setting Event Time Constraints on page 1597. 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.

Keysight, EX-1004
Page 1086

Flow Data Table Functions (Continued)

To...	You can...
view a host's host profile	click the host profile icon ( ) that appears next to the IP address. For more information, see Using Host Profiles on page 944.
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 1495.
sort flow data	click the column title. Click the column title again to reverse the sort order.
constrain the columns that appear	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, enable or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, use the Expand arrow ( ) to expand the search constraints, then click the column name under Disabled Columns. **IMPORTANT!** When you constrain flow events in a table, the packets and bytes from identical events are summed. However, if you are using a custom workflow and did not add a Count column, the events are listed individually and packets and bytes are **not** summed.
drill down to the next page in the workflow	use one of the following methods: • 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.

Keysight, EX-1004
Page 1087

Flow Data Table Functions (Continued)

To...	You can...
delete flow data from the system	use one of the following methods: • To delete some flow events or summaries, select the check boxes next to the events or summaries you want to delete, then click **Delete**. • To delete all flow events or summaries in the current constrained view, click **Delete All**, then confirm you want to delete all the events of summaries.
navigate within and between workflow pages	find more information in Using Workflow Pages on page 1593.
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 1614.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the flow data view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for flow data	click **Search**. For more information, see Searching for Flow Data on page 1094.

Keysight, EX-1004
Page 1088

Flow Data Fields

For any drill-down pages based on individual flows that you created in a custom workflow, and for the table view of flow data, the fields that can appear are described in the following table.

IMPORTANT! To view data on individual flows, you must set the flow data mode to **Enabled** in your RNA detection policy. For more information, see What is an RNA Detection Policy? on page 889.

Flow Data Fields

Field	Description
First Packet	The date and time the first packet of the session was seen.
Last Packet	The date and time the last packet of the session was seen.
Initiator IP	The IP address (and host name, if DNS resolution is enabled) of the host that initiated the session.
Responder IP	The IP address (and host name, if DNS resolution is enabled) of the host that responded to the session initiator.
Initiator User	The ID of a user logged to the host that initiated the session.
Responder User	The ID of a user logged in to the host that responded to the session initiator.
Protocol	The protocol used during the detected session.
Initiator Port	The port used by the host that initiated the session.
Responder Port	The port used by the host that responded to the session initiator.
TCP Flags	The TCP flags detected in the flow.
Flow Type	Whether the flow was detected by a 3D Sensor with RNA, or was exported by a NetFlow-enabled device.

Keysight, EX-1004
Page 1089

Flow Data Fields (Continued)

Field	Description
Service	One of: • the name of the service used in the flow • `pending`, if RNA cannot identify the service for one of several reasons For detailed information on how the Defense Center determines and displays service identities, see the Sourcefire 3D System Identification of Services table on page 1387.
NetBIOS Domain	The name of the NetBIOS domain used in the flow transaction.
Application Type	The type of client application used in the flow, if applicable.
Application	The name of the application used by the client in the flow, if applicable.
Version	The version of the application used by the client in the flow, if applicable.
Payload Type	The general type of HTTP content (such as audio, video, or webmail) in the flow, if applicable. For a list of detected payload types, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.
Payload	If the service is http, the specific type of content detected by RNA. For example, if the payload type is video, the payload itself could be WMV, QuickTime, MP4, and so on. For a list of detected payloads, see Detected Payloads on page 2039. Keep in mind, however, that the list can change depending on the version of the Sourcefire 3D System and the VDB that you have installed. 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 payload detectors.

Keysight, EX-1004
Page 1090

Flow Data Fields (Continued)

Field	Description
Initiator Pkts	The total number of packets transmitted by the host that initiated the session.
Responder Pkts	The total number of packets transmitted by the host that responded to the session initiator.
Initiator Bytes	The total number of bytes transmitted by the host that initiated the session.
Responder Bytes	The total number of bytes transmitted by the host that responded to the session initiator.
URL	If the service is http and the **Capture HTTP URLs** option is enabled in your RNA detection policy, this column displays the URL requested by the monitored host during the session.
Source Device	The IP address of the NetFlow-enabled device that exported the flow data for the flow. If the flow was detected by a 3D Sensor with RNA, this field contains a value of RNA.
Detection Engine	The detection engine that collected the flow event, or, for flows exported by NetFlow-enabled devices, the detection engine that processed the NetFlow data.
Count	The number of flows 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 flow is listed individually and packets and bytes are not summed.

Flow Summary Fields

Platform: DC

Component: RNA

For any drill-down pages based on flow summary data that you created in a custom workflow, and for the table view of flow summary data, the fields that can appear are described in the following table.

Flow summaries are based on flow data aggregated over five-minute intervals. If a session spans two or more intervals, the flow is considered a *long-running flow*.

Keysight, EX-1004
Page 1091

When calculating the number of flows in a flow summary, RNA increments the count only for the five-minute interval in which a long-running flow was initiated.

Also, when calculating the number of packets and bytes transmitted by the initiator and responder in a long-running flow, RNA does not report the number of packets and bytes that were actually transmitted during each five-minute interval. Instead, RNA assumes a constant rate of transmission and calculates estimated figures based on the total number of packets and bytes transmitted, the length of the flow, and what portion of the flow occurred during each five minute interval.

Note that because NetFlow-enabled devices generate unidirectional flows, a summary's flow count is incremented by two for every flow based on NetFlow data. For more information on flow summary data see Understanding Flow Summary Data on page 1076.

Flow Summary Fields

Field	Description
Time	The ending time of the five-minute interval that RNA used to aggregate the flows in the summary.
Initiator IP	The IP address (and host name, if DNS resolution is enabled) of the host that initiated the sessions aggregated in the summary.
Responder IP	The IP address (and host name, if DNS resolution is enabled) of the host that responded to the session initiator.
Initiator User	The ID of the user logged in to the host that initiated the sessions aggregated in the summary.
Responder User	The ID of the user logged in to the host that responded to the session initiator.
Protocol	The protocol used during the detected sessions aggregated in the summary.
Responder Port	The port used by the host that responded to the session initiator.
Flow Type	Whether the flows aggregated in the summary were detected by a 3D Sensor with RNA, or were exported by a NetFlow-enabled device.

Keysight, EX-1004
Page 1092

Flow Summary Fields (Continued)

Field	Description
Service	One of: • the name of the service used in the flow summary • `pending`, if RNA cannot identify the service for one of several reasons For detailed information on how the Defense Center determines and displays service identities, see the Sourcefire 3D System Identification of Services table on page 1387.
Initiator Pkts	The total number of packets transmitted by the host that initiated the sessions aggregated in the summary.
Responder Pkts	The total number of packets transmitted by the host that responded to the session initiator.
Initiator Bytes	The total number of bytes transmitted by the host that initiated the sessions aggregated in the summary.
Responder Bytes	The total number of bytes transmitted by the host that responded to the session initiator.
Source Device	The IP address of the NetFlow-enabled device that exported the flow data for the flow. If the flow was detected by a 3D Sensor with RNA, this field contains a value of `RNA`.
Detection Engine	The detection engine that collected the flow event, or, for flows exported by NetFlow-enabled devices, the detection engine that processed the NetFlow data.
Flows	The number of flows contained in the flow summary you are viewing. For long running flows, that is, flows that span multiple flow summary intervals, only the first flow summary interval is incremented.

Viewing Flow Data

Platform: DC

Component: RNA

The page you see when you access flow data differs depending on the workflow you use. You can use one of the predefined workflows or create a custom

Keysight, EX-1004
Page 1093

workflow that displays only the information that matches your specific needs. For information on creating a custom flow data workflow, see Creating Custom Flow Data Workflows on page 1621.

To view flow data:

Access: Any RNA/Admin

▶ Select **Analysis & Reporting > RNA > Flow Data**.

The first page of the default flow data workflow appears. There are two possibilities:

- The workflow page displays a **graph**. See Understanding Flow Data Graphs on page 1078 for information on the actions you can perform.
- The workflow page displays a **table**. See Understanding Flow Data Tables on page 1086 for information on the actions you can perform.

To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

Searching for Flow Data

Platform: DC

Component: RNA

You can search for specific flows by using one of the predefined searches delivered with the Defense Center or by using your own search criteria. The predefined searches serve as examples and can provide quick access to important information about your network. Default searches are:

- Possible Database Access, which searches for TCP traffic flows to ports that are commonly used by database servers.
- Standard HTTP, which searches for TCP traffic flows to standard web server ports.
- Standard Mail, which searches for TCP traffic flows to standard mail server ports.

Keysight, EX-1004
Page 1094

- Standard SSL, which searches for TCP traffic flows to the standard SSL port (443).
- Unauthorized SMTP, which searches for traffic to the standard SMTP port (25) on hosts that are not authorized SMTP servers from hosts inside the network.

  This search **must** be customized for your network before you can run it. Configure the Initiator IP with the address you want to use for the internal network (for example, `172.16.1.1/24`) and the Responder IP as the negated values of your SMTP servers (for example, `!172.16.1.28,!172.16.1.27`).

You may want to modify specific fields within the default searches to customize them for your network environment, then save them to re-use later. The search criteria you can use are described in the following table.

Flow Data Search Criteria

Field	Search Criteria Rules
First Packet or Last Packet	Type either: • the date and time that the first packet of the session was seen (**First Packet**) • the date and time that the last packet of the session was seen (**Last Packet**) See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time.
Initiator IP, Responder IP, or Initiator/Responder IP	Type the IP address (or host name, if DNS resolution is enabled) of either: • the host that initiated the session (**Initiator IP**) • the host that responded to the session initiator (**Responder IP**) • either host involved in the session (**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 1557 for a full description of the syntax allowed for IP addresses.
Initiator User, Responder User, or Initiator/Responder User	Type the User ID of either: • the user of the host that initiated the session (**Initiator User**) • the user of the host that responded to the session initiator (**Responder User**) • the user of either host involved in the session (**Initiator/Responder User**) Use the full User ID or a partial User ID during searches. Note that partial name recognition is permitted (for example, smith yields jsmith, msmith for John Smith and Mary Smith).
Protocol	Type the name or number of the transport protocol used in the flow as listed in http://www.iana.org/assignments/protocol-numbers.

Keysight, EX-1004
Page 1095

Flow Data Search Criteria (Continued)

Field	Search Criteria Rules
Initiator Port or Responder Port	Type either: • the port used by the host that initiated the session (**Initiator Port**) • the port used by the host that responded to the session initiator (**Responder Port**) See Specifying Ports in Searches on page 1558 for information about specifying ports.
TCP Flags	Type a list of comma-separated TCP flags to view all flows that have at least one of those flags. You can also select the **Only** checkbox to search for flows that have any of the flags you specify as their only TCP flag.
Flow Type	Type either: • `RNA` to search for flows detected by 3D Sensors with RNA • `NetFlow` to search for flows exported by a NetFlow-enabled device
Service	Type the name of the service used in the flow. For a list of services detected by RNA, see Detected Services on page 2027. **TIP!** Type `pending` to search for services that RNA cannot identify. For detailed information on how the Defense Center determines and displays service identities, see the Sourcefire 3D System Identification of Services table on page 1387.
NetBIOS Domain	Specify the name of the NetBIOS domain used in the flow transaction. You can use an asterisk (*) as a wildcard character in this field.
Application Type	Specify the type of client application used in the flow.
Application	Specify the name of the client application used by the host involved in the flow. For a list of applications detected by RNA, see Detected Client Applications on page 2036. You can use an asterisk (*) as a wildcard character in this field.
Version	Specify the version of the client application used by the host involved the flow. You can use an asterisk (*) as a wildcard character in this field.
Payload Type	Specify the payload type, that is, the general type of HTTP content (such as audio, video, or webmail) in the flows you want to view. For a list of payload types that you can specify, see Detected Payloads on page 2039.
Payload	Specify the payload, that is, the specific type of HTTP content (such as MPEG, WMV, or Gmail) in the flows you want to view. For a list of payloads that you can specify, see Detected Payloads on page 2039.

Keysight, EX-1004
Page 1096

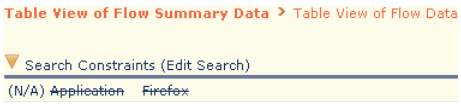
Flow Data Search Criteria (Continued)

Field	Search Criteria Rules
Initiator Pkts or Responder Pkts	Type either: • the number of packets transmitted by the hosts on the monitored network segment (**Initiator Pkts**) • the number of packets received by the hosts on the monitored network segment (**Responder Pkts**) You can precede the number with greater than (>), greater than or equal to (>=), less than (<), less than or equal to (<=), or equal to (=).
Initiator Bytes or Responder Bytes	Type either: • the number of bytes transmitted by the hosts on the monitored network segment (**Initiator Bytes**) • the number of bytes received by the hosts on the monitored network segment (**Responder Bytes**) You can precede the number with greater than (>), greater than or equal to (>=), less than (<), less than or equal to (<=), or equal to (=).
URL	Specify all or part of a captured URL. This field is case-insensitive. You can use an asterisk (*) as a wildcard character in this field. Specify whether RNA captures HTTP URLs in flow events when you configure the RNA detection policy. For more information, see What is an RNA Detection Policy? on page 889.
Source Device	Type the IP address of a NetFlow-enabled device that exported the flows you want to view. To search for flows detected by 3D Sensors with RNA, type **RNA**. See Specifying IP Addresses in Searches on page 1557 for a full description of the syntax allowed for IP addresses.
Detection Engine	Specify the name of the detection engine, detection engine group, sensor, or sensor group that generated the flow event you want to view. For more information, see Specifying Detection Engine/Appliance Names on page 1559.
Flows	Type the number of flows contained in a flow summary for the flow summaries you want to view. You can precede the number with greater than (>), greater than or equal to (>=), less than (<), less than or equal to (<=), or equal to (=). **IMPORTANT!** To view meaningful results for searches using this field, you must use a workflow that includes flow summary data. For more information on flow summaries, see Understanding Flow Summary Data on page 1076.

Keysight, EX-1004
Page 1097

When searching for flows, you should keep in mind that flows detected by 3D Sensors with RNA and flows exported by NetFlow-enabled devices contain different information. For example, flows detected by 3D Sensors with RNA do not contain TCP flag information. For more information, see How is NetFlow Data Different from RNA Flow Data? on page 1142.

Also, because you use the same page in the Defense Center web interface to search both flow summary data and individual flows, you must pay careful attention to whether your search constraints apply to the table you are searching. For example, client application data is not available in flow summaries. If you search for flow events based on the detected client application in the flow and then view the results in a flow summary event view, the Defense Center displays flow data as if you had not constrained it at all. Search constraints that do not apply to the current event view are labeled as not applicable (N/A) and are marked with a strikethrough, as shown in the following graphic.



For more information on flow summaries, see Understanding Flow Summary Data on page 1076. For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

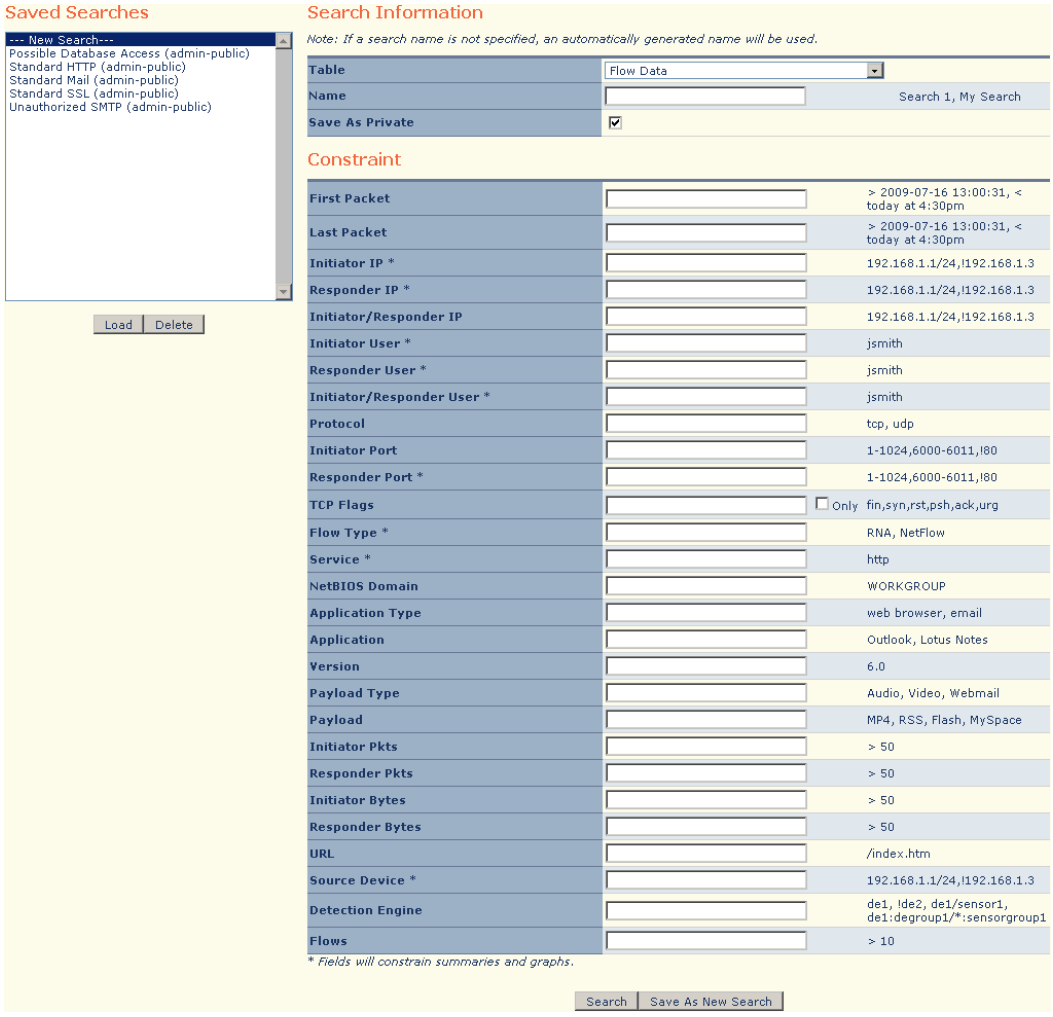
Keysight, EX-1004
Page 1098

To search for flow data:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Searches > Flow Data**.

   The Flow Data search page appears.



TIP! To search the database for a different kind of event, select it from the **Table** drop-down list.

Keysight, EX-1004
Page 1099

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 Flow Data Search Criteria table. If you enter multiple criteria, the Defense Center returns only the records that match all the criteria.

   Fields marked with an asterisk (*) constrain flow data graphs and flow summaries, as well as individual flows.

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 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 flow data workflow, constrained by the current time range. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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.

Manipulating Flow Data Graphs

Platform: DC

Component: RNA

You can manipulate flow data 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, services, ports, and detection engines

Because traffic profiles are based on flow data (see Creating Traffic Profiles on page 1118), you can view traffic profiles as line graphs. You can manipulate these

Keysight, EX-1004
Page 1100

graphs in the same way as you would any other flow graph, with some restrictions. Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

For more information, see:

- Viewing Aggregated Flow Data on page 1101 explains how to get more information about the data points on the graph, or to display the host profile of a host whose statistics are being graphed.
- Manipulating a Flow Data Graph on a Workflow Page on page 1103 explains how to constrain the data that appears on a flow data graph without advancing the workflow to the next page.
- Drilling Down through a Flow Data Workflow on page 1105 explains how to constrain the data that appears on a flow data graph while advancing the workflow to the next page.
- Recentering and Zooming on Line Graphs on page 1108 explains how to recenter a line graph around any point in time.
- Changing the Graph Type on page 1109 explains how to change between bar graphs and pie chart, and between standard line graphs and velocity graphs.
- Selecting Data to Graph on page 1109 explains how to change the data displayed on a flow data graph by changing its x- or y-axis.
- Selecting Datasets on page 1113 explains how to display several values on the y-axis for each x-axis data point on line graphs and bar graphs.
- Detaching Flow Data Graphs on page 1114 explains how you can detach a flow data graph into a new browser window and perform further analysis without affecting the default time range for the Defense Center.
- Exporting Flow Data on page 1115 explains how to export the flow data used to construct a graph as a CSV (comma-separated values) file.

Viewing Aggregated Flow Data

Platform: DC

Component: RNA

Flow data graphs are based on aggregated data over five-minute intervals, also called *flow summaries*. You can get more information about the specific flow summaries used to construct a flow data graph. For example, on a graph of flows over time, you may want to know exactly how many flows were detected over a specific interval.

Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs,

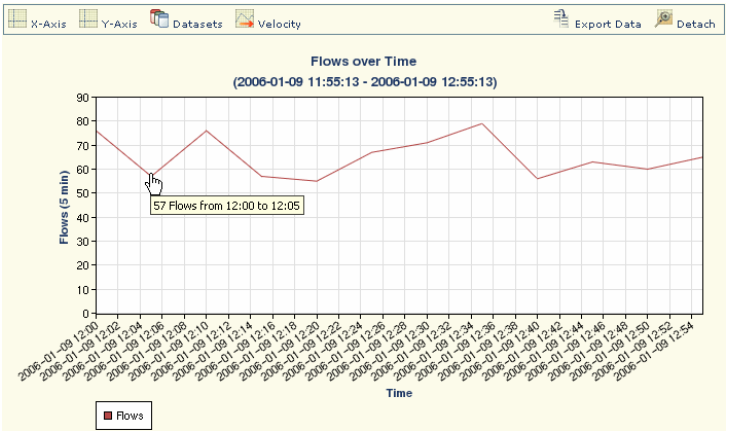
Keysight, EX-1004
Page 1101

which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

To get detailed information on aggregated flow data:

Access: Any RNA/P&R Admin/Admin

- Position your cursor over a point on a line graph (as shown in the following graphic), a bar in a bar graph, or a wedge in a pie chart. A tool tip appears with detailed information about the data used to construct that portion of the graph.



TIP! To see all the aggregated data used to build the flow data graph, view the table view of flow summary data; if you are using one of the workflows delivered with the Defense Center, click **Table View of Flow Summary Data** at the top of the page. You can also export the graph as a CSV file, as described in Exporting Flow Data on page 1115.

Keysight, EX-1004
Page 1102

Manipulating a Flow Data Graph on a Workflow Page

Platform: DC

Component: RNA

When you open a flow data workflow, the data is initially constrained only by a time range. You can constrain flow data graphs with additional criteria without advancing the workflow to the next page.

TIP! Constraining flow 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 flow data, use the **X-Axis** and **Y-Axis** menus. For more information, see Selecting Data to Graph on page 1109.

Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

To constrain flow data:

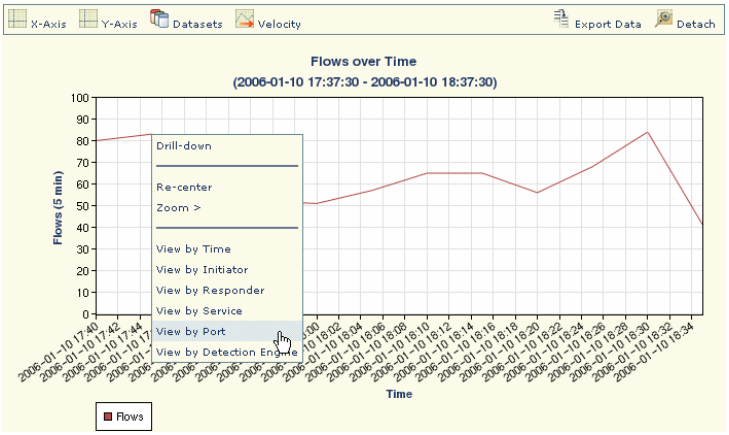
Access: Any RNA/P&R Admin/Admin

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.

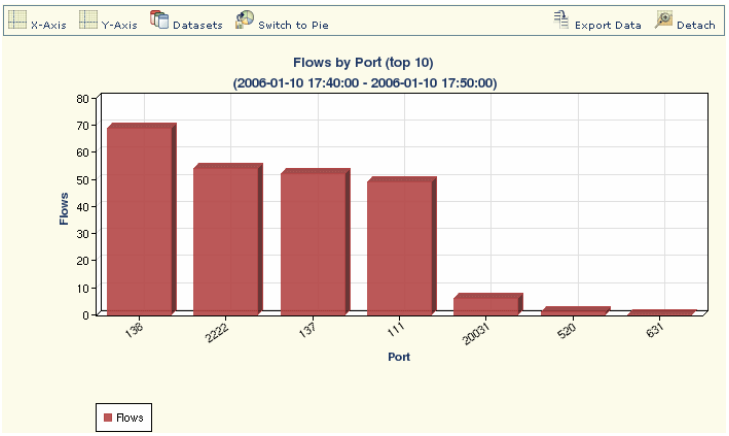
   You can constrain flow data based any of the criteria listed in the X-Axis Functions table on page 1111.

Keysight, EX-1004
Page 1103

For example, consider a graph of flows over time. If you constrain a point on the graph by port:

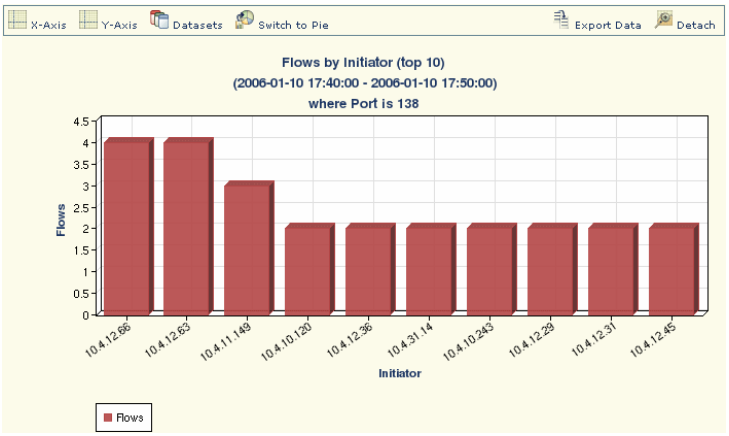


A bar graph appears, showing the 10 most active ports based on the number of detected flow events, but constrained by the ten-minute time span that is centered on the point you clicked.



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.

Keysight, EX-1004
Page 1104



IMPORTANT! Note that unless you are working with a detached graph, constraining flow data in this manner changes the time range. For more information on detached graphs, see Detaching Flow Data Graphs on page 1114.

Drilling Down through a Flow Data Workflow

Platform: DC

Component: RNA

When you open a flow data workflow, the data is initially constrained only by a time range. You can constrain flow data graphs while advancing the workflow to the next page.

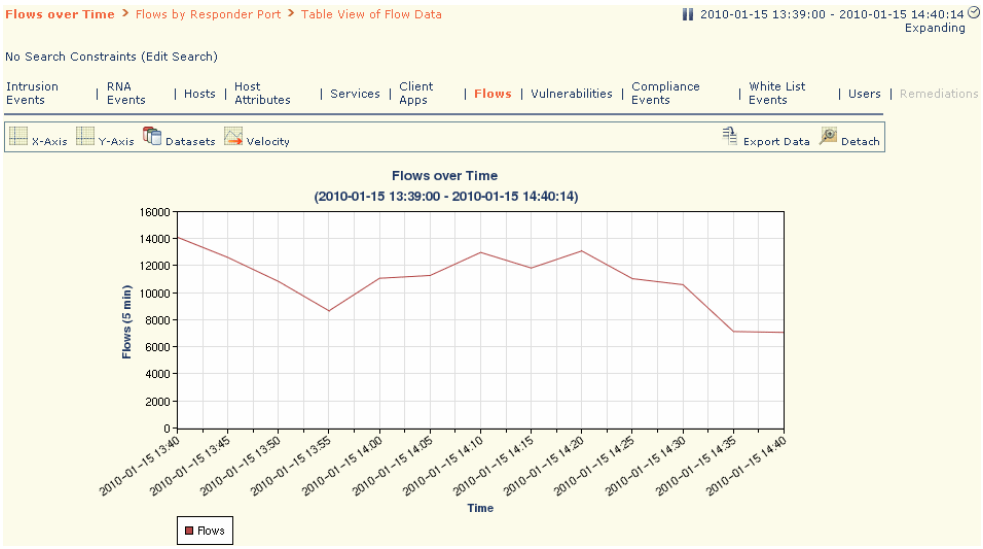
Keysight, EX-1004
Page 1105

To drill down in a flow data workflow:

Access: Any RNA/Admin

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**.

   For example, consider the following custom workflow.

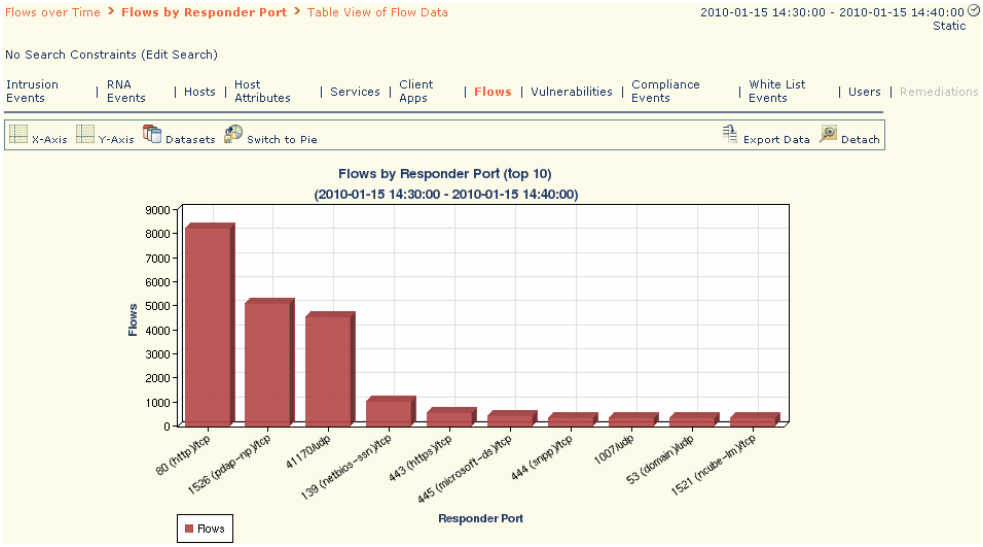


This workflow has three pages:

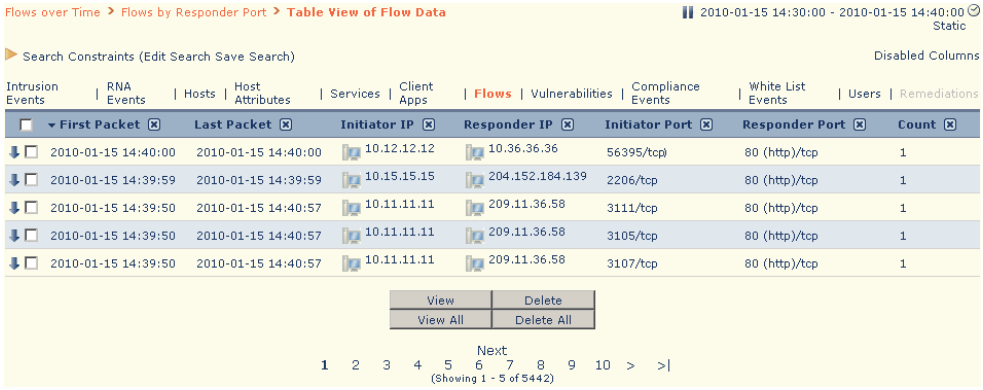
- **Flows over Time**, which is a line graph of the total number of flows over time
- **Flows by Responder Port**, which is a bar graph of the 10 most active responder ports based on the number of detected flow events
- the table view of flow data

If you drill down from the first page to the second page of the workflow, the time span for the **Flows by Responder Port** bar graph is now a 10 minute range, centered on the point you clicked.

Keysight, EX-1004
Page 1106



If you drill down again, the workflow advances to the table view. The time span is still constrained to the same 10 minute range, and the table view shows only those flow events on the port represented by the bar you clicked.



IMPORTANT! If you used the RNA detection policy to configure your 3D Sensors to transmit flow data to the Defense Center as flow summaries only, you **cannot** examine the individual flow events on which flow data graphs are based. In other words, you cannot drill down to the flow data table view from a graph based on flow summaries. For more information, see Understanding Flow Summary Data on page 1076.

Keysight, EX-1004
Page 1107

Recentering and Zooming on Line Graphs

Platform: DC

Component: RNA

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 Flow Data Graphs on page 1114.

Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

To recenter using the default time range:

Access: Any RNA/P&R Admin/Admin

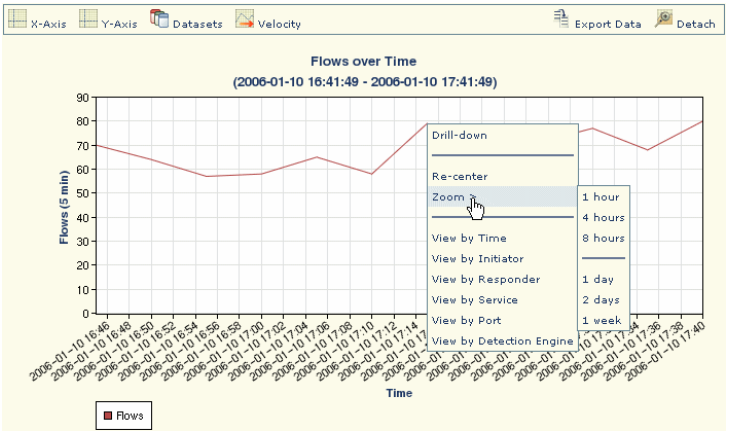
- Click the point on the line graph where you want to recenter the graph, and click **recenter**.

  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: Any RNA/P&R Admin/Admin

1. Click the point where you want to recenter the graph and click **Zoom**.



Keysight, EX-1004
Page 1108

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.

Changing the Graph Type

Platform: DC

Component: RNA

You can switch between standard and velocity line graphs, and between bar graphs and pie charts. For more information, see Understanding Graph Types on page 1079.

Follow the directions in the Changing Graph Types table to change the graph type. Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

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. For more information, see Understanding Datasets on page 1082.
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 Data to Graph

Platform: DC

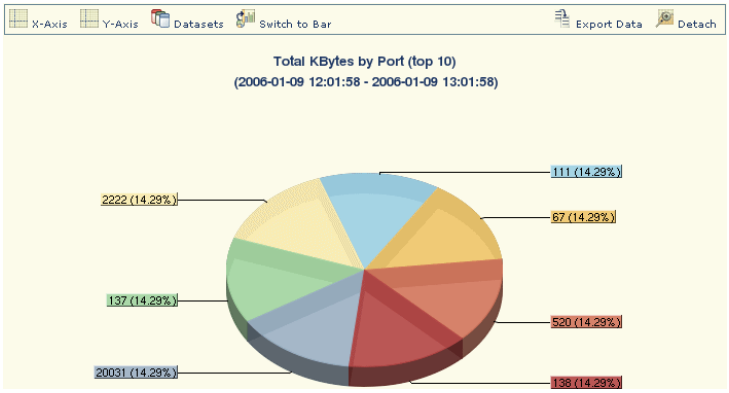
Component: RNA

You can display different data on a flow data graph by changing either the x-axis, the y-axis, or both.

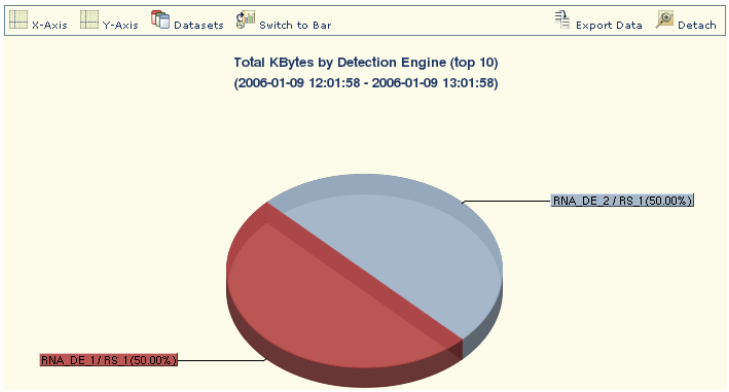
Note that on a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable. For example, consider a pie

Keysight, EX-1004
Page 1109

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 segment during a certain interval. The wedges of the pie represent the percent of the data that was detected on each port.



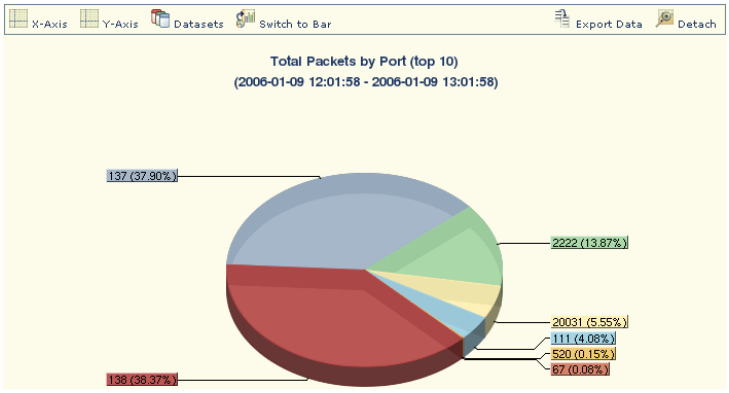
If you change the x-axis of the chart to **Detection Engine**, the pie chart still represents the total kilobytes of data transmitted, but the wedges of the pie represent the percentage of the data that was detected by each detection engine.



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

Keysight, EX-1004
Page 1110

segment 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 flow data graph. Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

X-Axis Functions

To graph flow data...	You can...
by the 10 most active detection engines on the monitored network segment based on the number of detected flow events	click **X-Axis** and select **Detection Engine**.
by whether the flows were detected by a 3D Sensor with RNA, or were exported by a NetFlow-enabled device	click **X-Axis** and select **Flow Type**.
by the 10 most active hosts on the monitored network segment based on the number of flow events where the host initiated the flow transaction	click **X-Axis** and select **Initiator IP**.
by the 10 most active users on the monitored network segment based on the number of flow events where the host where the user is logged in initiated the flow transaction	click **X-Axis** and select **Initiator User**.
by the 10 most active hosts on the monitored network segment based on the number of flow events where the host was the responder in the flow transaction	click **X-Axis** and select **Responder IP**.
by the 10 most active ports on the monitored network segment based on the number of detected flow events where the host was the responder in the flow transaction	click **X-Axis** and select **Responder Port**.

Keysight, EX-1004
Page 1111

X-Axis Functions (Continued)

To graph flow data...	You can...
by the 10 most active users on the monitored network segment based on the number of flow events where the host where the user is logged in was the responder in the flow transaction	click **X-Axis** and select **Responder User**.
by the 10 most active services on the monitored network segment based on the number of detected flow events	click **X-Axis** and select **Service**.
by the 10 most active source devices, which include NetFlow-enabled devices that exported the flow data for the flows, plus a source device named "RNA" for all flows detected by 3D Sensors with RNA	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 flow data graph. Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

Y-Axis Functions

To...	You can...
graph the number of flows on the monitored network segment by the criterion you chose for the x-axis	click **Y-Axis** and select **Flows**.
graph the total kilobytes transmitted on the monitored network segment 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 segment 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 segment 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 segment by the criterion you chose for the x-axis	click **Y-Axis** and select **Unique Hosts**.
graph the total number of unique services detected on the monitored network segment by the criterion you chose for the x-axis	click **Y-Axis** and select **Unique Services**.
graph the total number of unique users detected on the monitored network segment by the criterion you chose for the x-axis	click **Y-Axis** and select **Unique Users**.

Keysight, EX-1004
Page 1112

Selecting Datasets

Platform: DC

Component: RNA

On line and bar graphs you can display several values on the y-axis for each x-axis data point. Pie graphs can only display one dataset. For more information, see Understanding Datasets on page 1082.

The Dataset Options table describes the datasets you can display on the x-axis of a flow data graph.

Dataset Options

If the y-axis displays...	You can select as datasets...
Flows	the default only, which is the number of flows detected on the monitored network segment (**Flows**) This is the only option for traffic profile graphs.
KBytes	combinations of: • the total kilobytes transmitted on the monitored network segment (**Total KBytes**) • the number of kilobytes transmitted by the hosts on the monitored network segment (**Initiator KBytes**) • the number of kilobytes received by the hosts on the monitored network segment (**Responder KBytes**)
KBytes Per Second	the default only, which is the total kilobytes per second transmitted on the monitored network segment (**Total KBytes Per Second**)
Packets	combinations of: • the total packets transmitted on the monitored network segment (**Total Packets**) • the number of packets transmitted by the hosts on the monitored network segment (**Initiator Packets**) • the number of packets received by the hosts on the monitored network segment (**Responder Packets**)

Keysight, EX-1004
Page 1113

Dataset Options (Continued)

If the y-axis displays...	You can select as datasets...
Unique Hosts	combinations of: • the number of unique hosts that initiated sessions that were detected on the monitored network segment (**Unique Initiators**) • the number of unique hosts that responded to flow transaction that were detected on the monitored network segment (**Unique Responders**)
Unique Services	the default only, which is the number of unique services detected on the monitored network segment (**Unique Services**)
Unique Users	combinations of: • the number of unique users logged into hosts that initiated sessions that were detected on the monitored network segment (**Unique Initiator Users**) • the number of unique users logged into hosts that responded to flow transaction that were detected on the monitored network segment (**Unique Responder Users**)

To select the datasets displayed on a flow data graph:

Access: Any RNA/Admin

▶ Click **Datasets** and select the datasets you want to graph.

The datasets you can select are described in the Dataset Options table.

Detaching Flow Data Graphs

Platform: DC

Component: RNA

If you want to perform further analysis on a flow data 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 flow data graphs that you can on embedded flow data 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.

Keysight, EX-1004
Page 1114

To detach a graph:

Access: Any RNA/Admin

▶ Click **Detach**.

Exporting Flow Data

Platform: DC

Component: RNA

You can easily share flow data with others by exporting it as a CSV (comma-separated values) file. Note that to view traffic profile graphs you must have either Policy & Response Administrator or Administrator access. Compare this with other flow graphs, which you can view with RNA Analyst, RNA Analyst (Read-Only), or Administrator access.

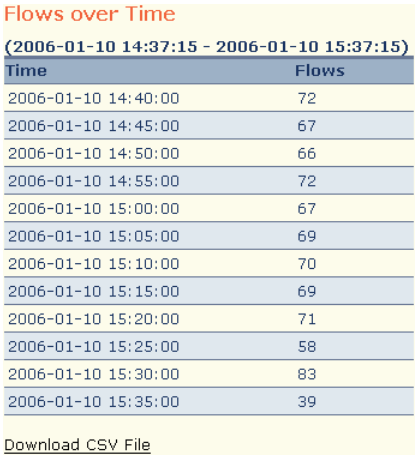
TIP! You can also save a flow data graph as an image by right-clicking on the graph and selecting **Save Image As** (Firefox and Netscape) or **Save Picture As** (Internet Explorer).

To export flow data:

Access: Any RNA/P&R Admin/Admin

1. Click **Export Data**.

   A pop-up window appears, displaying a table view of the data on your graph.



Flows over Time

(2006-01-10 14:37:15 - 2006-01-10 15:37:15)

Time	Flows
2006-01-10 14:40:00	72
2006-01-10 14:45:00	67
2006-01-10 14:50:00	66
2006-01-10 14:55:00	72
2006-01-10 15:00:00	67
2006-01-10 15:05:00	69
2006-01-10 15:10:00	70
2006-01-10 15:15:00	69
2006-01-10 15:20:00	71
2006-01-10 15:25:00	58
2006-01-10 15:30:00	83
2006-01-10 15:35:00	39

Download CSV File

2. Click **Download CSV File** and save the file.

Keysight, EX-1004
Page 1115

Viewing the Flow Summary Page

Platform: DC

Component: RNA

The Flow Summary page provides graphs of the activity on your monitored network segment organized by different criteria.

IMPORTANT! The Flow Summary page is visible only to Restricted Event Analysts. Most other users see the Flow Summary dashboard; see Working with Dashboards on page 92. (External Database users see neither dashboards nor the Flow Summary page.) Note that if you create a Restricted Event Analyst and hide flow data properties, no data is visible; see Modifying Restricted Event Analyst Access Properties on page 1685.

Contrast the Flow Summary page with flow summaries, which are comprised of flow data aggregated over five-minute intervals. RNA uses flow summaries to create the graphs on the Flow Summary page. For more information on flow summaries, see Understanding Flow Summary Data on page 1076.

The following table describes the information displayed on the Flow Summary page.

Flow Summary Graphs

Title	Description
Flows over Time	Displays a graph of the total number of flows on the monitored network segment over the interval that you select.
Flows by Initiator	Displays a graph of the 10 most active hosts on the monitored network segment based on the number of flow events where the host initiated the flow transaction.
Flows by Responder	Displays a graph of the 10 most active hosts on the monitored network segment based on the number of flow events where the host was the responder in the flow transaction.
Flows by Port	Displays a graph of the 10 most active ports on the monitored network segment based on the number of detected flow events.
Flows by Service	Displays a graph of the 10 most active services on the monitored network segment based on the number of detected flow events.
Traffic over Time	Displays a graph of the total kilobytes transmitted on the monitored network segment over the interval that you select.

Keysight, EX-1004
Page 1116

Flow Summary Graphs (Continued)

Title	Description
Traffic by Initiator	Displays a graph of the 10 most active hosts on the monitored network segment based on the number of kilobytes transmitted by the host.
Traffic by Responder	Displays a graph of the 10 most active hosts on the monitored network segment based on the number of kilobytes received by the host.
Traffic by Port	Displays a graph of the 10 most active ports on the monitored network segment based on the number of kilobytes transmitted.
Traffic by Service	Displays a graph of the 10 most active services on the monitored network segment based on the number of kilobytes transmitted.

The Flow Summary Page Functions table describes the different actions you can perform on the Flow Summary page.

Flow Summary Page Functions

To...	You can...
modify the time and date range for the Flow Summary page	find more information in Setting Event Time Constraints on page 1597.
learn more about flow graphs	find more information in Understanding Flow Data on page 1072.
manipulate flow graphs	find more information in Manipulating Flow Data Graphs on page 1100.
detach one of the flow graphs from the page	click **View** on the graphs you want to detach. For more information on detached graphs, see Detaching Flow Data Graphs on page 1114.

Note that you can perform almost all the same actions on flow summary graphs that you can perform on flow data graphs. However, because the graphs on the Flow Summary page are based on aggregated data, you cannot examine the individual flow events on which the graphs are based. In other words, you cannot drill down to a flow data table view from a flow summary graph.

To view flow summary graphs for your monitored network segment:

Access: Restricted Event Analyst

1. Select **Analysis & Reporting > Event Summary > Flow Summary**.

   The Flow Summary page appears for the current time range on your Defense Center.

Keysight, EX-1004
Page 1117

2. From the **Select Detection Engine** list, select the detection engine whose summary you want to view, or select **All** to view a summary of all detection engines.

Creating Traffic Profiles

Platform: DC

Component: RNA

A traffic profile is just that—a profile of the traffic on your network, based on flow data collected over a time span that you specify. You can use flow data collected by your 3D Sensors with RNA, flow 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.

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 flow data collected over the last week. You can change the PTW to be as short as an hour or as long as several weeks.

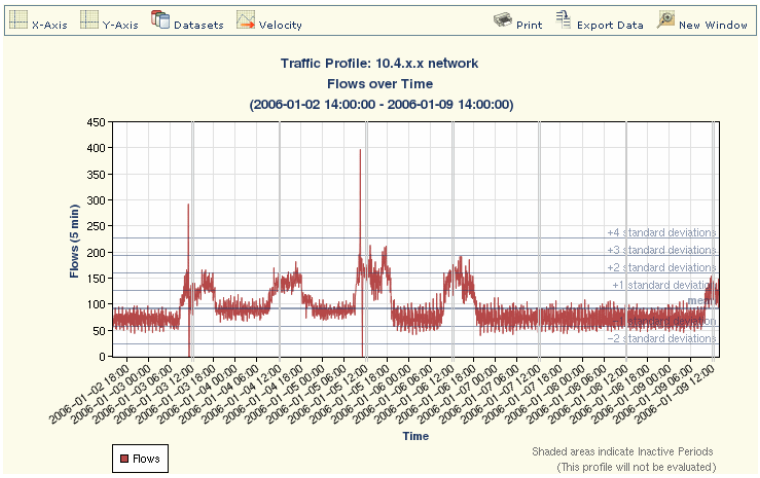
When you first activate a traffic profile, it collects and evaluates flow data according to the criteria you have set, for a learning period equal in time to the PTW. The Defense Center does not evaluate rules you have written against the traffic profile until the learning period is complete.

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 flow 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 flow 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 flow data.

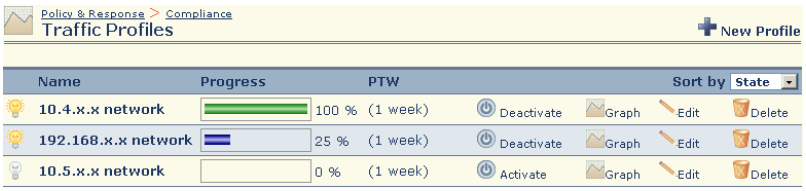
Keysight, EX-1004
Page 1118

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 12:30AM.



After you create and activate a traffic profile and its learning period is complete, you can create compliance 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 flows) 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 compliance 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 Compliance Policies on page 1214.

You create traffic profiles on the Traffic Profiles page.



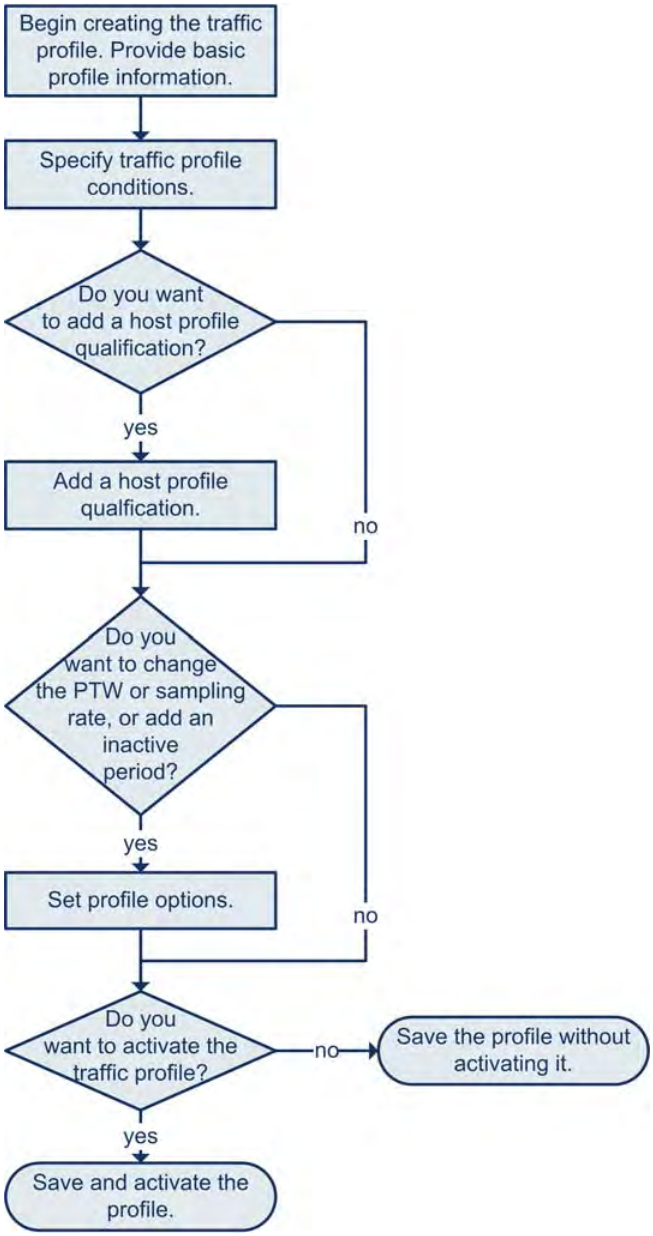
The light bulb icon next to each profile indicates whether the profile is active. If you want to base a compliance rule on a traffic profile change, you must activate the profile. If the light bulb icon is lit, the profile is active. If it is dark, the profile is inactive. For more information, see Activating and Deactivating Traffic Profiles on page 1129.

The progress bar shows the status of the traffic profile's learning period. When the progress bar reaches 100%, compliance rules written against the profile will trigger.

Keysight, EX-1004
Page 1119

You can sort traffic profiles by state (active versus inactive) or alphabetically by name using the **Sort by** drop-down list.

The following diagram illustrates the process for creating a traffic profile.



Keysight, EX-1004
Page 1120

For more information, see:

- Providing Basic Profile Information on page 1121
- Specifying Traffic Profile Conditions on page 1122
- Adding a Host Profile Qualification on page 1124
- Setting Profile Options on page 1127
- Saving a Traffic Profile on page 1129
- Activating and Deactivating Traffic Profiles on page 1129
- Editing a Traffic Profile on page 1130
- Understanding Condition-Building Mechanics on page 1130

Providing Basic Profile Information

Platform: DC

Component: RNA

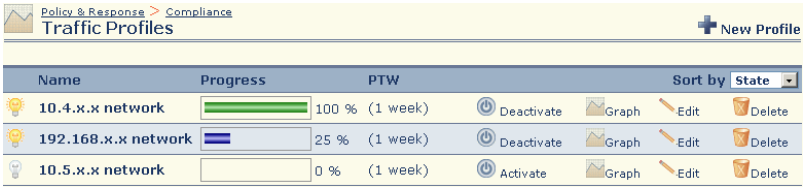
When you create a traffic profile, you must give it a name and, optionally, a short description.

To begin creating a traffic profile:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Traffic Profiles**.

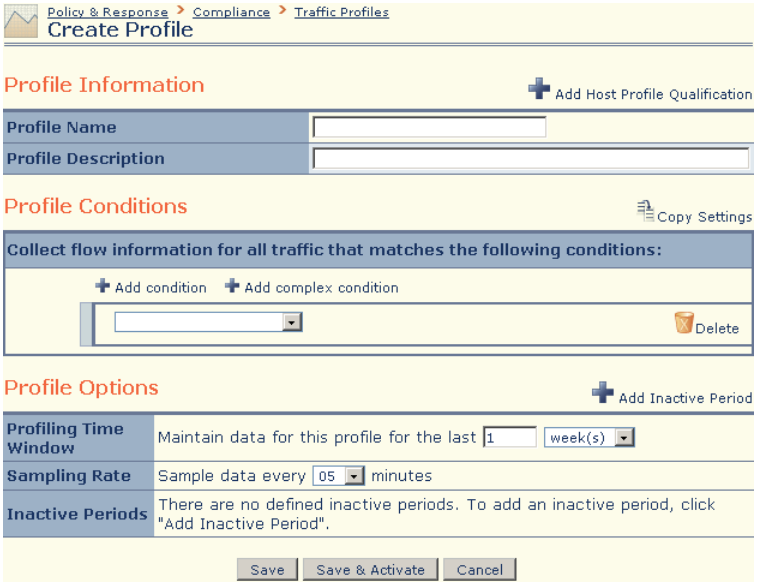
   The Traffic Profiles page appears.



Keysight, EX-1004
Page 1121

2. Click **New Profile**.

   The Create Profile page appears.



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

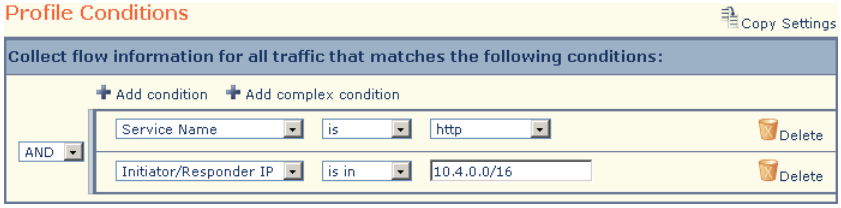
Platform: DC

Component: RNA

Profile conditions constrain the kinds of flow 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.

Keysight, EX-1004
Page 1122

For example, the traffic profile conditions in the following graphic collects HTTP flows on the 10.4.x.x subnet.



You build traffic profile conditions in the **Profile Conditions** section of the Create Profile page. See Understanding Condition-Building Mechanics on page 1130 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 1123.

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

Platform: DC

Component: RNA

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 flow is the initiator and which is the responder. When RNA 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 How is NetFlow Data Different from RNA Flow Data? on page 1142.

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 1557 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	Type the port number.
Protocol	Type `TCP` or `UDP` as the protocol.

Keysight, EX-1004
Page 1123

Syntax for Profile Conditions (Continued)

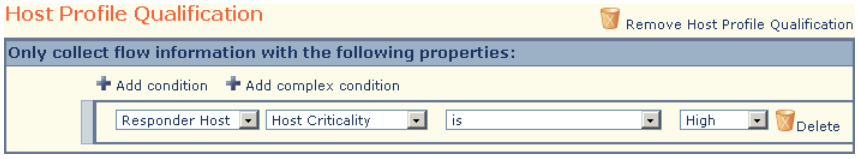
If you specify...	Select an operator, then...
Flow Type	Select whether you want to use flow data collected by 3D Sensors with RNA or by NetFlow-enabled devices in the traffic profile. If you do not specify a flow type, the traffic profile includes both.
NetFlow Device	Select the NetFlow-enabled device 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 system settings), the NetFlow Device drop-down list is blank.
Service Name	Select one or more service names from the list of available services.

Adding a Host Profile Qualification

Platform: DC

Component: RNA

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 flow data only for hosts that are assigned a host criticality of **high**.



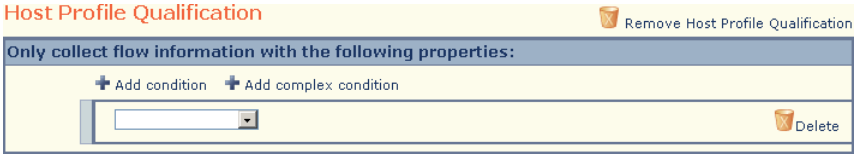
To use a host profile qualification, the host must exist in the database and the host profile property you want to use as a qualification must already be included in the host profile. For example, if you configure a compliance policy rule to trigger when an intrusion event is generated on a host running 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: P&R Admin/Admin

1. On the Create Rule page, click **Add Host Profile Qualification**.

   The Host Profile Qualification section appears.



Keysight, EX-1004
Page 1124

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 1130 for information building conditions.

   The syntax you can use to build conditions is described in Syntax for Host Profile Qualifications on page 1125.

   TIP! To remove a host profile qualification, click **Remove Host Profile Qualification**.

Syntax for Host Profile Qualifications

Platform: DC

Component: RNA

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 type, continue building your host profile qualification condition, as described in the Syntax for Host Profile Qualifications table.

Although you can configure the RNA detection 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 flow data exported by NetFlow-enabled devices, keep in mind that NetFlow records do not contain information about which host in the flow is the initiator and which is the responder. When RNA 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 How is NetFlow Data Different from RNA Flow Data? on page 1142.

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 Name	Select one or more operating system names from the drop-down list.

Keysight, EX-1004
Page 1125

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 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 Assigning a Host Criticality Value to a Host on page 986.
VLAN ID	Type the VLAN ID number of the host.
Service > Service Name	Select one or more services from the drop-down list.
Service > Service Port	Type the service port number.
Service > Service Protocol	Select one or more protocols from the drop-down list.
Client Application > Application Type	Select one or more application types from the drop-down list.
Client Application > Application	Select one or more applications from the drop-down list.
Client Application > Application Version	Type the application version.
MAC Address > MAC Address	Type all or part of the MAC address of the host.

Keysight, EX-1004
Page 1126

Syntax for Host Profile Qualifications (Continued)

If you specify...	Select an operator, then...
MAC Address > MAC Type	Select whether the MAC type is **ARP/DHCP Detected**. That is, select whether RNA positively identified the MAC address as belonging to the host (**is ARP/DCHCP Detected**), whether RNA is seeing many hosts with that MAC address because, for example, there is a router between the sensor 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: • 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. For more information on host attributes, see Working with User-Defined Host Attributes on page 987.

Setting Profile Options

Platform: DC

Component: RNA

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.

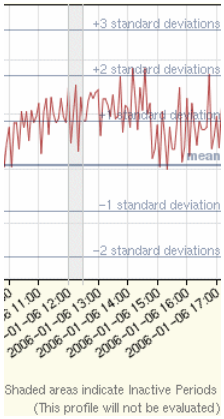
Also, traffic profiles are based on aggregated flow data. By default, traffic profiles generate statistics on flow events generated RNA over five-minute intervals. However, you can set this sampling rate anywhere between the default five minutes and one hour.

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.

Keysight, EX-1004
Page 1127

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, as displayed in the following graphic.



Access: P&R Admin/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.

Keysight, EX-1004
Page 1128

Saving a Traffic Profile

Platform: DC

Component: RNA

Use the following procedure to save a traffic profile.

To save a traffic profile:

Access: P&R Admin/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

Platform: DC

Component: RNA

When you want it to begin profiling the traffic on a monitored network segment, you must activate a traffic profile.

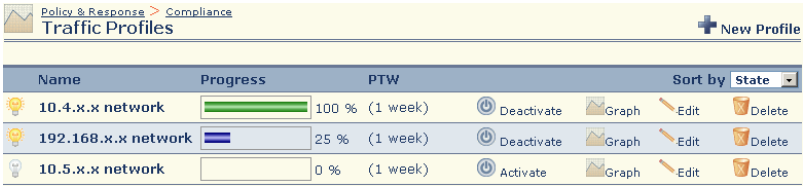
Deactivate a profile when you want it to stop collecting and evaluating flow data. Rules written against deactivated traffic profiles do not trigger. In addition, deactivating a traffic profile deletes all data collected and aggregated by the profile. If you later reactivate a deactivated traffic profile, you must wait the length of its PTW before rules written against it will trigger.

To activate or deactivate a traffic profile:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Traffic Profiles**.

   The Traffic Profiles page appears.



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**.

Keysight, EX-1004
Page 1129

Editing a Traffic Profile

Platform: DC

Component: RNA

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 1129.

To edit a traffic profile:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > 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.

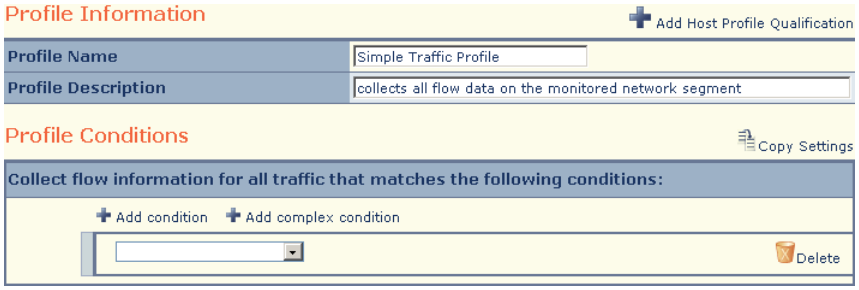
Understanding Condition-Building Mechanics

Platform: DC

Component: RNA

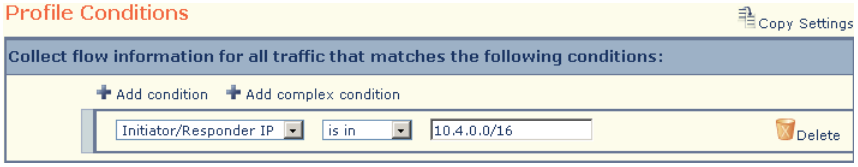
You build traffic profiles by specifying the conditions they use to collect data. You can create either simple conditions or more elaborate constructs with nested conditions.

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.

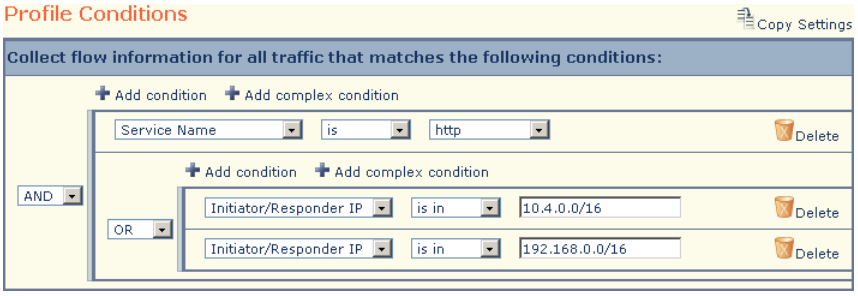


Keysight, EX-1004
Page 1130

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 1131
- Adding and Linking Conditions on page 1133
- Using Multiple Values in a Condition on page 1136

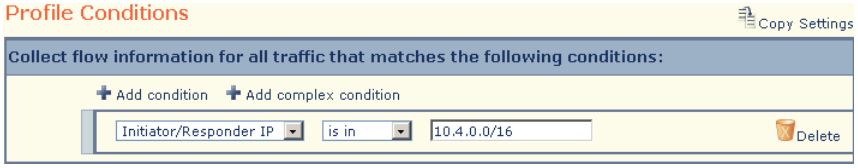
Building a Single Condition

Platform: DC

Component: RNA

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`.



Keysight, EX-1004
Page 1131

The following steps explain how to build this traffic profile condition.

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Traffic Profiles**.

   The Traffic Profiles page appears.

2. Click **New Profile**.

   The Create Profile page appears.

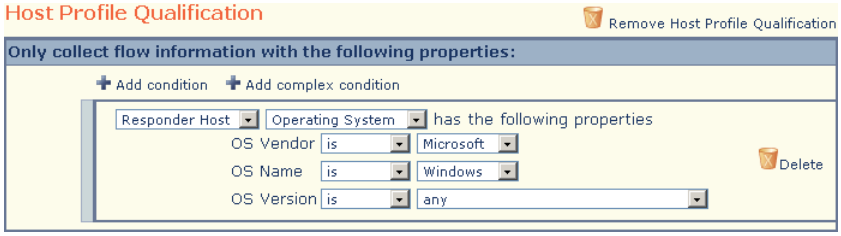
3. Under **Profile Conditions**, begin building the profile's single condition by selecting **Initiator/Responder IP** from the first (category) drop-down list.

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 59.

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 flow data only if the responding host in the detected flow is running a version of Microsoft Windows.



The following steps explain how to build this host profile qualification.

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > 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 flow.

Keysight, EX-1004
Page 1132

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 1136.

For information on the syntax for building traffic profile conditions and host profile qualifications, see:

- Syntax for Traffic Profile Conditions on page 1123
- Syntax for Host Profile Qualifications on page 1125

Adding and Linking Conditions

Platform: DC

Component: RNA

You can create simple traffic profile conditions and host profile qualifications, or you can create more elaborate constructs by combining and nesting conditions.

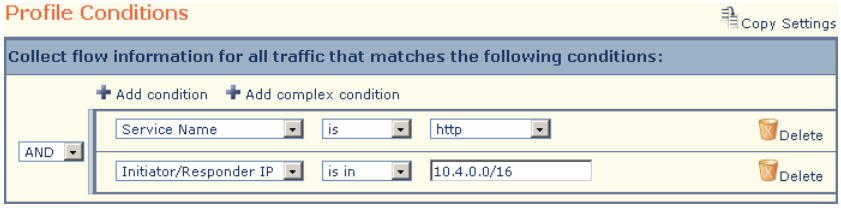
When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together.

- 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.

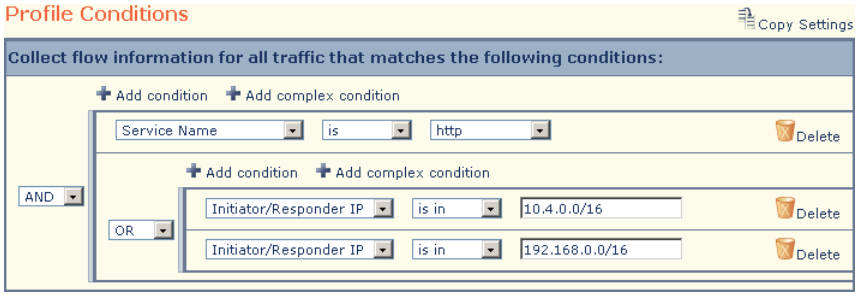
For example, the following traffic profile contains two conditions linked by **AND**. This means that the traffic profile collects flow data only if both conditions are

Keysight, EX-1004
Page 1133

true. In this example, it collects HTTP flows for all hosts with an IP addresses in the 10.4.x.x subnet.



In contrast, the following traffic profile, which collects flow 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	Service Name is http
B	IP Address is in 10.4.0.0/16
C	IP Address is in 192.168.0.0/16

Keysight, EX-1004
Page 1134

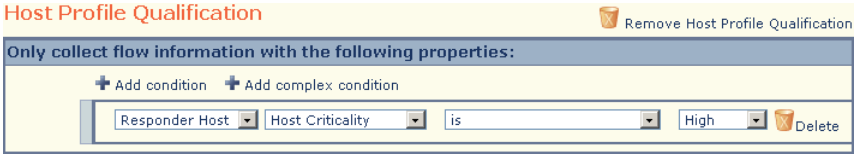
To add a single condition:

Access: P&R Admin/Admin

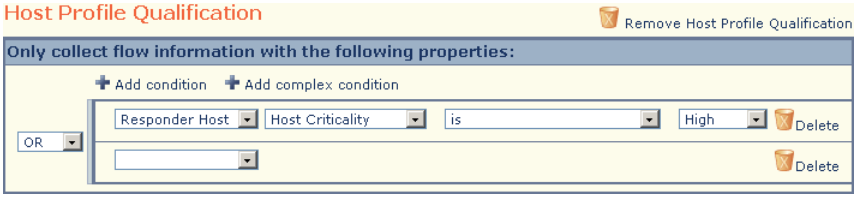
- To add a single condition, click **Add condition** above the current condition.

  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:



  The result is:



Keysight, EX-1004
Page 1135

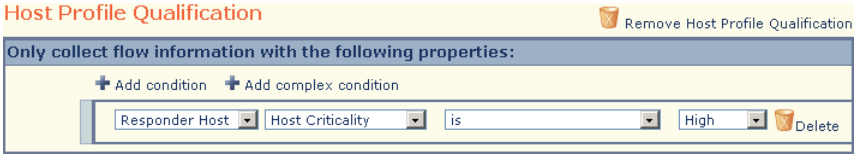
To add a complex condition:

Access: P&R Admin/Admin

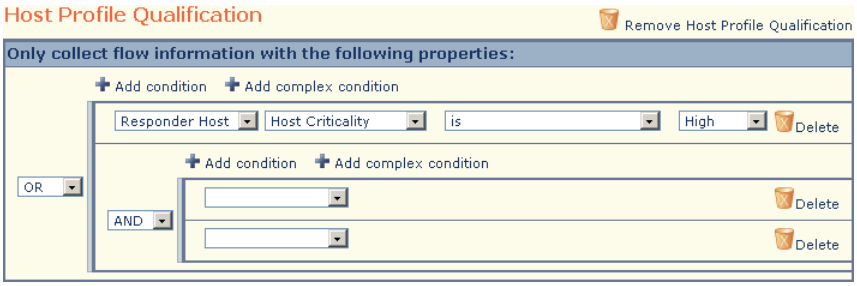
▶ Click **Add complex condition** above the current condition.

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 result is:



To link conditions:

Access: P&R Admin/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

Platform: DC

Component: RNA

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.

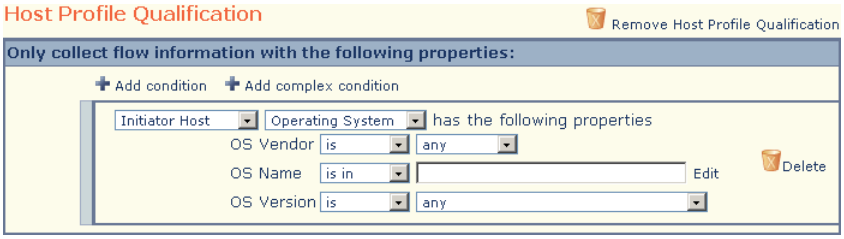
Keysight, EX-1004
Page 1136

To include multiple values in one condition:

Access: P&R Admin/Admin

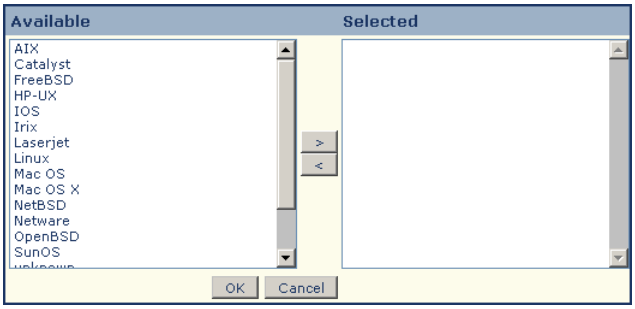
1. Build a condition, choosing **is in** or **is not in** as the operator.

   The drop-down list changes to a text field.



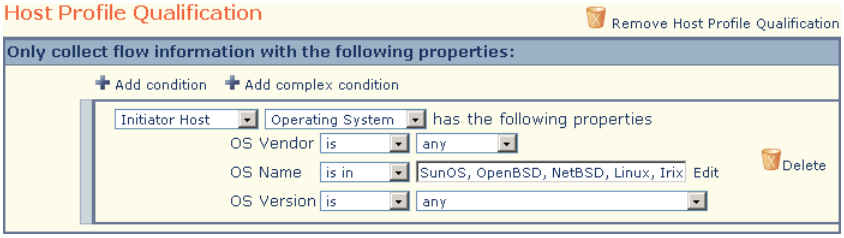
2. Click anywhere in the text field or on the **Edit** link.

   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.



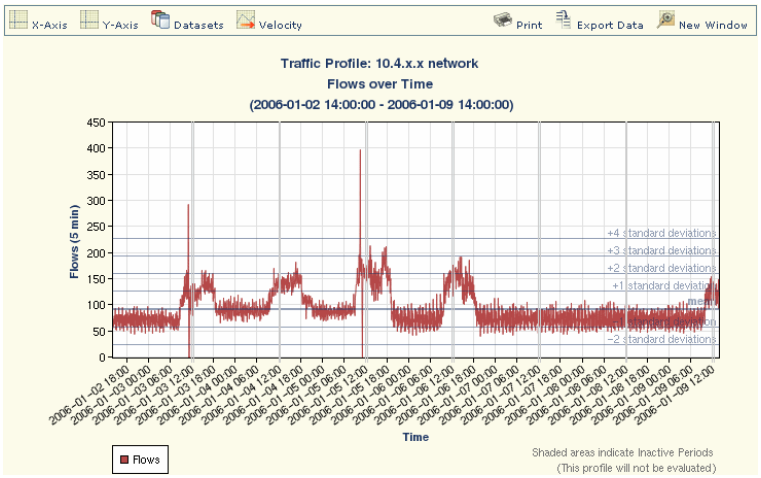
Keysight, EX-1004
Page 1137

Viewing Traffic Profiles

Platform: DC

Component: RNA

Because traffic profiles are based on flow 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 12:30AM.



You can perform almost all the same actions on traffic profile graphs that you can perform on flow data graphs. However, because traffic profiles are based on aggregated data (flow summaries), you cannot examine the individual flow events on which the graphs are based. In other words, you cannot drill down to a flow data table view from a traffic profile graph. See Viewing Flow Data on page 1093 for more information. In addition, traffic profiles appear as detached graphs. For more information, see Detaching Flow Data Graphs on page 1114.

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.

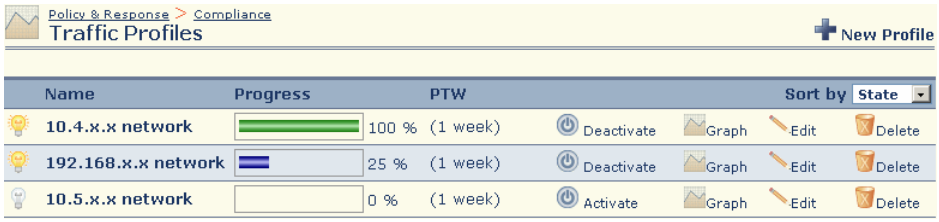
Keysight, EX-1004
Page 1138

To view a traffic profile graph for a traffic profile:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Traffic Profiles**.

   The Traffic Profiles page appears.



2. Next to the traffic profile for which you want to view the graph, click **Graph**.

   The graph for the traffic profile appears in a separate browser window.

Keysight, EX-1004
Page 1139

Chapter 26
Introduction to NetFlow

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 is a sequence of packets that represents a connection between a source and destination host, using specific ports, protocol, and service.) You can use this flow data to supplement the flow data gathered directly by your 3D Sensors with RNA. This is useful if you have NetFlow-enabled devices deployed on networks that your sensors cannot monitor, because you can use NetFlow data to monitor those networks.

Note that, as with RNA data gathered directly by your sensors, you must use a Defense Center to configure NetFlow data collection and to view the collected data. In addition, your deployment must include at least one 3D Sensor with RNA that can communicate with your NetFlow-enabled devices. Although you can use NetFlow-enabled devices exclusively to monitor your network, the Sourcefire 3D System uses RNA detection engines on 3D Sensors to analyze NetFlow data.

For more information, see the following sections:

- Understanding NetFlow on page 1141
- Planning your NetFlow Deployment on page 1148
- Configuring NetFlow Data Collection on page 1150

Keysight, EX-1004
Page 1140

Understanding NetFlow

Platform: DC

Component: RNA + NetFlow

You can use the Sourcefire 3D System to analyze exported NetFlow records and thereby supplement the data gathered directly by the detection engines on 3D Sensors with RNA. For more information, see the following sections:

- What Information Does NetFlow Collect and Provide? on page 1141
- How is NetFlow Data Different from RNA Flow Data? on page 1142
- What are the Prerequisites for NetFlow Data Collection? on page 1146

What Information Does NetFlow Collect and Provide?

Platform: DC

Component: RNA + NetFlow

NetFlow collects and reports records of active sessions involving the hosts on your network. RNA uses this data to generate flow events and, optionally, to update the network map with host and service information.

There are several differences between the NetFlow data and the flow data detected directly by detection engines on 3D Sensors with RNA. You **must** keep these differences in mind when performing any kind of analysis that includes or uses information derived from NetFlow records. For more information, see How is NetFlow Data Different from RNA Flow Data? on page 1142.

Flow Data

RNA generates a flow event when a connection is terminated between a monitored host and any other host. Flow events are one of the ways the Sourcefire 3D System alerts you to the activity on your network so that you can respond appropriately. RNA uses NetFlow data to generate flow events, just as it does for flows directly detected by 3D Sensors.

Flow events based on NetFlow data include information about the collected traffic, including:

- the IP address of the NetFlow-enabled device that exported the record of the flow
- the timestamps of the first and last packet of the transaction
- the IP addresses of the hosts involved in the flow
- the ports and protocol used by the flow
- any TCP flags detected in the flow

Keysight, EX-1004
Page 1141

- the user logged into the hosts involved in the flow (if RUA is enabled and the hosts are on the monitored network)
- the number of packets and bytes sent and received by the monitored hosts

You can view NetFlow-based flow events using a set of predefined workflows, or you can create custom workflows. You can choose to view NetFlow data in a table format or as graphs. You can also use the flow data to create and view a profile of your normal network traffic, called a traffic profile.

Further, you can use flow events and traffic profiles as the basis for compliance rules. Compliance rules are one piece of the Sourcefire 3D System's powerful policy and response feature that lets you respond in real time to threats on your network. Within a compliance policy, compliance rules help you describe the types of network activity that constitute policy violations for your organization. Using NetFlow-based flow events, you can create compliance rules that trigger when:

- a flow event occurs that meets specific criteria
- your network traffic deviates from your normal network traffic pattern as characterized in a traffic profile that was built using NetFlow data

Using policy and response, you can configure RNA to respond to suspicious activity in real-time by using compliance policies to send alerts or launch remediations. For more information, see:

- Working with Flow Data and Traffic Profiles on page 1071
- Configuring Compliance Policies and Rules on page 1212
- Configuring Responses for Compliance Policies on page 1281

Network Map (Host and Service Data)

RNA aggregates data for each host and service into the network map, which is an up-to-date, detailed representation of your network assets. The network map is continuously updated as network traffic indicates a change or a new host. You can configure your RNA detection policy to add hosts to the network map based on data exported by NetFlow-enabled devices. You can also configure your RNA detection policy to add services to the network map based on NetFlow data.

For more information, see Using the Network Map on page 923.

How is NetFlow Data Different from RNA Flow Data?

Platform: DC

Component: RNA + NetFlow

There are several differences between the data collected and exported by NetFlow-enabled devices and the flow data detected directly by the detection engines on 3D Sensors with RNA. It is important to keep these differences in mind when performing any kind of analysis that includes flow data, or that uses

Keysight, EX-1004
Page 1142

information derived from flow data. For more information, see the following sections.

Number of Flow Events Generated per Monitored Session

RNA generates flow events when connections between monitored hosts and any other hosts are terminated. For flows detected directly by the detection engines on 3D Sensors with RNA, RNA generates a single bidirectional flow event. However, because NetFlow-enabled devices export unidirectional flow data, RNA generates at least two flow events for each flow detected by NetFlow-enabled devices, depending on how you configured the devices.

If you configured your NetFlow-enabled devices to output records only when the monitored session ends, RNA generates two flow events for that session. On the other hand, if you configured your NetFlow-enabled devices to output records at a fixed interval even if a flow is still ongoing, RNA generates a flow event for each record exported by the device. For example, if you configure your NetFlow-enabled devices to output flow records for long-running flows every five minutes, and a particular flow lasts twelve minutes, RNA will generate six flow events for that session: one pair of events for the first five minutes, one pair for the second five minutes, and a final pair when the connection is terminated.

Further, flow summary counts (see Understanding Flow Summary Data on page 1076) are incremented by one for every flow event that RNA generates. This means that flow summaries based on NetFlow data can keep an inflated count of the number of flows that are actually occurring on your network.

Keep these points in mind when performing any analysis that includes statistics on the number of flows, for example:

- viewing flow graphs that display the number of flows over time
- viewing flow summary data
- viewing traffic profiles
- creating compliance rules based on traffic profile changes
- creating compliance rules based on flow duration

NetBIOS and Application Data

Because NetFlow-enabled devices do not capture NetBIOS information, flow events based on NetFlow data do not contain NetBIOS information. For the same reason, flow events based on NetFlow data do not include any information on flows' client applications (type, name, or version) or any URLs accessed during monitored sessions.

Keysight, EX-1004
Page 1143

Keep these points in mind when performing any analysis that requires NetBIOS or application data, including:

- searching for flows that have specific NetBIOS or application data
- creating compliance rules based on the NetBIOS and application data in flow events
- adding host profile qualifications to compliance rules or traffic profiles

Host and Operating System Data

When an RNA detection engine directly detects a flow, if you applied a host license to the Defense Center, RNA adds the hosts involved in the flow to the network map.

Although you can configure RNA to add hosts to the network map based on NetFlow records, those records do not include operating system or NetBIOS data for the hosts involved in the flow, nor can RNA 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.

Because there is no operating system information available for hosts added to the network map based on NetFlow data, the Sourcefire 3D System cannot determine which vulnerabilities might affect those hosts, unless you use the host input feature to manually set the hosts' operating system identity.

Keep these points in mind when performing any analysis that requires operating system or other host data that is not included in NetFlow records, including:

- working with hosts on the network map
- searching for hosts based on operating system, host type, or NetBIOS name
- creating compliance rules based on RNA's detection of a host with a particular operating system, host type, or NetBIOS name
- adding host profile qualifications to compliance rules or traffic profiles
- working with vulnerabilities

Service Data

When an RNA detection engine directly detects a flow, it attempts to identify the service involved in the monitored session by examining the content of the flow. If you applied a host license to the Defense Center, RNA adds the service to the network map.

Although you can configure RNA to add services to the network map based on NetFlow records, RNA uses the port-service correlation in `/etc/sf/services` to extrapolate service names because flow records exported by NetFlow do not contain service information. This means that services running on non-standard ports can be unidentified or misidentified. In addition, RNA is unable to determine vendors and versions for services added to the network map based on NetFlow

Keysight, EX-1004
Page 1144

data. You can, however, manually provide service information using the host input feature.

Keep these points in mind when performing analysis that requires service data, including:

- working with services on the network map
- searching for services
- creating compliance rules based on RNA's detection of a particular service or on flow events involving a particular service
- creating traffic profiles based on flow events that involve a particular service
- adding host profile qualifications to compliance rules or traffic profiles

Initiator and Responder Information

For flows detected directly by the detection engines on 3D Sensors with RNA, RNA can identify which host is the initiator, or source, and which is the responder, or destination.

However, flow data exported by NetFlow does not contain initiator or responder information. When RNA 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, RNA considers the host using the lower-number port to be the responder. If only one of the hosts is using a well-known port, RNA 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 service information in `/etc/sf/services` on the 3D Sensor.

Keep these points in mind when performing analysis that requires knowing which host is the initiator and which is the responder, including:

- searching for flows or flow summaries
- using RUA to determine the user identities of the users logged into the hosts involved in a flow
- creating compliance events or traffic profiles based on the initiator or responder information in flow events
- adding flow trackers to compliance rules
- adding host profile qualifications to compliance rules or traffic profiles

Keysight, EX-1004
Page 1145

TCP Flag Information

Flow data exported by NetFlow-enabled devices contains information on the TCP flags set for the flow. Flows detected directly by the detection engines on 3D Sensors with RNA do not contain TCP flag information.

Keep this point in mind when performing any analysis that requires TCP flag information, including:

- searching for flows or flow summaries
- creating compliance rules based on the TCP flag information in flow events

What are the Prerequisites for NetFlow Data Collection?

Platform: DC

Component: RNA + NetFlow

Before you use the Defense Center to configure an RNA detection policy, which controls which networks are monitored by specific NetFlow-enabled devices, there are several prerequisites you must meet. Because NetFlow data supplements the data gathered by 3D Sensors with RNA, the prerequisites for NetFlow data collection include those for RNA. There are also several NetFlow-specific prerequisites, as described in the following sections. For more detailed information on RNA-only prerequisites, see What are the Prerequisites for RNA Data Collection? on page 887.

Configure and Deploy at Least One 3D Sensor with RNA

Make sure your deployment includes at least one 3D Sensor with RNA that can communicate with your NetFlow-enabled devices. The sensing interface for at least one RNA detection engine must be connected to a network where it can collect the data that your NetFlow-enabled devices export; the Sourcefire 3D System uses RNA detection engines on 3D Sensors to analyze NetFlow data. For more information, see Planning your NetFlow Deployment on page 1148.

Keysight, EX-1004
Page 1146

Configure NetFlow-Enabled Devices to Export NetFlow Records

Make sure you enable the NetFlow feature on the routers or other NetFlow-enabled devices you plan to use to export flow data. In addition, there are a few points to keep in mind:

- You must configure your NetFlow-enabled devices to export NetFlow version 5 data.
- Your NetFlow-enabled devices must be set to export records to a destination network where the sensing interface of a 3D Sensor with RNA is connected. Because the sensing interfaces on 3D Sensors do not have IP addresses, the Sourcefire 3D System does not support the direct collection of NetFlow records.
- Sourcefire recommends that you configure your NetFlow-enabled devices to output flow data only when monitored sessions close. If you configure your NetFlow-enabled devices to output flow records at fixed intervals, analyzing flow data can be more complicated; see Number of Flow Events Generated per Monitored Session on page 1143.
- If enabled, the Sampled NetFlow feature that is available on some NetFlow-enabled devices collects NetFlow statistics on a only subset of packets that pass through the devices. Although turning on this feature can improve CPU utilization on the NetFlow-enabled device, it can affect the data you are collecting for analysis by the Sourcefire 3D System.

Apply a Host License to the Defense Center

Use the system settings to apply a host license to your Defense Center. This ensures that you can use RNA to monitor hosts on your network, including hosts discovered by NetFlow-enabled devices. For more information, see Managing Feature Licenses on page 1787.

Note that if you plan to use only NetFlow-enabled devices to monitor your network, you do not need to apply host license to the Defense Center. However, because RNA will not build the network map or store information about the hosts on your monitored network, your analysis will be limited to working with the flow data exported by your NetFlow-enabled devices.

Configure RNA Settings in the System Policy

Use the system policy to control the kinds and amount of RNA data (including data exported by NetFlow-enabled devices) stored in the Defense Center's database, and therefore determine the data that other parts of the Sourcefire 3D System can use. For more information, see Configuring RNA Settings on page 1756.

Keysight, EX-1004
Page 1147

Planning your NetFlow Deployment

Platform: DC

Component: RNA + NetFlow

Sourcefire RNA allows your organization to confidently monitor and protect your network using a combination of forensic analysis, behavioral profiling, and built-in alerting and remediation. 3D Sensors with RNA passively observe your organization's network traffic and analyze it to provide you with a complete profile of your network.

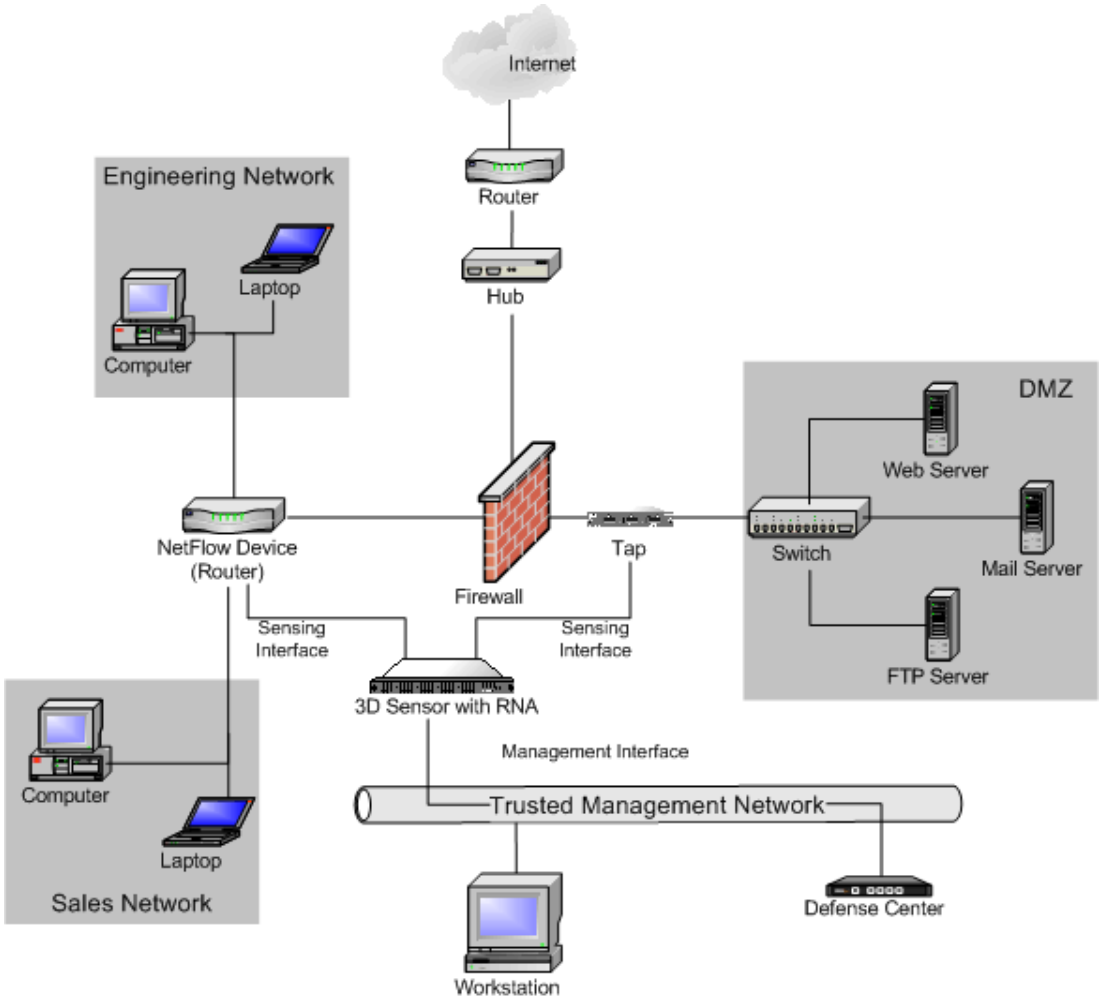
If your organization uses NetFlow-enabled devices to collect IP traffic information, you can use the NetFlow records generated by those devices to supplement the data gathered directly by 3D Sensors with RNA. This is useful if you have NetFlow-enabled devices deployed on networks that your 3D Sensors with RNA cannot see, because you can use the data exported by those NetFlow-enabled devices to monitor those networks.

Note that your deployment **must** includes at least one 3D Sensor with RNA that can communicate with your NetFlow-enabled devices. Although you can use NetFlow-enabled devices exclusively to monitor your network, the Sourcefire 3D System uses RNA detection engines on 3D Sensors to analyze NetFlow data.

Also note that Sourcefire recommends that you **not** monitor the same network segment with NetFlow-enabled devices and 3D Sensors with RNA. Although you can configure the system policy to drop duplicate flows detected by 3D Sensors with RNA, and to drop duplicate flows that are exported by NetFlow-enabled devices, you cannot selectively drop a flow that is detected by both a sensor and a NetFlow-enabled device. That is, if you are monitoring the same network segment with a NetFlow-enabled device and a 3D Sensor with RNA, RNA will generate at least three flow events for each detected flow—one that represents the flow as detected by the sensor, and at least two for the flow as detected by NetFlow-enabled device. You **cannot** drop these duplicate events.

Keysight, EX-1004
Page 1148

The following diagram illustrates a sample Sourcefire 3D System deployment using a NetFlow-enabled device. Although the 3D Sensor with RNA can directly observe the traffic to and from the DMZ, the sensor cannot see the traffic to the engineering and sales networks. However, a router configured as a NetFlow-enabled device can see the traffic. The NetFlow-enabled device broadcasts flow data such that a sensing interface on the 3D Sensor with RNA can passively observe it.



In this scenario, you could set up one passive interface set for both sensing interfaces, then use a single RNA detection engine to analyze NetFlow data and to observe DMZ traffic. Alternately (depending on the model of the 3D Sensor), you could configure two interface sets and then use separate RNA detection engines to analyze the network traffic—one for NetFlow data and one for DMZ traffic.

Keysight, EX-1004
Page 1149

Configuring NetFlow Data Collection

Platform: DC

Component: RNA + NetFlow

After you configure and deploy your NetFlow-enabled devices, you can use the Defense Center to configure the Sourcefire 3D System to collect and analyze NetFlow data. Then, you can begin using this data to supplement your RNA analysis.

To configure NetFlow data collection:

1. Make sure you have met the prerequisites described in What are the Prerequisites for NetFlow Data Collection? on page 1146.
2. Use the system settings to apply a NetFlow license to your Defense Center; see Licensing NetFlow on page 1150.
3. Use the system settings to specify the NetFlow-enabled devices you are going to use; see Configuring the Sourcefire 3D System to see NetFlow Devices on page 1151.
4. Create, save, and apply an RNA detection policy that monitors networks with your NetFlow-enabled devices; see Using NetFlow in an RNA Detection Policy on page 1152.

Licensing NetFlow

Platform: DC

Component: RNA + NetFlow

For you to use NetFlow-enabled devices in your Sourcefire 3D System deployment, your Defense Center must have a NetFlow feature license. This license specifies the number of NetFlow licenses you can use.

Note that in a high-availability environment, both Defense Centers in the pair must have NetFlow feature licenses. If one Defense Center does not have a NetFlow license, it will ignore NetFlow data.

If your organization has purchased a NetFlow feature license but you do not have the license file, you can request it from the Defense Center web interface. Before beginning, you should have the 12-digit feature license serial number provided by Sourcefire. If you do not have the serial number, you can find it by logging into the Sourcefire Support Site (https://support.sourcefire.com/), clicking **Account**, then clicking **Products & Contracts**. The serial number appears in the Sourcefire Software & Licenses section.

TIP! After you add a NetFlow feature license, you can view it in the system settings. For more information, see Viewing Feature Licenses on page 1790.

Keysight, EX-1004
Page 1150

To add a NetFlow feature license:

Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **License**.

   The License page appears.

3. Click **Add New License**.

   The Add Feature License page appears.

4. Do you already have the license file?
   - If yes, skip to step 6.
   - If no, click **Get License**.

   The Licensing Center web site appears.

   IMPORTANT! If your web browser cannot access the Internet, you must switch to a host that can access it. Copy the license key at the bottom of the page and browse to https://keyserver.sourcefire.com/.

5. Follow the on-screen instructions for a feature license to obtain your license file, which will be sent to you in an email.

6. Copy the license file from the email you received, paste it into the **License** field, and click **Submit License**.

   If the license is valid, the NetFlow license is installed on the Defense Center.

7. Continue with the next procedure, Configuring the Sourcefire 3D System to see NetFlow Devices.

Configuring the Sourcefire 3D System to see NetFlow Devices

Platform: DC

Component: RNA + NetFlow

Before you can begin collecting NetFlow data, you must use the system settings to specify the NetFlow-enabled devices you are going to use. You must configure these NetFlow-enabled devices to export NetFlow version 5 data.

You can specify as many NetFlow-enabled devices as are permitted by your NetFlow feature license. When you configure the RNA detection policy to monitor networks with NetFlow-enabled devices, you can use only those NetFlow-enabled devices you have added to the Defense Center.

Keysight, EX-1004
Page 1151

To add NetFlow-enabled devices for flow data collection:

Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **NetFlow Devices**.

   The NetFlow Devices page appears.



3. Click Add Device to add a NetFlow-enabled device.



4. In the **IP Address** field, enter the IP address of the NetFlow-enabled device you want to use to collect flow data.
5. To add additional NetFlow-enabled devices, repeat steps 3 and 4.

TIP! To remove a NetFlow-enabled device, click **Delete** next to the device you want to remove. Keep in mind that if you remove a NetFlow-enabled device from the system settings, you should also remove it from your RNA detection policy. For more information, see Editing an RNA Detection Policy on page 921.

---

6. Click **Save**.

   The list of NetFlow-enabled devices is saved.

7. Continue with the next procedure, Using NetFlow in an RNA Detection Policy.

Using NetFlow in an RNA Detection Policy

Platform: DC

Component: RNA + NetFlow

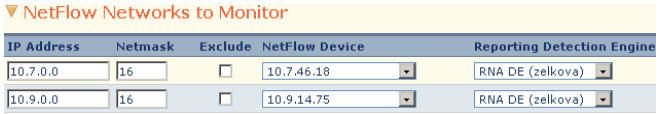
RNA detection policies control which networks are monitored with NetFlow-enabled devices. Note that there are several points you must keep in

Keysight, EX-1004
Page 1152

mind when configuring your RNA detection policy using NetFlow-enabled devices:

- Although you can use NetFlow-enabled devices exclusively to monitor your network, your deployment must include at least one 3D Sensor with RNA. This sensor must be deployed such that one of its sensing interfaces is connected to the destination network where your NetFlow-enabled device exports flow records. This ensures that the RNA detection engines on the sensor can collect and analyze the NetFlow data. For more information, see Planning your NetFlow Deployment on page 1148.
- Subnet detection (see How Do I Manage my RNA Deployment with Subnet Detection? on page 897) is not supported for networks you are monitoring with NetFlow-enabled devices. If you want to use NetFlow-enabled devices to monitor all or part of your network, you must know where and how they are physically deployed in order to monitor the appropriate network segments. You must explicitly configure at least one RNA detection engine to analyze the data exported by your NetFlow-enabled devices. The detection engine that analyzes the NetFlow data for a network segment is called the *reporting detection engine* for that segment.
- Flow filtering, that is, setting a data collection mode that affects how flow data is collected (see Filter Flows as Needed on page 893), is not supported for networks you are monitoring with NetFlow-enabled devices. You can, however, entirely exclude networks from monitoring.
- The port exclusion feature (see Exclude Ports as Needed on page 893) only affects data collected by 3D Sensors with RNA. You cannot configure NetFlow-enabled devices to exclude ports from monitoring.

The following graphic shows an example NetFlow configuration in an RNA detection policy. The RNA detection engine on the sensor named `zelkova` is the reporting detection engine for the 10.7.0.0/16 and 10.9.0.0/16 subnets, which are being monitored by two separate NetFlow-enabled devices: 10.7.46.18 and 10.9.14.75, respectively.



RNA detection policies also include a number of settings that control how RNA events and flow data are collected by RNA, including whether you want to generate hosts and services from NetFlow data. If you enable these options, when a flow exported by a NetFlow-enabled device involves a previously undetected host or service running on a host, that host or service is added to the network map. Note that the information about hosts and services generated from NetFlow data is limited; for more information, see What Information Does NetFlow Collect and Provide? on page 1141.

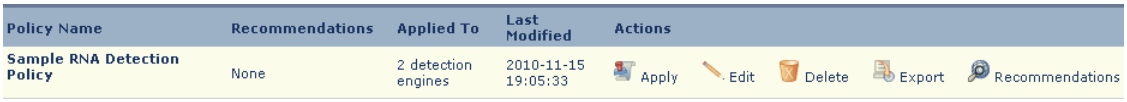
Keysight, EX-1004
Page 1153

To configure NetFlow monitoring in an RNA detection policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > RNA > Detection Policy**.

   The Detection Policy page appears.



2. Depending on whether you want to add NetFlow monitoring to an existing RNA detection policy, or you want to configure NetFlow monitoring as part of a new policy, you have the following options:
   - To add NetFlow monitoring to an existing RNA detection policy, click **Edit** next to the policy you want to modify.
   - To configure NetFlow monitoring as part of a new RNA detection policy, click **Create Policy**.

   In either case, the Create Policy page appears.

3. Configure the non-NetFlow-specific portions of the RNA detection policy. This includes:
   - providing basic policy information, such as the policy name and description, and configuring the non-NetFlow-specific detection policy options; see Configuring RNA Detection Policy Settings on page 902
   - optionally, specifying the networks you want to monitor directly with 3D Sensors with RNA; see Specifying Networks to Monitor with 3D Sensors on page 905
   - optionally, excluding traffic that travels from or to specific ports from monitoring by RNA; see Specifying Ports to Exclude on page 907

     Note that the port exclusion feature only affects data collected by 3D Sensors with RNA. You cannot configure NetFlow-enabled devices to exclude ports from monitoring.

4. Optionally, enable the NetFlow-specific detection policy options.
   - Select **Generate Hosts from NetFlow Data** if you want RNA to add hosts to the network map based on NetFlow data.
   - Select **Generate Services from NetFlow Data** if you want RNA to add services to the network map based on NetFlow data.

5. Begin specifying the networks you want to monitor with NetFlow-enabled devices by clicking **Add Network** under NetFlow Networks to Monitor.

6. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (in CIDR notation) that represents the network segment you want to monitor or exclude from monitoring.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

Keysight, EX-1004
Page 1154

7. To exclude the network from monitoring, select **Exclude**.
8. Under **NetFlow Device**, select the NetFlow-enabled device that you want to use to monitor the network you specified.

   The NetFlow-enabled devices listed are those you specified in the system settings.
9. Under **Reporting Detection Engine**, select the RNA detection engine that you want to use to analyze the data exported by the NetFlow-enabled device you chose.

   You must make sure that the 3D Sensor where the detection engine resides can see the data exported by the NetFlow-enabled device.
10. To add additional networks, repeat steps 5 through 9.

TIP! To remove a network, click **Delete** next to the network you want to remove.

---

11. Click **Save Policy**.

    The RNA detection policy is saved. You must apply the policy to start monitoring the hosts in the networks you specified. See Applying an RNA Detection Policy on page 919 for more information.

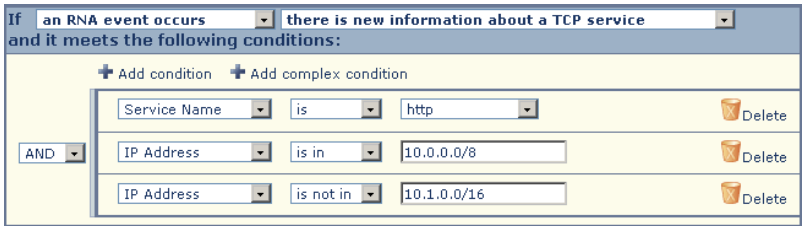
Keysight, EX-1004
Page 1155

Chapter 27
Using RNA as a Compliance Tool

A *compliance white list* (or *white list*) is a set of criteria that allows you to specify the operating systems, client applications, services, 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 the HTTP service, 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 the HTTP service.

Note that you could create a compliance rule that performs this function by configuring the rule so that it triggers when:

- RNA discovers new information about a service
- the service name is http
- the IP address of the host involved in the RNA event is not in your web farm



However, compliance 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. Compliance rules are also wider in scope, allowing you to generate a compliance event when any specific flow, intrusion, RNA, RUA, or host input event occurs and the event meets any criteria that you

Keysight, EX-1004
Page 1156

specify. On the other hand, white lists are specifically meant to help you evaluate the operating systems, client applications, services, 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, client applications, services, 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 compliance policy, when RNA detects that a host is violating the white list, the Defense Center logs a white list event—which is a special kind of compliance event— to the database. Further, you can configure the Defense Center to trigger responses (remediations and alerts) automatically when it detects a white list violation.

IMPORTANT! Although you can configure the RNA detection policy to add hosts and services to the network map based on data exported by NetFlow-enabled devices, the available information about these hosts and services 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 What Information Does RNA Provide? on page 880.

Because the Defense Center 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 the HTTP service in violation of your policy, and take appropriate action.

Then, using the policy and response feature, you can configure the Defense Center to alert you whenever a host that is not in your web farm starts running the HTTP service.

In addition, the Sourcefire 3D 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 Status Dashboard to monitor recent system-wide compliance activity, including white list events and summary views of the overall white list compliance of your network.

IMPORTANT! You must use a Defense Center to create compliance white lists and to view white list events. You cannot create white lists on a Master Defense Center, but you can view white list events forwarded from managed Defense Centers.

Keysight, EX-1004
Page 1157

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 1158
- Creating Compliance White Lists on page 1167
- Managing Compliance White Lists on page 1188
- Working with Shared Host Profiles on page 1189
- Working with White List Events on page 1197
- Working with White List Violations on page 1204

In addition, see the following chapters and sections for more information:

- Creating Compliance Policies on page 1262 explains how to create and configure compliance policies that include compliance white lists, and explains how to assign responses and priorities to the white lists.
- Using Host Profiles on page 944 explains how to use a host's profile to determine whether it is violating any white lists.
- Using Dashboards on page 61 explains how to obtain an at-a-glance view of your current system status, including white list compliance activity.

Understanding Compliance White Lists

Platform: DC

Component: RNA

A *compliance white list* is a set of criteria that specify the operating systems, client applications, services, 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 service 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, client applications, services, and protocols are allowed to run on the targets.

After you create a white list and add it to an active compliance policy, the Defense Center 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 Defense Center generates a *white list event* when it detects that a valid target is violating the white list.

Keysight, EX-1004
Page 1158

For more information, see the following sections:

- Understanding White List Targets on page 1159 explains how white lists only target the hosts that you specify.
- Understanding White List Host Profiles on page 1160 explains the different kinds of profiles that describe which client applications, services, and protocols are allowed to run on your network.
- Understanding White List Evaluations on page 1164 explains how RNA 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 1165 explains how RNA detects and the Defense Center notifies you of white list violations.

Understanding White List Targets

Platform: DC

Component: RNA

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 IP addresses. 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 987 and Assigning a Host Criticality Value to a Host on page 986.
- 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 10.5.x.x network and one that excludes the 10.5.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.

Keysight, EX-1004
Page 1159

The targets for your white list are listed on left of the Create White List page.



Note that the default white list uses a target of 0.0.0.0/0, which represents 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 1171.

Understanding White List Host Profiles

Platform: DC

Component: RNA

Once 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, client applications, services, 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 Accessing Compliance White List Host Profiles table 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

Keysight, EX-1004
Page 1160

For more information, see the following sections:

- Understanding the Global Host Profile on page 1161
- Understanding Host Profiles for Specific Operating Systems on page 1161
- Understanding Shared Host Profiles on page 1163

Understanding the Global Host Profile

Platform: DC

Component: RNA

Every white list contains a global host profile, which specifies the services, client 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 NetBIOS services, you can configure the global host profile to allow the services regardless of the operating system on which they were 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 1175.

Understanding Host Profiles for Specific Operating Systems

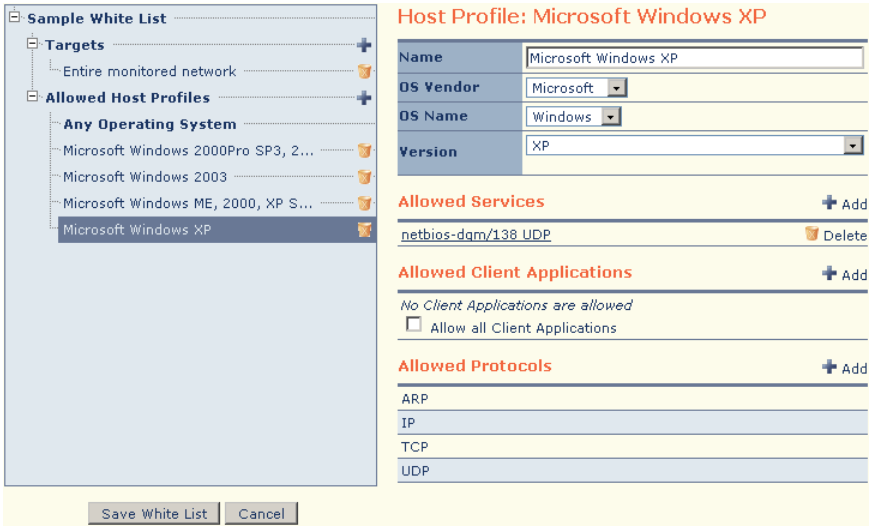
Platform: DC

Component: RNA

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

Keysight, EX-1004
Page 1161

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 XP.

After you create a host profile for a particular operating system, you can specify the services, client applications, and protocols that are allowed to run on target hosts running that operating system. For example, you could allow the SSH service to run on Linux hosts on port 22. You could also restrict the particular vendor and version of the service 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 RNA has not yet gathered enough information to identify their operating systems. *Unknown* hosts are hosts whose traffic has been analyzed by RNA, but whose operating systems do not match any of RNA's known fingerprints.

For more information, see Creating Host Profiles for Specific Operating Systems on page 1176.

Keysight, EX-1004
Page 1162

Understanding Shared Host Profiles

Platform: DC

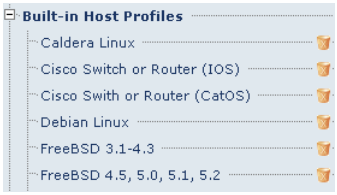
Component: RNA

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, client applications, services, 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 ( ).



Built-in host profiles use built-in services, protocols, and client applications. You can use these elements as-is in both the default white list and in any custom

Keysight, EX-1004
Page 1163

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 service, protocol, or client application, it affects every white list that uses it.

For more information on shared host profiles, Working with Shared Host Profiles on page 1189.

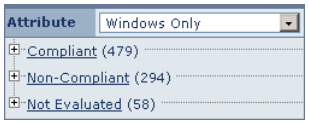
Understanding White List Evaluations

Platform: DC

Component: RNA

After you create white list host profiles and save the white list, you can add the white list to a compliance policy, just as you would a compliance rule. For more information, see Configuring Compliance Policies and Rules on page 1212.

Once you activate the compliance policy, the Defense Center 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 Defense Center 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 Defense Center completes its processing, more hosts move from Not Evaluated to either Compliant or Non-Compliant. The Defense Center can evaluate approximately 100 hosts per second.

Additionally, a host may be marked as Not Evaluated if the Defense Center has insufficient information to determine whether the host is in compliance. For example, this may occur if RNA has detected a new host but has not yet gathered

Keysight, EX-1004
Page 1164

relevant information on the operating system, client applications, services, 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 935.

Understanding White List Violations

Platform: DC

Component: RNA

After the initial white list evaluation, the Defense Center 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 compliance event, and are logged to the Defense Center compliance 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 1197.

White list violations occur when RNA generates an event that indicates that a host is out of compliance. Similarly, RNA events can indicate that a previously non-compliant host is now compliant, although the Defense Center does **not** generate a white list event when this occurs.

The following RNA events can change the compliance of a host:

- RNA detects a change in a host's operating system
- RNA detects an identity conflict for a host's operating system or a service on the host
- RNA detects a new TCP service port (for example, a port used by SMTP or web services) active on a host, or a new UDP service running on a host
- RNA detects a change in a discovered TCP or UDP service running on a host, for example, a version change due to an upgrade
- RNA detects a new client application running on a host
- RNA drops a client application from its database due to inactivity
- RNA 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
- RNA 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 application, protocol, or service to a host

Keysight, EX-1004
Page 1165

- delete a client application, protocol, or service 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 RNA detects that the host is now running Mac OS X, the Defense Center 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
- RNA 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 the AIM service, and you then delete the AIM service from the services network map or from a services event view, hosts running the AIM service become compliant. However, if RNA detects the service again, the Defense Center generates a white list event and the hosts become non-compliant.

Note that if RNA 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, RNA detects that port 21, using the TCP protocol, has become active on one of the white list targets, but RNA is unable to determine whether the traffic is FTP. In this scenario, the white list does not trigger until either RNA 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 Defense Center 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 client applications, services, and protocols to be rediscovered, which can trigger white list events. For more information, see Purging the RNA and RUA Databases on page 2014.

Finally, you can configure the Defense Center to trigger responses automatically when it detects a white list violation. Responses include remediations (such as blocking a host at the firewall or router), alerts (email, SNMP, and syslog alerts), or

Keysight, EX-1004
Page 1166

combination of alerts and remediations. For more information, see Configuring Responses for Compliance Policies on page 1281.

Creating Compliance White Lists

Platform: DC

Component: RNA

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 RNA has detected on the network segment. By default, these host profiles allow all of the client applications, services and protocols that RNA 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 IP addresses. 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, client applications, services, 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 your 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 compliance policy. The Defense Center 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 1158.

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 1026.

To create a compliance white list:

Access: P&R Admin/Admin

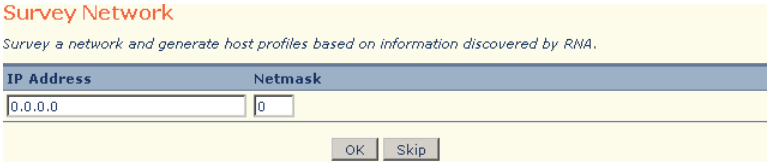
1. Select **Policy & Response > Compliance > White List**.

   The White List page appears.

Keysight, EX-1004
Page 1167

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 1168.
   - To create a white list without surveying your network, continue with the next step.
4. 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 1171.
5. 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 1175.
6. Click **Save White List** to save your white list.

   The white list is saved. You can now add it to an active compliance 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 Compliance Policies on page 1262.

Surveying Your Network

Platform: DC

Component: RNA

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 RNA database about the services, client applications, and protocols running on the different detected operating systems. Then, the Defense Center creates one host profile within the white list for each detected operating system. By default, these host profiles allow all of the detected client applications, services and protocols that RNA 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

Keysight, EX-1004
Page 1168

or delete the host profiles that the survey created to suit your needs, and also you can 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 client applications, services, 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 compliance policy, and the survey changes either your targets or host profiles, when 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.

To begin creating a compliance white list by surveying your network:

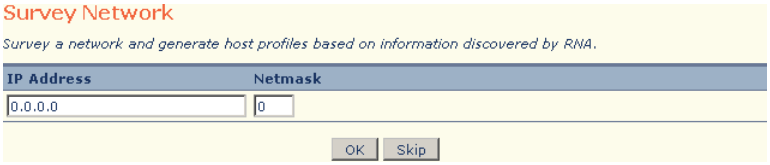
Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > White List**.

   The White List page appears.

2. Click **New White List**.

   The Survey Network page appears.



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 CIDR notation) that represents the IP addresses you want to survey.

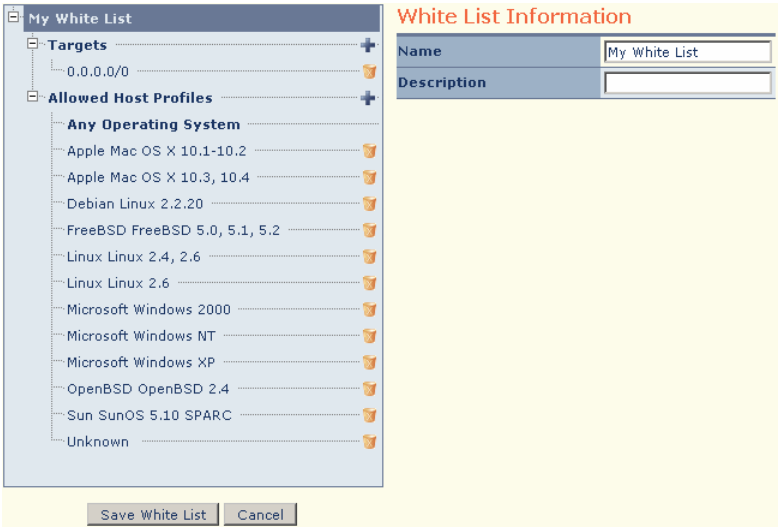
   You should make sure that you specify a network that you configured RNA to monitor in What is an RNA Detection Policy? on page 889. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

TIP! To survey the entire monitored network use the default value of `0.0.0.0/0`.

Keysight, EX-1004
Page 1169

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 client applications, services, 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.

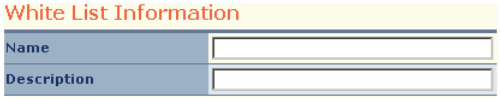
7. Continue with Providing Basic White List Information.

Providing Basic White List Information

Platform: DC

Component: RNA

You must give each white list a name, and, optionally, short description.



Keysight, EX-1004
Page 1170

To provide basic white list information:

Access: P&R Admin/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. Continue with Configuring Compliance White List Targets.

Configuring Compliance White List Targets

Platform: DC

Component: RNA

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 IP addresses. 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 1159.

When you are finished creating compliance white list targets, continue with Configuring Compliance White List Host Profiles on page 1175.

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 create, modify, and delete targets, see:

- Creating a Target on page 1171
- Modifying Existing Targets on page 1174
- Deleting Existing Targets on page 1174

Creating a Target

Platform: DC

Component: RNA

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.

Keysight, EX-1004
Page 1171

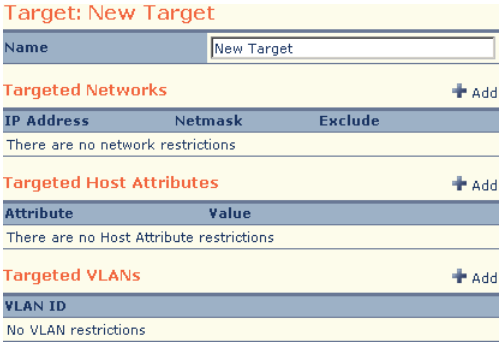
Note that if you add a target to a white list that is used by an active compliance policy, once you save the white list, the new target hosts are evaluated for compliance. However, this evaluation does not generate white list events.

To create a compliance white list target:

Access: P&R Admin/Admin

1. On the Create White List Page, next to **Targets**, click the Add icon ( ).

   The settings for the new target appear.



> **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 1168. 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.



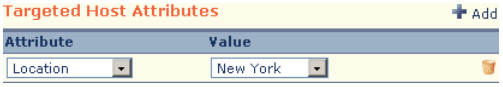
Keysight, EX-1004
Page 1172

4. In the **IP Address** and **Netmask** fields, enter the IP address and network mask (as in CIDR notation) that represents the IP addresses you want to target or exclude from targeting.

   You should make sure that you specify a network that you configured RNA to monitor in What is an RNA Detection Policy? on page 889. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

   TIP! To target the entire monitored network, use `0.0.0.0/0`.

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.



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.



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 0-4095 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.

Keysight, EX-1004
Page 1173

Modifying Existing Targets

Platform: DC

Component: RNA

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 compliance 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 Defense Center changes the white list host attribute of previously valid targets to Not Evaluated.

To modify an existing target:

Access: P&R Admin/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 Creating a Target on page 1171.

Deleting Existing Targets

Platform: DC

Component: RNA

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 compliance policy, the Defense Center changes the white list host attribute of previously valid targets to Not Evaluated.

To delete a white list target:

Access: P&R Admin/Admin

1. Next to the target you want to delete, click the Delete icon ( ).

2. At the prompt, confirm that you want to delete the target.

   The target is deleted.

Keysight, EX-1004
Page 1174

Configuring Compliance White List Host Profiles

Platform: DC

Component: RNA

Host profiles in a compliance white list specify which operating systems, client applications, services, 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 services, client 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 services, client 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 1160.

When you are finished creating compliance white list host profiles, you can add the white list to an active compliance 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 information on how to create, modify, and delete compliance white list host profiles, see:

- Configuring the Global Host Profile on page 1175
- Creating Host Profiles for Specific Operating Systems on page 1176
- Adding a Shared Host Profile to a Compliance White List on page 1184
- Modifying Existing Host Profiles on page 1184
- Deleting Existing Host Profiles on page 1187

Configuring the Global Host Profile

Platform: DC

Component: RNA

Every white list contains a global host profile, which specifies the services, client 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 1161.

Keysight, EX-1004
Page 1175

To configure the global host profile:

Access: P&R Admin/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 services you want to allow, follow the directions in Adding a Service to a Host Profile on page 1178.
3. To specify the client applications you want to allow, follow the directions in Adding a Client Application to a Host Profile on page 1180.
4. To specify the protocols you want to allow, follow the directions in Adding a Protocol to a Host Profile on page 1182.

   Note that ARP, IP, TCP, and UDP are always allowed.

Creating Host Profiles for Specific Operating Systems

Platform: DC

Component: RNA

Host profiles for specific operating systems indicate not only which operating systems are allowed to run on your network, but also the services, client 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 1161.

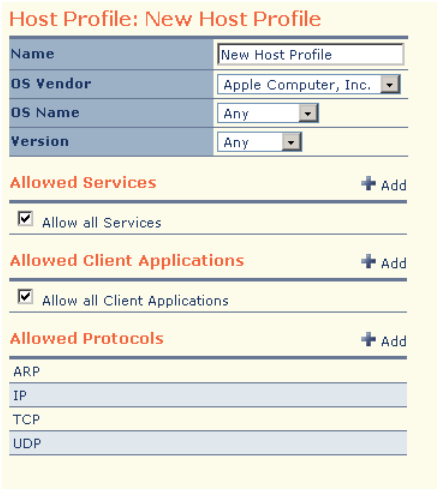
Keysight, EX-1004
Page 1176

To create a new compliance white list host profile for a specific operating system:

Access: P&R Admin/Admin

1. Next to Allowed Host Profiles, click the Add icon ( ).

   The settings for the new host profile appear.



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 services you want to allow. You have three options:
   - To allow all services, leave the **Allow all Services** check box selected.
   - To allow no services, clear the **Allow all Services** check box.
   - To allow specific services, follow the directions in Adding a Service to a Host Profile on page 1178.
5. Specify the client applications you want to allow. You have three options:
   - To allow all applications, leave the **Allow all Client Applications** check box selected.
   - To allow no applications, clear the **Allow all Client Applications** check box.
   - To allow specific applications, follow the directions in Adding a Client Application to a Host Profile on page 1180.
6. 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 1182. Note that ARP, IP, TCP, and UDP are always allowed.

Keysight, EX-1004
Page 1177

Adding a Service to a Host Profile

Platform: DC

Component: RNA

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 services to run on specific operating systems. You can also configure a white list to allow certain services to run on any valid target; these are called globally allowed services.

For any allowed service, you can either specify the type of service that you want to allow—aim, ftp, and ssh are examples of service types—or you can allow a custom service by specifying a service type of `any`. You must also specify the protocol the allowed service uses (TCP or UDP). You can allow the service on any port, or restrict it to a port that you specify.

Optionally, you can require that the service have a specific vendor or version. For example, you could allow the SSH service to run on Linux hosts on port 22. You could also restrict the particular vendor and version of the service to OpenSSH 4.2.

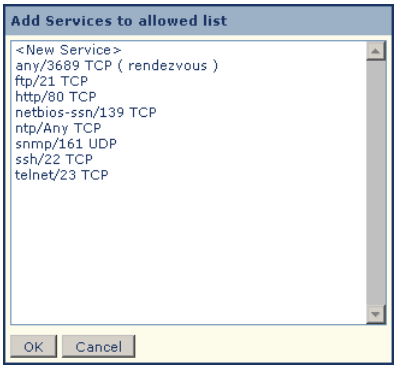
To add a service to a compliance white list host profile:

Access: P&R Admin/Admin

1. While you are creating or modifying a white list host profile, click **Add** next to Allowed Services (or next to Globally Allowed Services if you are modifying the `Any Operating System` host profile).

   A pop-up window appears. The services listed are:

   - services that you created within the white list
   - services that existed in the network map when you surveyed your networks as described in Surveying Your Network on page 1168
   - services that are used by other host profiles in the white list, which may include built-in services created by the VRT for use in the default white list

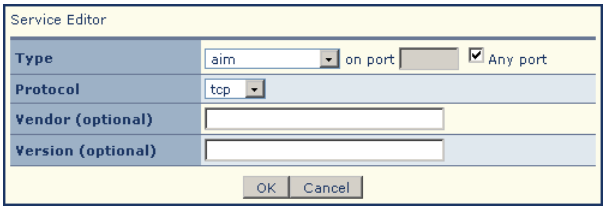


Keysight, EX-1004
Page 1178

2. You have two options:
   - To add a service already in the list, select it and click **OK**. Use Ctrl or Shift while clicking to select multiple services. You can also click and drag to select multiple adjacent services.

     The service is added. Note that if you added a built-in service, its name appears in italics. You can skip the rest of the procedure, or optionally, to change any of the service's values (such as the port or protocol), click the service you just added to display the service editor.
   - To add a new service, select **<New Service>** and click **OK**.

     The service editor appears.



3. From the **Type** drop-down list, select the service type. For custom services, select **any**.
4. Specify the service port. You have two options:
   - To allow the service to run on any port, check the **Any port** check box.
   - To allow the service to run only on a specific port, type the port number in the **port** field.
5. From the **Protocol** drop-down list, select the service protocol: **TCP** or **UDP**.
6. Optionally, in the **Vendor** and **Version** fields, specify a vendor and version for the service.

   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 a table view of services or in the services network map.
7. Click **OK**.

   The service is added. Note that you must save the white list for your changes to take effect.

   If you added a service to a white list that is used by an active compliance 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.

Keysight, EX-1004
Page 1179

Adding a Client Application to a Host Profile

Platform: DC

Component: RNA

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 applications to run on any valid target; these are called globally allowed client applications.

For any allowed client application, you must specify the name of the application as well as its type; **web browser** and **instant messenger** are examples of application types. Optionally, you can require that the application be a specific version. For example, you could allow only Microsoft Internet Explorer 6.0 to run on Microsoft Windows hosts.

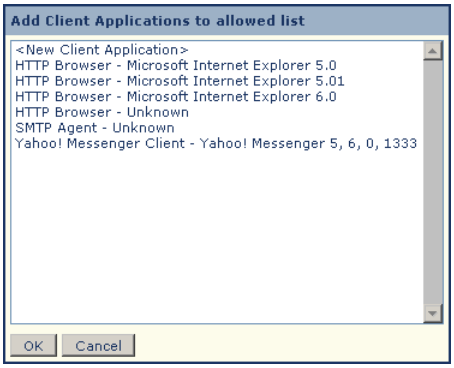
To add a client application to a compliance white list host profile:

Access: P&R Admin/Admin

1. While you are creating or modifying a white list host profile, click **Add** next to Allowed Client Applications (or next to Globally Allowed Client Applications if you are modifying the `Any Operating System` host profile).

   A pop-up window appears. The client applications listed are:

   - client applications that you created within the white list
   - client applications that were running on hosts in the network map when you surveyed your networks as described in Surveying Your Network on page 1168
   - client applications that are used by other host profiles in the white list, which may include built-in client applications created by the VRT for use in the default white list

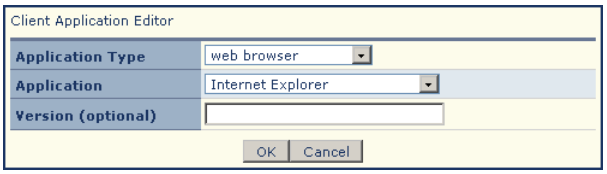


Keysight, EX-1004
Page 1180

2. You have two options:
   - To add a client application already in the list, select it and click **OK**. Use Ctrl or Shift while clicking to select multiple applications. You can also click and drag to select multiple adjacent applications.

     The client application is added. Note that if you added a built-in client application, its name appears in italics. You can skip the rest of the procedure, or optionally, to change any of the application's values (such as its version), click the application you just added to display the application editor.
   - To add a new client application, select **<New Client Application>** and click **OK**.

     The client application editor appears.



3. From the **Application Type** and **Application** drop-down lists, select the client application type and name.
4. Optionally, in the **Version** fields, specify a version for the client application.

   If you do not specify a version, the white list allows all versions as long as the type and name match. Note that if you restrict the version, you must specify it exactly as it would appear in a table view of client applications.
5. Click **OK**.

   The client application is added. Note that you must save the white list for your changes to take effect.

   If you added an application to a white list that is used by an active compliance 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.

Keysight, EX-1004
Page 1181

Adding a Protocol to a Host Profile

Platform: DC

Component: RNA

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.

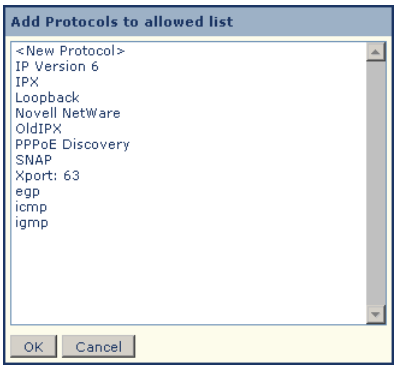
To add a protocol to a compliance white list host profile:

Access: P&R Admin/Admin

1. While you are creating or modifying a white list host profile, click **Add** 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 1168
   - 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

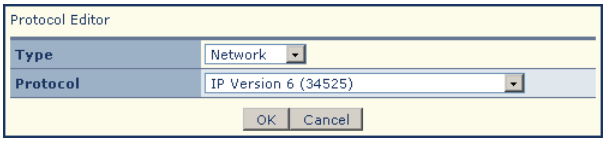


Keysight, EX-1004
Page 1182

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 compliance 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.

Keysight, EX-1004
Page 1183

Adding a Shared Host Profile to a Compliance White List

Platform: DC

Component: RNA

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.

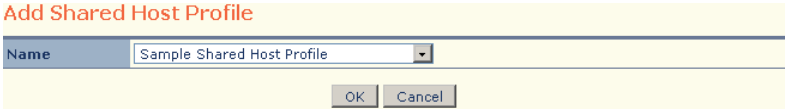
You can add any of the built-in shared host profiles to your compliance white lists, or you can add shared host profiles that you created. For more information, see Understanding Shared Host Profiles on page 1163 and Creating Shared Host Profiles on page 1190.

To add a shared host profile to a compliance white list:

Access: P&R Admin/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 1184.

Modifying Existing Host Profiles

Platform: DC

Component: RNA

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 compliance 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

Keysight, EX-1004
Page 1184

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.

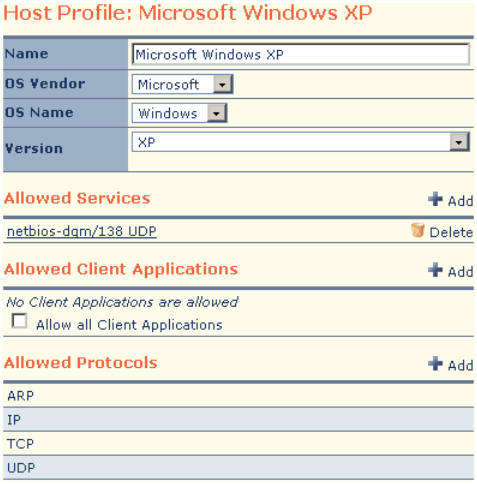
TIP! As with other shared host profiles, you can edit the built-in host profiles used by the default white list. You can also reset them to their factory defaults. For more information, see Resetting Built-In Host Profiles to Their Factory Defaults on page 1195.

To modify an existing host profile:

Access: P&R Admin/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.



Keysight, EX-1004
Page 1185

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 1192.
   - 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 table below.

To...	You can...
rename the host profile	type a new name in the **Name** field.
change the operating system for the host profile	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.
add a service	follow the directions in Adding a Service to a Host Profile on page 1178.
add a client application	follow the directions in Adding a Client Application to a Host Profile on page 1180.
add a protocol	follow the directions in Adding a Protocol to a Host Profile on page 1182.
allow all services	under Allowed Services, select the **Allow all Services** check box. Note that the check box does not appear until you delete any services you have previously allowed.
allow all client applications	under Allowed Client Applications, select the **Allow all Client Applications** check box. Note that the check box does not appear until you delete any services you have previously allowed.

Keysight, EX-1004
Page 1186

To...	You can...
modify a service, client application, or protocol	click the element you want to modify. For more information on the properties you can change, see: • Adding a Service to a Host Profile on page 1178 • Adding a Client Application to a Host Profile on page 1180 • Adding a Protocol to a Host Profile on page 1182 **IMPORTANT!** The changes you make to a service, client application, or protocol are reflected in every host profile that uses that element.
delete a service, client application, or protocol	next to the element you want to delete, click **Delete**.
survey your network	click **Survey Network**. Surveying your network can add additional allowed client applications, services, 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 1168.

Deleting Existing Host Profiles

Platform: DC

Component: RNA

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 compliance 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: P&R Admin/Admin

1. On the Create White List page, next to the host profile you want to delete, click the Delete icon ( ).
2. At the prompt, confirm that you want to delete the host profile.

   The host profile is deleted.

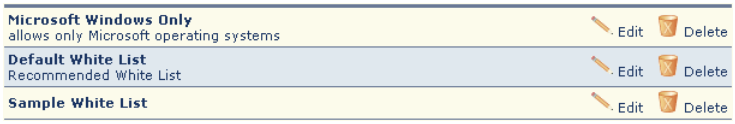
Keysight, EX-1004
Page 1187

Managing Compliance White Lists

Platform: DC

Component: RNA

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.



For more information, see:

- Creating Compliance White Lists on page 1167
- Modifying a Compliance White List on page 1188
- Deleting a Compliance White List on page 1189
- Working with Shared Host Profiles on page 1189

Modifying a Compliance White List

Platform: DC

Component: RNA

When you modify a compliance white list that is included in an active compliance policy, the Defense Center re-evaluates the target hosts. Note that the Defense Center 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 compliance policy and a previously compliant host becomes non-compliant with the updated white list.

To modify an existing compliance white list:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > White List**.
   The White List page appears.
2. Next to the white list you want to modify, click **Edit**.
   The Create White List page appears.
3. Make modifications as necessary and click **Save White List**.
   The white list is updated.

Keysight, EX-1004
Page 1188

Deleting a Compliance White List

Platform: DC

Component: RNA

You cannot delete a compliance white list that you are using in one or more compliance 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 Modifying a Compliance Policy on page 1270.

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: P&R Admin/Admin

1. Select **Policy & Response > Compliance > White List**.

   The White List page appears.

2. Next to the white list you want to delete, click **Delete**.

   The white list is deleted.

Working with Shared Host Profiles

Platform: DC

Component: RNA

Shared host profiles specify which operating systems, client applications, services, 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 1163.

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 services, protocols, or client applications, you can reset them to their factory defaults. For more information, see:

- Creating Shared Host Profiles on page 1190
- Modifying a Shared Host Profile on page 1192
- Deleting a Shared Host Profile on page 1195
- Resetting Built-In Host Profiles to Their Factory Defaults on page 1195

Keysight, EX-1004
Page 1189

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 1184.

Creating Shared Host Profiles

Platform: DC

Component: RNA

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 977.

To create a shared host profile:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > 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 RNA has accumulated about your network. This saves you from manually creating and configuring multiple shared host profiles.

   - To survey your network, click **Survey Network**. For more information, see Surveying Your Network on page 1168.

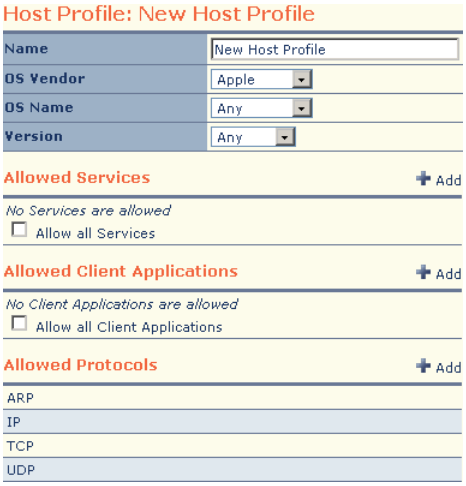
     RNA creates one or more baseline shared host profiles. You can edit or delete these shared host profiles as described in Modifying a Shared Host Profile on page 1192 and Deleting a Shared Host Profile on page 1195. To add any other shared host profiles you might need, continue with the next step.

   - To skip surveying you network, continue with the next step.

Keysight, EX-1004
Page 1190

4. Next to Shared Host Profiles, click the Add icon ( ).

   The settings for the new shared host profile appear.



5. In the **Name** field, type a descriptive name for the shared host profile.
6. 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.
7. Specify the services you want to allow. You have three options:
   - To allow all services, select the **Allow all Services** check box.
   - To allow no services, leave the **Allow all Services** check box cleared.
   - To allow specific services, next to Allowed Services, follow the directions in Adding a Service to a Host Profile on page 1178.
8. Specify the client applications you want to allow. You have three options:
   - To allow all applications, select the **Allow all Client Applications** check box.
   - To allow no applications, leave the **Allow all Client Applications** check box cleared.
   - To allow specific applications, follow the directions in Adding a Client Application to a Host Profile on page 1180.
9. 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 1182. Note that ARP, IP, TCP, and UDP are always allowed.
10. 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.

Keysight, EX-1004
Page 1191

Modifying a Shared Host Profile

Platform: DC

Component: RNA

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 compliance 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: P&R Admin/Admin

1. Select **Policy & Response > Compliance > White List**.

   The White List page appears.

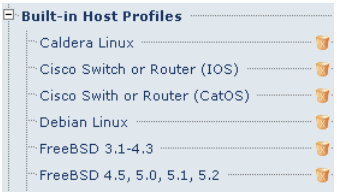
2. Click **Edit Shared Profiles**.

   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.

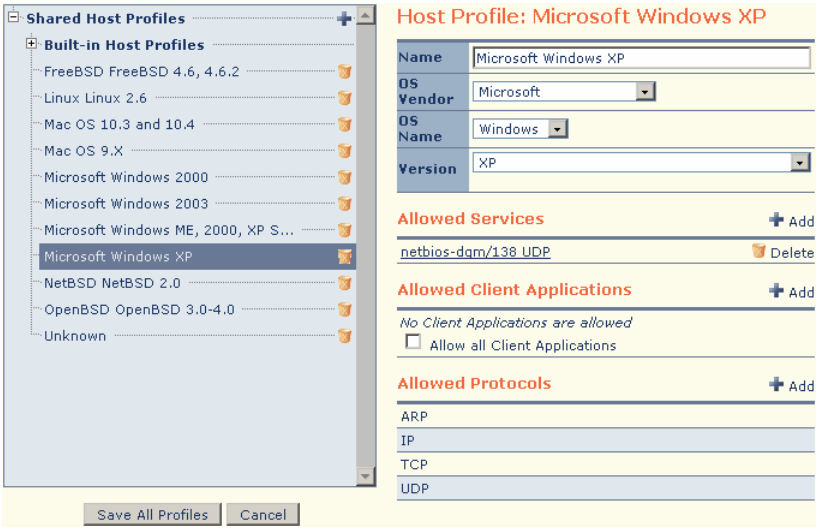


   - If no, continue with the next step.

Keysight, EX-1004
Page 1192

4. Click the name of the shared host profile you want to modify.

   The host profile appears.



5. 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 a service	follow the directions in Adding a Service to a Host Profile on page 1178.
add a client application	follow the directions in Adding a Client Application to a Host Profile on page 1180.
add a protocol	follow the directions in Adding a Protocol to a Host Profile on page 1182.
allow all services	under Allowed Services, select the **Allow all Services** check box. Note that the check box does not appear until you delete any services you have previously allowed.

Keysight, EX-1004
Page 1193

To...	You can...
allow all client applications	under Allowed Client Applications, select the **Allow all Client Applications** check box. Note that the check box does not appear until you delete any services you have previously allowed.
modify a service, client application, or protocol	click the element you want to modify. For more information on the properties you can change, see: • Adding a Service to a Host Profile on page 1178 • Adding a Client Application to a Host Profile on page 1180 • Adding a Protocol to a Host Profile on page 1182 **IMPORTANT!** The changes you make to a service, client application, or protocol are reflected in every host profile that uses that element.
delete a service, client application, or protocol	next to the element you want to delete, click **Delete**.
survey your network	click **Survey Network**. Surveying your network can add additional allowed client applications, services, 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 1168.

6. Click **Save all Profiles** to save your changes.

   The shared host profile is deleted and removed from all compliance white lists that use it.

Keysight, EX-1004
Page 1194

Deleting a Shared Host Profile

Platform: DC

Component: RNA

If the shared host profile you delete belongs to one or more white lists used in an active compliance 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 1195.

To delete a shared host profile:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > 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

Platform: DC

Component: RNA

The default white list uses a special category of shared host profiles, called *built-in host profiles*. Built-in host profiles use built-in services, protocols, and client applications. 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.

Keysight, EX-1004
Page 1195

If you make changes to the built-in profiles, services, protocols, or client applications 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, services, protocols, and client applications that you modified are reset to their factory defaults.
- **All** built-in host profiles, services, protocols, and client applications that you deleted are restored.
- **All** white lists (including the default white list) that are used by active compliance policies and that used any of the reset built-in host profiles, services, protocols, or client applications 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, services, protocols, and client applications:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > 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, services, protocols, and client applications are reset to factory defaults. Any white list that is used by an active compliance policy and that used any of the reset built-in host profiles, services, protocols, or client applications is re-evaluated.

Keysight, EX-1004
Page 1196

Working with White List Events

Platform: DC/MDC

Component: RNA

When RNA generates an RNA event that indicates that a host is out of compliance with a white list that is included in an activated compliance policy, a white list event is generated. White list events are a special kind of compliance event, and are logged to the compliance 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 1747. Note that white list events are stored in the compliance event database.

For more information, see the following sections:

- Viewing White List Events on page 1197
- Understanding the White List Events Table on page 1200
- Searching for Compliance White List Events on page 1201

Viewing White List Events

Platform: DC/MDC

Component: RNA

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 1618.

Keysight, EX-1004
Page 1197

The Compliance White List Event Actions table below 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 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 1495.
sort and constrain events on the current workflow page	find more information in Sorting Drill-down Workflow Pages on page 1612.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1613.
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 1593.
learn more about the columns that appear	find more information in Understanding the White List Events Table on page 1200.
modify the time and date range for displayed events	find more information in see Setting Event Time Constraints on page 1597.
drill down to the next page in the workflow, constraining on a specific value	use one of the following methods: • 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.

Keysight, EX-1004
Page 1198

Compliance White List Event Actions (Continued)

To...	You can...
delete white list events from the system	use one of the following methods: • 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. See Purging the RNA and RUA Databases on page 2014 for information on deleting all RNA events from the database.
navigate to other event views to view associated events	find more information in Navigating between Workflows on page 1614.
temporarily use a different workflow	click **Workflows**. For more information, see Selecting Workflows on page 1590.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for white list events	click **Search**. For more information, see Searching for Compliance White List Events on page 1201.

To view compliance white list events:

Access: Any Analyst/Admin

▶ Select **Analysis & Reporting > Compliance > White List Events**.

The first page of the default white list events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

TIP! If you are using a custom workflow that does not include the table view of white list events, from the **Workflows** menu on the toolbar, select **White List Events**.

Keysight, EX-1004
Page 1199

Understanding the White List Events Table

Platform: DC/MDC

Component: RNA

You can use the compliance policy feature to build *compliance policies* that let your Defense Center respond in real time to threats on your network. Compliance policies describe the type of activity that constitutes a policy violation, which include compliance white list violations. For more information on compliance policies, see Configuring Compliance Policies and Rules on page 1212.

When a compliance white list is violated, the Defense Center generates a white list event. The fields in the white list events table are described in the Compliance White List Event Fields 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 a service white list violation (a violation that occurred as a result of a non-compliant service). For other types of white list violations, this field is blank.

Description	A description of how the white list was violated. For example:
	`Client Application “AOL Instant Messenger” is not allowed.`
	
	Violations that involve a service indicate the service 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:
	`Service "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 compliance policy that was violated, that is, the compliance policy that includes the white list.
White List	The name of the white list.

Keysight, EX-1004
Page 1200

Compliance White List Event Fields (Continued)

Field	Description
Priority	The priority specified by the policy or white list that triggered the policy violation. For information on setting compliance rule and policy priorities, see Providing Basic Policy Information on page 1264 and Setting Rule and White List Priorities on page 1266.
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 Assigning a Host Criticality Value to a Host on page 986.
Detection Engine	The name of the detection engine that generated the event that triggered 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

Platform: DC/MDC

Component: RNA

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 search criteria you can use are described in the Compliance White List Event Search Criteria table.

Compliance White List Event Search Criteria

Field	Search Criteria Rules
Policy	Enter all or part of the name of a compliance policy to return all events caused by violations of white lists included in that policy.
White List	Enter all or part of the name of a white list to return all events caused by violations of that white list.

Keysight, EX-1004
Page 1201

Compliance White List Event Search Criteria (Continued)

Field	Search Criteria Rules
Description	Enter all or part of the white list event description. You can use an asterisk (*) as a wildcard character in this field. For example, enter `*ssh*` to return all white list events that involve the SSH service. For more information, see Using Wildcards and Symbols in Searches on page 1555.
Priority	Specify the priority of the white list event, which is determined either by the priority of the white list in a compliance policy or by the priority of the compliance policy itself. Note that the white list priority overrides the priority of its policy. Enter `none` for no priority. For information on setting compliance rule and policy priorities, see Providing Basic Policy Information on page 1264 and Setting Rule and White List Priorities on page 1266.
IP Address	Specify the IP address of a host that has become non-compliant with a white list. See Specifying IP Addresses in Searches on page 1557 for detailed information about how to enter IP addresses.
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 RNA event that triggered a service white list violation (a violation that occurred as a result of a non-compliant service). See Specifying Ports in Searches on page 1558 for information about the syntax for specifying ports.
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 Assigning a Host Criticality Value to a Host on page 986.
Detection Engine	Specify the name of the detection engine, detection engine group, sensor, or sensor group that generated the RNA event that indicated a host's noncompliance with a white list. For more information, see Specifying Detection Engine/ Appliance Names on page 1559.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

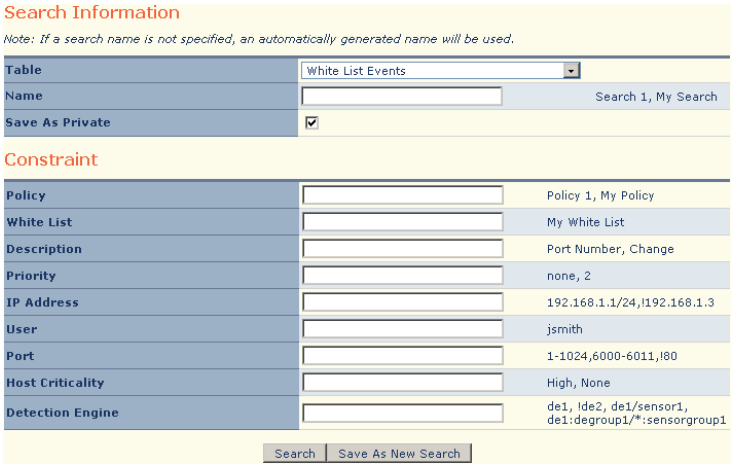
Keysight, EX-1004
Page 1202

To search for compliance white list events:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Searches > White List Events**.

   The White List search page appears.



   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.

   Search criteria are described in the Compliance White List Event Search Criteria table on page 1201. 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 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.

Keysight, EX-1004
Page 1203

5. 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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

Platform: DC

Component: RNA

The Defense Center keeps track of the ways in which hosts on your network violate the compliance white lists in active compliance policies. You can search and view these records.

For more information, see the following sections:

- Viewing White List Violations on page 1204
- Understanding the White List Violations Table on page 1207
- Searching for White List Violations on page 1208

Viewing White List Violations

Platform: DC

Component: RNA

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

Keysight, EX-1004
Page 1204

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 1618.

The Compliance White List Violations Actions table below 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 an IP address	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 1612.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1613.
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 1593.
learn more about the columns that appear	find more information in Understanding the White List Violations Table on page 1207.

Keysight, EX-1004
Page 1205

Compliance White List Violations Actions (Continued)

To...	You can...
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
navigate to other event views to view associated events	find more information in Navigating between Workflows on page 1614.
temporarily use a different workflow	click **Workflows**. For more information, see Selecting Workflows on page 1590.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for white list violations	click **Search**. For more information, see Searching for White List Violations on page 1208.

Keysight, EX-1004
Page 1206

To view compliance white list violations:

Access: Any Analyst/Admin

- Select **Analysis & Reporting > Compliance > White List Violations**.

  The first page of the default white list violations workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of white list violations, from the **Workflows** menu on the toolbar, select **White List Violations**.

Understanding the White List Violations Table

Platform: DC

Component: RNA

You can use the compliance policy feature to build *compliance policies* that let your Defense Center respond in real time to threats on your network. Compliance policies describe the type of activity that constitutes a policy violation, which include compliance white list violations. For more information on compliance policies, see Configuring Compliance Policies and Rules on page 1212.

When a compliance white list is violated, the Defense Center records the violation. 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 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: • operating system (**os**) • service (**service**) • client application (**client app**) • protocol (**protocol**)

Keysight, EX-1004
Page 1207

Compliance White List Violation Fields (Continued)

Field	Description
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 a service white list violation (a violation that occurred as a result of a non-compliant service). For other types of white list violations, this field is blank.
Protocol	The protocol, if any, associated with the event that triggered a service white list violation (a violation that occurred as a result of a non-compliant service). For other types of white list violations, this field is blank.
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

Platform: DC

Component: RNA

You can search for specific compliance white list violations. You may want to create searches customized for your network environment, then save them to

Keysight, EX-1004
Page 1208

re-use later. The search criteria you can use are described in the Compliance White List Violations Search Criteria table.

Compliance White List Violations Search Criteria

Field	Search Criteria Rules
Time	Specify the date and time that the white list was violated. See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time.
IP Address	Specify the IP address of a host that has become non-compliant with a white list. See Specifying IP Addresses in Searches on page 1557 for detailed information about how to enter IP addresses.
White List	Enter all or part of the name of a white list to return all violations of that white list.
Type	Enter the type of white list violation: • enter `os` (or `operating system`) to search for violations based on operating systems • enter `service` to search for violations based on services • enter `clientapp` (or `client app`) to search for violations based on client applications • enter `protocol` to search for violations based on protocols
Information	Enter all or part of the white list violation information. You can use an asterisk (*) as a wildcard character in this field. For example, enter `*Apple*` to return all violations that occurred because you set up a white list that disallows Apple, Inc. operating systems. For more information, see Using Wildcards and Symbols in Searches on page 1555.
Port	Specify the port, if any, associated with the RNA event that triggered a service white list violation (a violation that occurred as a result of a non-compliant service). See Specifying Ports in Searches on page 1558 for information about the syntax for specifying ports.
Protocol	Specify the protocol, if any, associated with the RNA event that triggered a service white list violation (a violation that occurred as a result of a non-compliant service).

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

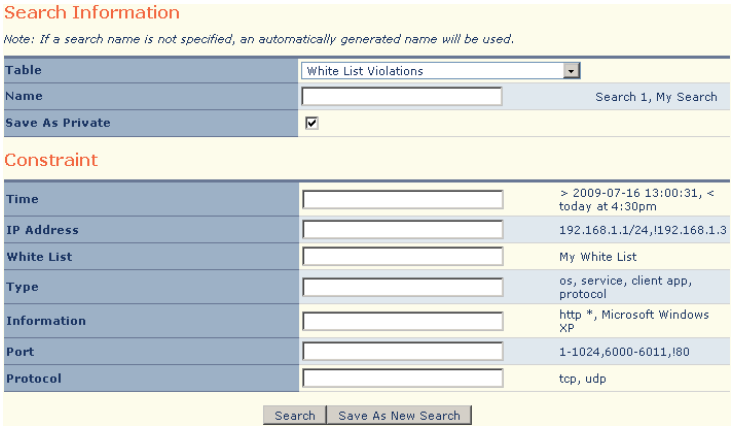
Keysight, EX-1004
Page 1209

To search for compliance white list events:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Searches > White List Violations**.

   The White List Violations search page appears.



   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.

   Search criteria are described in the Compliance White List Violations Search Criteria table on page 1209. 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 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.

Keysight, EX-1004
Page 1210

5. 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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.

Keysight, EX-1004
Page 1211

Chapter 28
Configuring Compliance Policies and Rules

You can use the powerful policy and response feature to build *compliance policies* that let the Defense Center respond in real time to threats to your network. Compliance policies are populated by *compliance rules* and *compliance white lists*.

A compliance rule triggers when a specific flow, intrusion, RNA, RUA, or host input event occurs and the event meets criteria that you specify. You can also configure a compliance rule to trigger 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, comprise criteria that specify which operating systems, client applications, services, and protocols are allowed to run on your network; when a host violates those criteria, the white list triggers.

A *policy violation* occurs when the activity on your network triggers either a compliance rule or white list. When this occurs, the Defense Center logs either a compliance event or a white list event to the database, depending on the type of policy violation. You can search, view, and delete compliance events and compliance white list events.

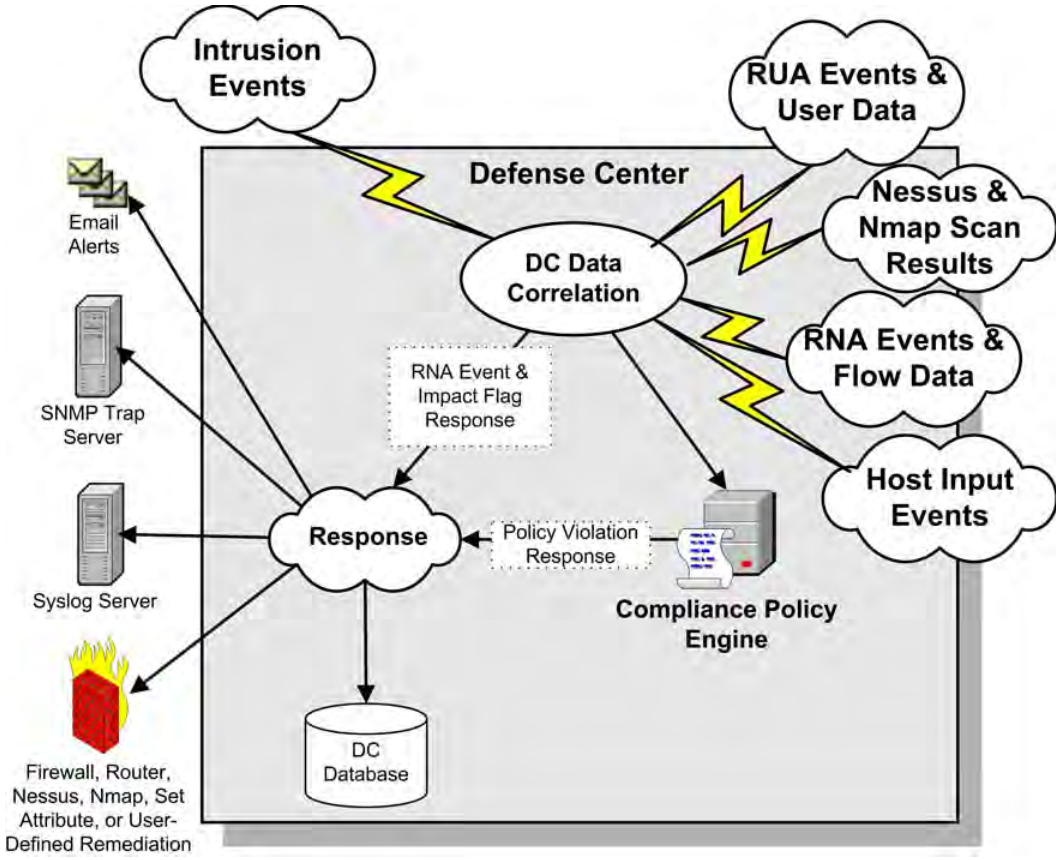
Within compliance policies, you can configure the Defense Center to trigger responses automatically when it detects a policy violation. Responses include remediations (such as blocking a host at the firewall or router when it violates a policy), alerts (email, SNMP, and syslog alerts), or combination of alerts and

Keysight, EX-1004
Page 1212

remediations. For more information on responses, see Configuring Responses for Compliance Policies on page 1281.

IMPORTANT! This chapter focuses on compliance policies, compliance rules, and compliance events. For more information on compliance white lists, see Using RNA as a Compliance Tool on page 1156.

The following graphic illustrates the event notification and the policy and response process:



IMPORTANT! You must use a Defense Center to create compliance policies and to view compliance events. On the Master Defense Center, you cannot create compliance policies, but you can view compliance events that were forwarded from managed Defense Centers.

Keysight, EX-1004
Page 1213

For more information on creating compliance rules and compliance policies, see the following sections:

- Creating Rules for Compliance Policies on page 1214
- Managing Rules for Compliance Policies on page 1260
- Creating Compliance Policies on page 1262
- Managing Compliance Policies on page 1269
- Working with Compliance Events on page 1271

Creating Rules for Compliance Policies

Platform: DC

Component: None

Before you create a compliance policy, you should create compliance rules or compliance white lists (or both) to populate it.

IMPORTANT! This section describes how to create compliance rules. For information on creating compliance white lists, see Creating Compliance White Lists on page 1167.

A *compliance rule* triggers (and generates a compliance event) when your network traffic meets criteria that you specify. You can build compliance rules to trigger when specific flow, intrusion, RNA, RUA, or host input events occur and the event meets those criteria. You can also build compliance rules to trigger when your network traffic deviates from your normal network traffic pattern as characterized in an existing traffic profile.

You can build complex rules with multiple conditions; you can even configure a compliance rule to trigger another rule, provided that both rules use the same type of network discovery event as their basis. In addition, you can constrain some compliance rules with information from the host profile of a host involved in the event. This constraint is called a *host profile qualification*.

You can further constrain a compliance rule so that once its initial criteria are met, RNA begins tracking certain flows. Then, a compliance event is generated only if the tracked flows meet additional criteria. This configuration is called a *flow tracker*.

IMPORTANT! You **cannot** add a host profile qualification or a flow tracker to a rule that triggers on a traffic profile change. In addition, you cannot add a host profile qualification to a rule that triggers on the detection of a new IP host.

If you employ the Real Time User Awareness (RUA) feature, a user qualification can be added to the compliance rule to track certain users or groups of users. For

Keysight, EX-1004
Page 1214

example, you could constrain a compliance 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.

Finally, you configure compliance rules to have snooze periods and inactive periods. When a compliance 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). During *inactive periods*, the compliance rule does not trigger.

When you create compliance rule trigger criteria, host profile qualifications, user qualifications, and flow trackers, you can use simple conditions or you can create more elaborate constructs by combining and nesting conditions. The syntax you can use within conditions varies depending on the element you are creating, but the mechanics remains consistent. For more information, See Understanding Rule-Building Mechanics on page 1251.

To create a compliance rule:

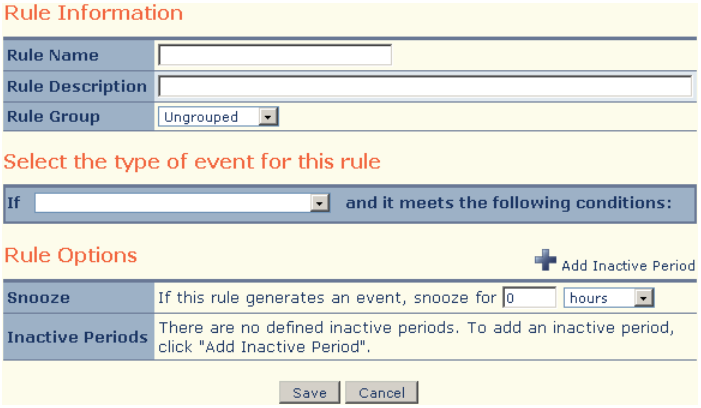
Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Rule Management**.

   The Rule Management page appears.

2. Click **Create Rule**.

   The Create Rule page appears.



3. Provide basic rule information, such as the rule name, description, and group.

   See Providing Basic Rule Information on page 1216.

4. Specify the basic criteria on which you want the rule to trigger.

   See Specifying Compliance Rule Trigger Criteria on page 1216.

5. Optionally, add a host profile qualification to the rule.

   See Adding a Host Profile Qualification on page 1233.

Keysight, EX-1004
Page 1215

6. Optionally, add a flow tracker to the rule.

   See Adding a Flow Tracker on page 1237.

7. Optionally, add a user qualification to the rule.

   See Adding a User Qualification on page 1248.

8. Optionally, add an inactive period or snooze period (or both) to the rule.

   See Adding Snooze and Inactive Periods on page 1250.

9. Click **Save Rule**.

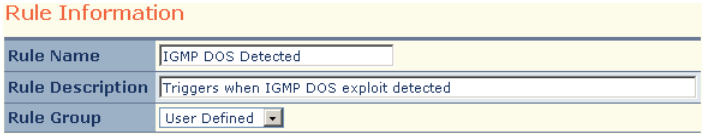
   The rule is saved. You can now use the rule within compliance policies or within other compliance rules that trigger on the same event type.

Providing Basic Rule Information

Platform: DC

Component: None

You must give each compliance rule a name, and, optionally, a short description. You can also place the rule in a rule group.



To provide basic rule information:

Access: P&R Admin/Admin

1. On the Create Rule page, in the **Rule Name** field, type a name for the rule.
2. In the **Rule Description** field, type a description for the rule.
3. Optionally, select a group for the rule from the **Rule Group** drop-down list.

   For more information on rule groups, see Managing Rules for Compliance Policies on page 1260.

4. Continue with the procedure in the next section, Specifying Compliance Rule Trigger Criteria.

Specifying Compliance Rule Trigger Criteria

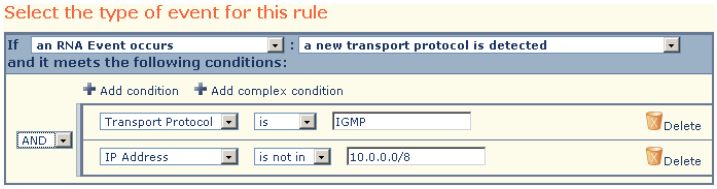
Platform: DC

Component: None

A simple compliance rule requires only that any flow, intrusion, RNA, RUA, or host input event occur; you do not need to provide more specific conditions. Also, compliance rules based on traffic profiles require only one condition. In contrast, compliance rules can be complex, with multiple nested conditions. For example,

Keysight, EX-1004
Page 1216

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.



To specify compliance rule trigger criteria:

Access: P&R Admin/Admin

1. Select the type of event on which you want to base your rule.

   When you build a compliance rule, first you must 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**:

   - Requires: DC + RNA Select **a flow event occurs** to trigger the rule when flow data meets specific criteria.
   - Requires: DC + IPS Select **an intrusion event occurs** to trigger your rule when a specific intrusion event occurs.
   - Requires: DC + RNA Select **an RNA event occurs** to trigger the rule when a specific RNA event occurs. When triggering a compliance rule on an RNA event, you must also choose the type of event you want to use. You can choose from a subset of the RNA events described in Understanding RNA Network Discovery Event Types on page 1005; you cannot, for example, trigger a compliance rule on a confidence change. You can, however, choose **there is any type of event** to trigger the rule when any kind of RNA event occurs.
   - Requires: DC + RUA Select **an RUA event occurs** to trigger the rule when a new user identity is detected or a user logs in to a host.
   - Requires: DC + RNA Select **a host input event occurs** to trigger the rule when a specific host input event occurs. When triggering a compliance 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 RNA Host Input Event Types on page 1010.
   - Requires: DC + RNA Select **a traffic profile changes** to trigger the compliance rule when network traffic deviates from your normal network traffic pattern as characterized in an existing traffic profile.

Keysight, EX-1004
Page 1217

2. Specify the rule's conditions.

   The syntax you can use within compliance 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 1251.

   The syntax you can use to build conditions is described in the following sections.

   - Syntax for Flow Events on page 1218
   - Syntax for Intrusion Events on page 1221
   - Syntax for RNA Events on page 1223
   - Syntax for RUA Events on page 1226
   - Syntax for Host Input Events on page 1227
   - Syntax for Traffic Profile Changes on page 1229

   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 a new TCP service, 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 next sections to add a host profile qualification, a flow tracker, and inactive and snooze periods.

   - Adding a Host Profile Qualification on page 1233
   - Adding a Flow Tracker on page 1237
   - Adding Snooze and Inactive Periods on page 1250

Syntax for Flow Events

Platform: DC

Component: RNA

The Syntax for Flow Events table describes how to build a compliance rule condition when you choose a flow event as the base event.

Note that if you configure your 3D Sensors to transmit flow data to the Defense Center as flow summaries only, when you build a compliance rule with a flow event as the base event, you should use data that the flow summaries contain. A trigger criterion based on information that is not included in a flow summary automatically evaluates to true.

For example, flow summaries do not include information on the initiator port. Therefore, if you build a compliance rule that requires that the initiator must be using a specific port, the rule will trigger for every flow summary that meets the other criteria in the rule. If that is the only criterion in the rule, the rule will trigger for every flow summary that the Defense Center receives. For more information, see Understanding Flow Summary Data on page 1076.

Keysight, EX-1004
Page 1218

In addition, you should keep in mind that flows detected by 3D Sensors with RNA and flows exported by NetFlow-enabled devices contain different information. For example, flows detected by 3D Sensors with RNA do not contain TCP flag information. Therefore, if you want to specify that a flow event have a certain TCP flag to trigger a compliance rule, none of the flows detected by 3D Sensors with RNA will trigger the rule. As another example, NetFlow records do not contain information about which host in the flow is the initiator and which is the responder. When RNA 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 How is NetFlow Data Different from RNA Flow Data? on page 1142.

Syntax for Flow Events

If you specify...	Select an operator, then...
Flow Duration	Type the duration of the flow event, in seconds.
Initiator IP, Responder IP, or Initiator/Responder IP	Either type a single IP address or use CIDR notation to specify a range of IP addresses. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
Initiator Port or Responder Port	Type the port number.
Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Detection Engine	Select the detection engine that detected the flow (for flows detected by 3D Sensors with RNA), or processed the flow (for flow data exported by a NetFlow-enabled device).
Flow Type	Select whether you want to trigger the compliance rule based on whether the flow was detected by a 3D Sensor with RNA (**RNA**) or was exported by a NetFlow-enabled device (**NetFlow**).
NetFlow Device	Select the IP address of the NetFlow-enabled device that exported the flow data you want to use to trigger the compliance rule. If you did not add any NetFlow-enabled devices to your deployment (using the system settings), the NetFlow Device drop-down list is blank.
TCP Flags	Select a TCP flag that a flow must contain in order to trigger the compliance rule. **IMPORTANT!** Only flows exported by NetFlow-enabled devices contain TCP flag data.
Service Name	Select one or more service names from the list of available services.

Keysight, EX-1004
Page 1219

Syntax for Flow Events (Continued)

If you specify...	Select an operator, then...
Application	Select one or more applications from the list of available applications.
Application Type	Select one or more application types from the list of available application types. Application types include web browsers, email clients, instant messengers, and so on.
Application Version	Type the version number of the application.
Initiator Bytes, Responder Bytes, or Total Bytes	Type one of: • the number of bytes transmitted by the hosts on the monitored network segment (**Initiator Bytes**). • the number of bytes received by the hosts on the monitored network segment (**Responder Bytes**). • the number of bytes both transmitted and received by the hosts on the monitored network segment (**Total Bytes**).
Initiator Packets, Responder Packets, or Total Packets	Type one of: • the number of packets transmitted by the hosts on the monitored network segment (**Initiator Packets**). • the number of packets received by the hosts on the monitored network segment (**Responder Packets**). • the number of packets both transmitted and received by the hosts on the monitored network segment (**Total Packets**)

Keysight, EX-1004
Page 1220

Syntax for Intrusion Events

Platform: DC

Component: IPS

The Syntax for Intrusion Events table describes how to build a compliance rule condition when you choose an intrusion event as the base event.

Syntax for Intrusion Events

If you specify...	Select an operator, then...
Source IP, Destination IP, or Source/Destination IP	Either type a single IP address or use CIDR notation to specify a range of IP addresses. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
Generator ID	Select one or more preprocessors from the list of available preprocessors. See Using Advanced Settings in an Intrusion Policy on page 477 for more information about available preprocessors.
Classifications	Select one or more classifications from the list of available classifications.
Rule SID	Type a single Snort 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 Message	Type all or part of the rule message.
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 864.
Priority	Select the rule priority from the list of available priorities (**low**, **medium**, or **high**). 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.

Keysight, EX-1004
Page 1221

Syntax for Intrusion Events (Continued)

If you specify...	Select an operator, then...
Impact Flag	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: • 1 - red (Vulnerable) • 2 - orange (Potentially Vulnerable) • 3 - yellow (Currently Not Vulnerable) • 4 - blue (Unknown Target) • 5 - gray (Unknown) Note that impact flags (other than gray) are assigned to intrusion events only when you deploy IPS and RNA on the same network segment and manage both with the same Defense Center. **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 red (Vulnerable) impact flags for intrusion events involving those hosts, unless you use the host input feature to manually set the host operating system identity. For more information, see Using Impact Flags to Evaluate Events on page 307.
Source Port or Destination Port	Type the port number.
Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Detection Engine	Select one or more detection engines that may have generated the event.
Inline Result	Select either: • `dropped`, to specify whether the packet was 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 IPS 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 435, Setting Drop Behavior in an Inline Deployment on page 358, and Tap Mode on page 227.
VLAN ID	Type the innermost VLAN ID associated with the packet that triggered the intrusion event

Keysight, EX-1004
Page 1222

Syntax for RNA Events

Platform: DC

Component: RNA

If you base your compliance rule on an RNA event, you must first choose the type of RNA 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 RNA event types. For detailed descriptions of RNA event types, see Understanding RNA Network Discovery Event Types on page 1005.

Compliance Rule Trigger Criteria for RNA Events vs. RNA Event Types

This option in the drop-down list...	Corresponds to this RNA event type...
a client application has changed	Client Application Update
a client application timed out	Client Application 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
the OS or service identity for a host has a conflict	Identity Conflict
the OS or service identity for a host has timed out	Identity Timeout
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 application is detected	New Client Application
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

Keysight, EX-1004
Page 1223

Compliance Rule Trigger Criteria for RNA Events vs. RNA Event Types (Continued)

This option in the drop-down list...	Corresponds to this RNA event type...
a new network protocol is detected	New Network Protocol
a new TCP service is detected	New TCP Service
a new transport protocol is detected	New Transport Protocol
a new UDP service is detected	New UDP Service
the OS information for a host has changed	New OS
a TCP port closed	TCP Port Closed
a TCP port timed out	TCP Port 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 service	TCP Service Information Update
there is new information about a UDP service	UDP Service Information Update
a UDP port closed	UDP Port Closed
a UDP port timed out	UDP Port Timeout
a VLAN tag was updated	VLAN Tag Information Update

Note that you cannot trigger a compliance rule on confidence or hops changes, or when RNA 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 RNA event occurs.

After you choose the RNA event type, you can build compliance rule conditions as described in the Syntax for RNA Events table below. Depending on the type of RNA event you choose, you can build conditions using subsets of the criteria in the following table. For example, if you trigger your compliance rule when a new client application is detected, you can build conditions based on the IP or MAC

Keysight, EX-1004
Page 1224

address of the host, the application name, type, or version, and the detection engine that detected the event.

Syntax for RNA Events

If you specify...	Select an operator, then...
Application	Select one or more applications from the list of available applications.
Application Type	Select one or more application types from the list of available application types. Application types include web browsers, email clients, instant messengers, and so on.
Application Version	Type the version number of the application.
Detection Engine	Select the detection engine that generated the event.
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.
IP Address or New IP Address	Either type a single IP address or use CIDR notation to specify a range of IP addresses. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
MAC Address	Type all or part of the MAC address of the host. For example, if you know that a certain hardware manufacturer's devices' MAC addresses begins 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 RNA positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**), or whether RNA is seeing many hosts with that MAC address because, for example, there is a router between the sensor 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 network discovery event.
NETBIOS Name	Type the NetBIOS name of the host.
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 from the list of available names.
OS Vendor	Select one or more operating system vendor names from the list of available vendors.

Keysight, EX-1004
Page 1225

Syntax for RNA Events (Continued)

If you specify...	Select an operator, then...
OS Version	Select one or more operating system versions from the list of available versions.
Payload	Select one or more payloads from the list of available payloads.
Payload Type	Select one or more payload types from the list of available payload types.
Protocol or Transport Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Service Name	Select one or more services from the list of available services.
Service Port	Type the service port number.
Source	Select the source of the host input data (for operating system and service identity changes and timeouts) from the list of available sources.
Source Type	Select the type of the source for the host input data (for operating system and service identity changes and timeouts) from the list of available source types.
VLAN ID	Type the VLAN ID of the host involved in the event.

Syntax for RUA Events

Platform: DC

Component: RUA

If you base your compliance rule on an RUA event, choose the type of RUA event you want to use from a drop-down list. After you choose the RUA event type, you can build compliance rule conditions as described in the Syntax for RUA Events table.

Depending on the type of RUA event you choose, you can build conditions using subsets of the criteria in the following table. For example, if you trigger your compliance rule when a user logs in to a host, you can build conditions based on the IP host.

Syntax for RUA Events

If you specify...	Select an operator, then...
A new User Identity is detected	Type the username.
A user logs in to a host	Type the host IP address or the username.

Keysight, EX-1004
Page 1226

Syntax for Host Input Events

Platform: DC

Component: RNA

If you base your compliance 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 RNA Host Input Event Types on page 1010.

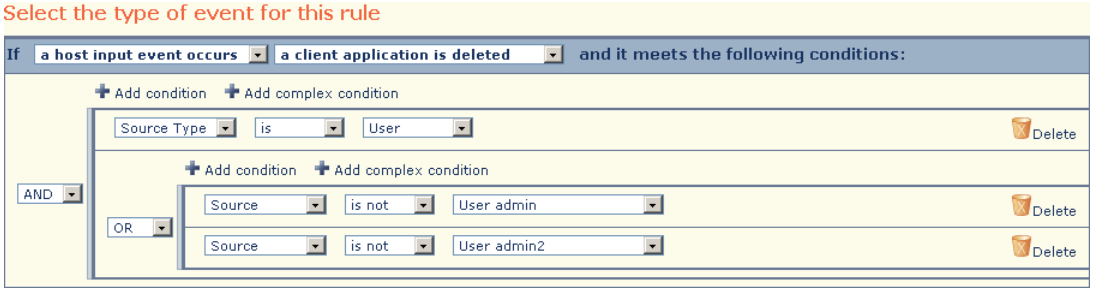
Compliance Rule Trigger Criteria for Host Input Events vs. Host Input Event Types

This option in the drop-down list...	Corresponds to this host input event type...
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
a client applicaiton is added	Add Client Application
a client application is deleted	Delete Client Application
host criticality is set	Set Host Criticality
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 service definition is set	Set Service Definition
a service is added	Add Service
a service is deleted	Delete Service
a vulnerability is marked invalid	Vulnerability Set Invalid
a vulnerability is marked valid	Vulnerability Set Valid

Keysight, EX-1004
Page 1227

Note that you cannot trigger a compliance 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 compliance rule conditions as described in the Syntax for Host Input Events table below. For example, if you wanted an alert if a user other than one of your administrative users changed client application data in your network map, you could set up a rule that notifies you when a client application is deleted from a host. You could specify that the rule triggers if the Source Type is User and the Source is neither admin or admin2, as shown in the following graphic.



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 compliance rule when a client application 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	Either type a single IP address or use CIDR notation to specify a range of IP addresses. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
Source Type	Select the type of the source for the host input data from the list of available source types.
Source	Select the source for the host input data from the list of available sources.

Keysight, EX-1004
Page 1228

Syntax for Traffic Profile Changes

Platform: DC

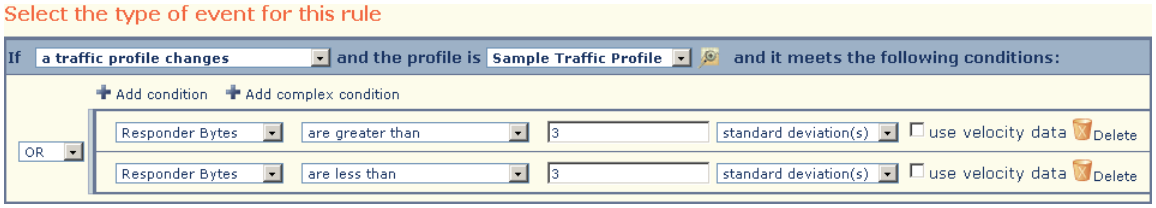
Component: RNA

If you base your compliance 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 1118.

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.



  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

Keysight, EX-1004
Page 1229

You can also use velocity data, which uses 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

For more information on velocity data, see Changing the Graph Type on page 1109.

To use velocity data:

Access: P&R Admin/Admin

▶ Check **use velocity data** next to the condition you are building.

The Syntax for Traffic Profile Changes table describes how to build a condition in a compliance rule when you choose a traffic profile change as the base event.

If your traffic profile uses flow data exported by NetFlow-enabled devices, there are two important points to keep in mind:

- For flows detected by 3D Sensors, RNA generates one bidirectional record, or flow event. On the other hand, NetFlow-enabled devices export unidirectional flow data, so RNA generates at least two flow events for each flow detected by NetFlow-enabled devices, depending on whether you configured the device to output records at a fixed interval even if a flow is still ongoing.
- NetFlow records do not contain information about which host in the flow is the initiator and which is the responder. When RNA 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.

Keysight, EX-1004
Page 1230

For more information, see How is NetFlow Data Different from RNA Flow Data? on page 1142.

Syntax for Traffic Profile Changes

If you specify...	Select an operator, then...	And then choose one of the following...
Number of Flows	Type the total number of flows detected on the monitored network segment. **or** Type the number of standard deviations either above or below the mean that the number of flows detected must be in to trigger the rule.	• flows • standard deviation(s)
Total Bytes, Initiator Bytes, or Responder Bytes	Type one of: • the total bytes transmitted on the monitored network segment (**Total Bytes**) • the number of bytes transmitted by the hosts on the monitored network segment (**Initiator Bytes**) • the number of bytes received by the hosts on the monitored network segment (**Responder Bytes**) **or** Type 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)

Keysight, EX-1004
Page 1231

Syntax for Traffic Profile Changes (Continued)

If you specify...	Select an operator, then...	And then choose one of the following...
Total Packets, Initiator Packets, or Responder Packets	Type one of: • the total packets transmitted on the monitored network segment (**Total Packets**) • the number of packets transmitted by the hosts on the monitored network segment (**Initiator Packets**) • the number of packets received by the hosts on the monitored network segment (**Responder Packets**) **or** Type 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	Type the number of unique hosts that initiated sessions that were detected on the monitored network segment. **or** Type 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	Type the number of unique hosts that responded to sessions that were detected on the monitored network segment. **or** Type 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)

Keysight, EX-1004
Page 1232

Adding a Host Profile Qualification

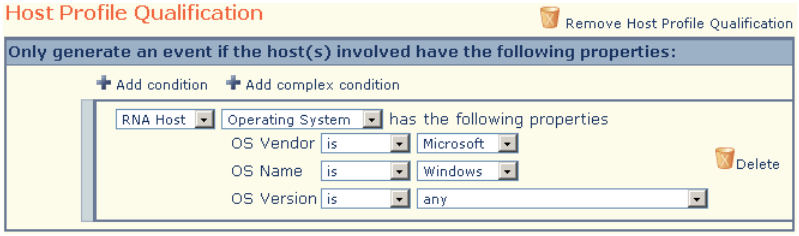
Platform: DC

Component: RNA

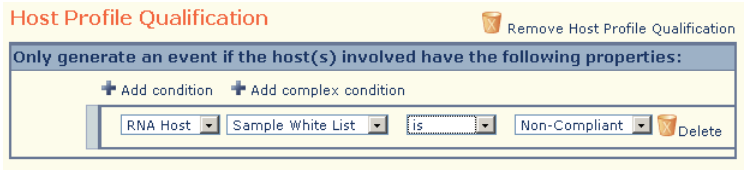
If you are using an flow, intrusion, RNA, RUA, or host input event to trigger your compliance rule, you can constrain the rule based on the host profile of a host involved in the event. This constraint is called a *host profile qualification*.

IMPORTANT! You **cannot** add a host profile qualification to a compliance rule that triggers on a traffic profile change or on the detection of a new IP host.

For example, as shown in the following graphic, you could constrain a compliance 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, the host profile qualification shown in the following graphic constrains a compliance rule so that it triggers only when the host is out of compliance with a white list.



Note that to use a host profile qualification, the host must exist in the RNA 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 compliance 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.

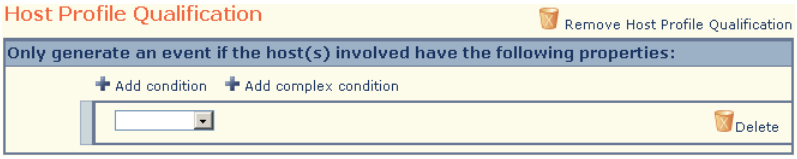
Keysight, EX-1004
Page 1233

To add a host profile qualification:

Access: P&R Admin/Admin

1. 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**.

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 Rule-Building Mechanics on page 1251 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 1234.

3. Optionally, continue with the procedures in the following sections to add a flow tracker and inactive and snooze periods.
   - Adding a Flow Tracker on page 1237
   - Adding Snooze and Inactive Periods on page 1250

Syntax for Host Profile Qualifications

Platform: DC

Component: RNA

When you build a host profile qualification condition, you must first select the host you want to use to constrain your compliance 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 flow event, select **Responder Host** or **Initiator Host**.
- If you are using an intrusion event, select **Destination Host** or **Source Host**.
- If you are using an RNA or host input event, select **RNA Host**.
- If you are using an RUA event, select **RNA 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 RNA detection policy to add hosts to the network map based on data exported by NetFlow-enabled devices, the

Keysight, EX-1004
Page 1234

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 flow data exported by NetFlow-enabled devices, keep in mind that NetFlow records do not contain information about which host in the flow is the initiator and which is the responder. When RNA 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 How is NetFlow Data Different from RNA Flow Data? on page 1142.

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 Name	Select one or more operating system names from the drop-down list.
Operating System > OS Vendor	Select one or more operating system vendor 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 Assigning a Host Criticality Value to a Host on page 986.
VLAN ID	Type the VLAN ID associated with the host.
Service > Service Name	Select one or more services from the drop-down list.
Service > Service Port	Type the service port number.
Service > Service Protocol	Select one or more protocols from the drop-down list.

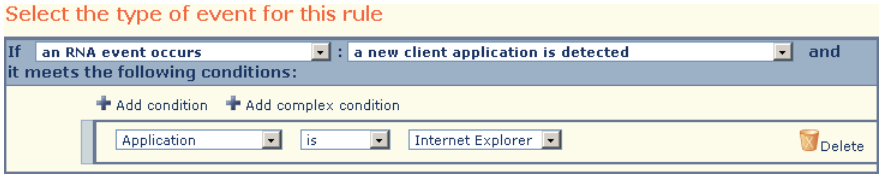
Keysight, EX-1004
Page 1235

Syntax for Host Profile Qualifications (Continued)

If you specify...	Select an operator, then...
Client Application > Application Type	Select one or more application types from the drop-down list. Application types include web browsers, email clients, instant messengers, and so on.
Client Application > Application	Select one or more applications from the drop-down list.
Client Application > Application Version	Type the application version.
MAC Address > MAC Address	Type all or part of the MAC address of the host. For example, if you know that a certain hardware manufacturer's devices' MAC addresses begins 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 RNA positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**), whether RNA is seeing many hosts with that MAC address because, for example, there is a router between the sensor 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: • 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. For more information on host attributes, see Working with User-Defined Host Attributes on page 987.

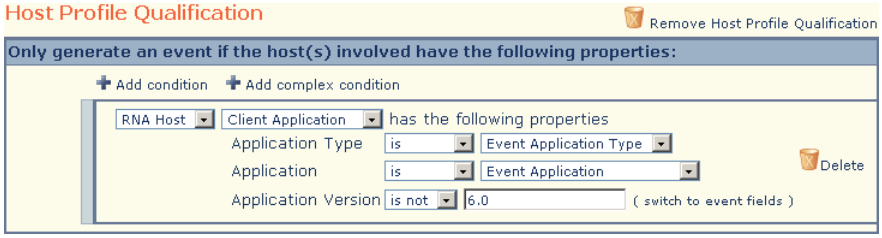
Keysight, EX-1004
Page 1236

Note that you can often use event data when constructing a host profile qualification. For example, assume your compliance rule triggers when RNA 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 6.0).

You could add a host profile qualification to this compliance rule so that the rule triggers only if the **Application Type** is the **Event Application Type** (that is, web browser) and the **Application** is the **Event Application** (that is, Internet Explorer), but the **Application Version** is not `6.0`.



TIP! To specify that the application version be the same as the one detected in the event, click **switch to event fields** and select **Event Application Version**. Click **switch to manual entry** to go back to manually specifying the version.

Adding a Flow Tracker

Platform: DC

Component: RNA

If you are using an flow, intrusion, RNA, RUA, or host input event to trigger your compliance rule, you can add a *flow tracker* to the rule. A flow tracker constrains a compliance rule so that once its initial criteria are met, RNA begins tracking certain flows. The Defense Center generates a compliance event for the rule only if the tracked flows meet additional criteria. You **cannot** add a flow tracker to a rule that triggers on a traffic profile change.

Unlike traffic profiles, which typically monitor a broad range of network traffic and run persistently, flow trackers typically monitor very specific traffic and, when

Keysight, EX-1004
Page 1237

triggered, run only for a finite, specified time. For more information on traffic profiles, see Creating Traffic Profiles on page 1118.

IMPORTANT! If you configure your 3D Sensors to transmit flow data to the Defense Center as flow summaries only, when you build a flow tracker, you should use data that the flow summaries contain. A flow tracker criterion based on information that is not included in a flow summary automatically evaluates to true. For example, flow summaries do not include information on the initiator port. Therefore, if you build a flow tracker that tracks every flow summary where the initiator is using a specific port, the Defense Center will track every flow summary that meets the other criteria in the flow tracker. If that is the only criterion in the flow tracker, the Defense Center will track every flow summary that it receives. For more information, see Understanding Flow Summary Data on page 1076.

When you configure a flow tracker, you must specify:

- which flows you want to track
- the conditions that the flows you are tracking must meet for the Defense Center to generate a compliance event
- the amount of time that can elapse before the flow tracker "expires," that is, the amount of time after which RNA stops tracking flows if the flow tracker conditions have not yet been met

When you add a flow tracker to a compliance rule, once your network traffic satisfies the basic conditions of the rule as well as of any host profile qualification you added, RNA begins tracking flows according to the criteria you specified.

If your network traffic meets the flow tracker's conditions before the time period you specify expires, the Defense Center generates a compliance event and immediately stops tracking flows for that rule instance. On the other hand, if time expires before your network traffic meets the conditions in the flow tracker, the Defense Center does not generate a compliance event. After time expires, RNA stops tracking flows for that rule instance.

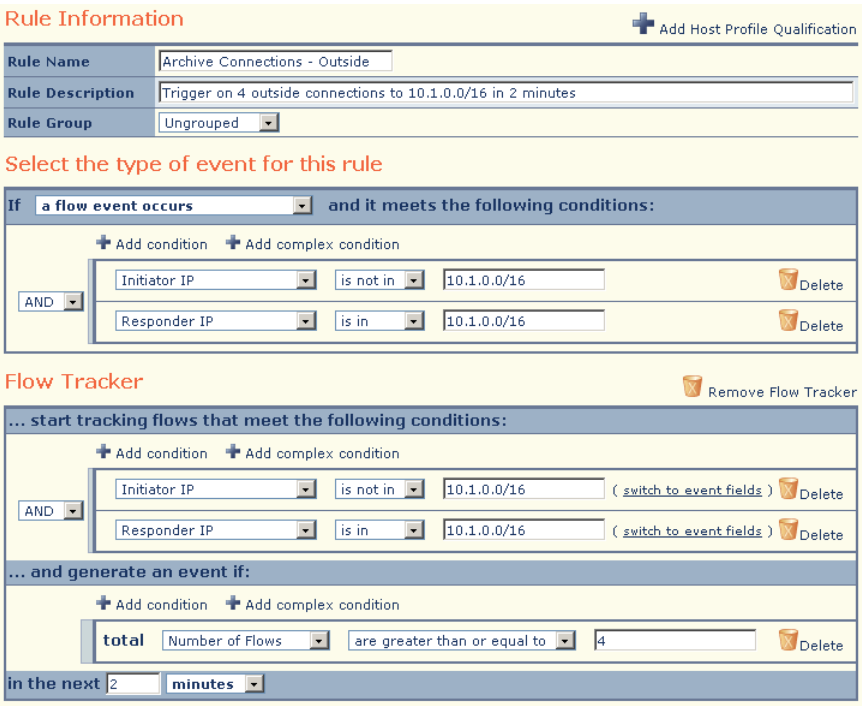
TIP! You can add a flow tracker to a simple compliance rule that requires only that any flow, intrusion, RNA, RUA, or host input event occurs.

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.

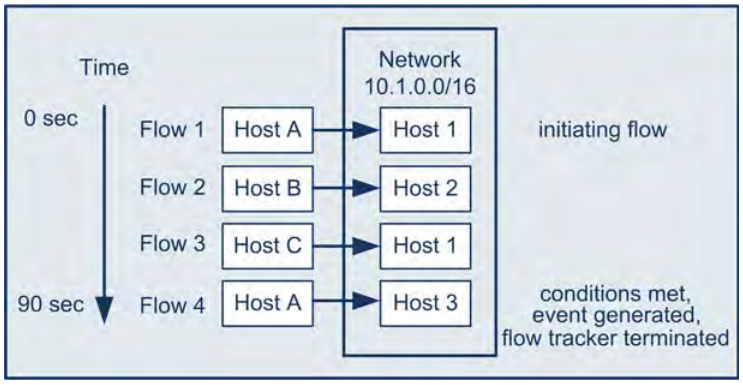
The rule shown in the following graphic specifies that once a connection occurs from outside the 10.1.0.0/16 network to inside the network, RNA begins tracking flows that meet that criterion. The Defense Center then generates a compliance

Keysight, EX-1004
Page 1238

event if RNA detects four flows (including the original flow) within two minutes that match that signature.



The following graphic shows how network traffic can trigger the above compliance rule.



In this example, RNA detected a flow that met the basic conditions of the compliance rule, that is, RNA detected a connection from a host outside the 10.1.0.0/16 network to a host inside the network. This created a flow tracker.

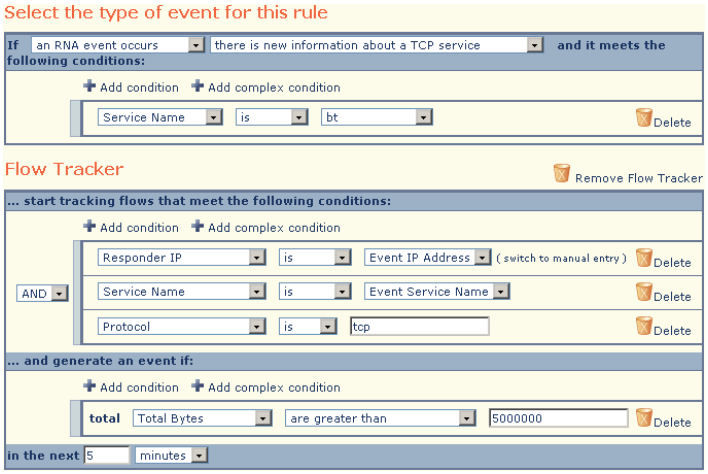
Keysight, EX-1004
Page 1239

RNA processes the flow tracker in the following stages:

1. RNA starts tracking flows when it detects a connection from Host A outside the network to Host 1 inside the network.
2. RNA detects two more flows that match the flow tracker signature: Host B to Host 2 and Host C to Host 1.
3. RNA detects a fourth qualifying flow when Host A connects to Host 3 within the two-minute time limit. The rule conditions are met.
4. The Defense Center generates a compliance event and RNA stops tracking flows.

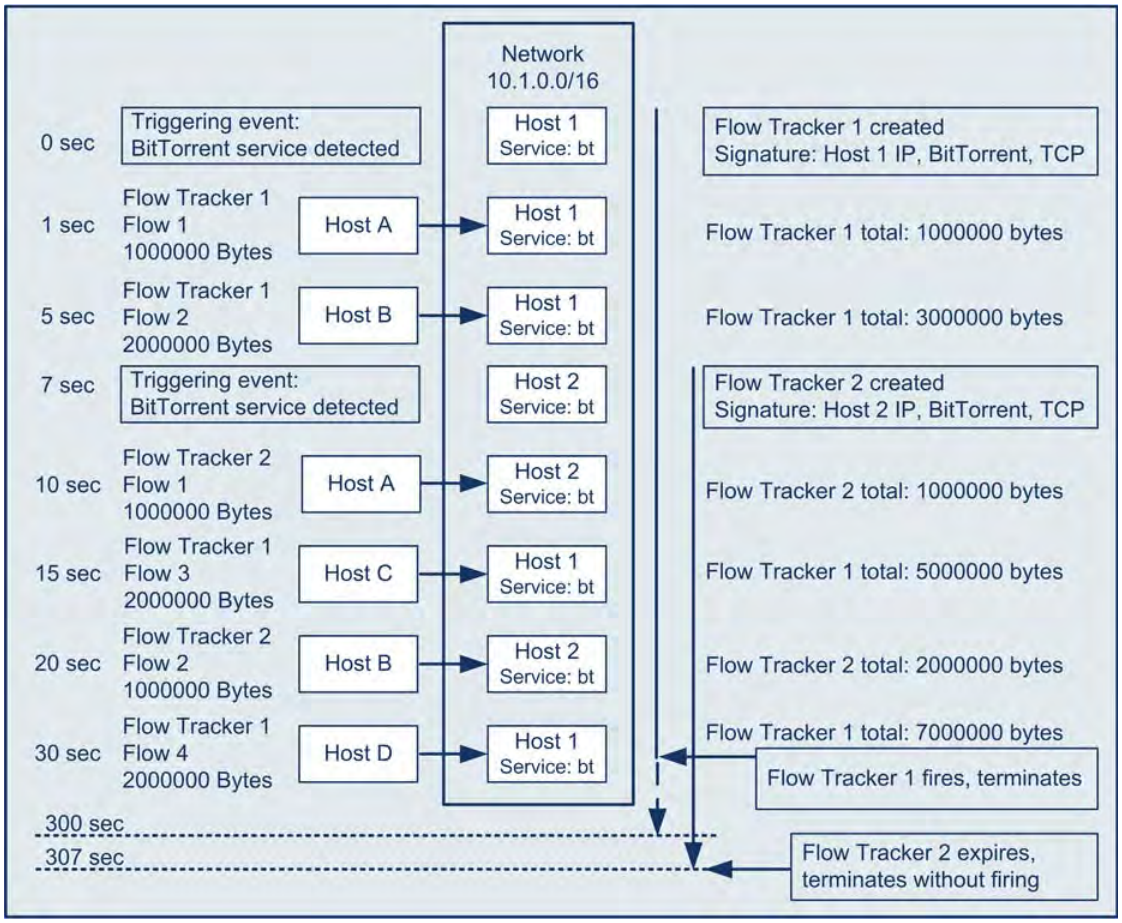
Note that when you add a flow tracker to a compliance rule, depending on the type of event you use to trigger the rule, the flow tracker is populated with default constraints. For example, the following graphic shows a flow tracker added to a rule that triggers when RNA detects the Bittorrent service on your monitored network.

The flow tracker constrains the rule so that it does not trigger until more than 5MB of data is transmitted by the host, using the service and protocol in the RNA event that initially violated the compliance rule, in the five minutes following the initial policy violation.



Keysight, EX-1004
Page 1240

The following graphic shows how network traffic can trigger the above compliance rule.



In this example, RNA detected two flows that meet the basic conditions of the compliance rule, that is, RNA detected the BitTorrent TCP service on two different hosts: Host 1 and Host 2. This created two flow trackers, Flow Tracker 1 and Flow Tracker 2, each with its own signature.

The signatures represent unique combinations of protocol, IP address, and service. That is, RNA created two flow trackers because the IP addresses of Host 1 and Host 2 differ. If RNA had detected the BitTorrent TCP service on a different host, it would have created a third flow tracker. However, it would **not** have created an additional flow tracker if it detected the BitTorrent UDP service on either Host 1 or Host 2. This is because the basic conditions of the compliance rule require the TCP protocol.

Keysight, EX-1004
Page 1241

RNA processes Flow Tracker 1 in the following stages:

1. RNA starts tracking flows at the 0-second marker when RNA detects the BitTorrent service on Host 1. Note that Flow 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. RNA detects that Host 1 transmitted data that matches the signature for Flow 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. RNA 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 compliance event and RNA stops tracking flows.

   Note that the Defense Center generates the compliance event *after* Host 1 transmits the entire 2MB to Host D, because it does not tally flow data until the session terminates.

   Also, note that the flow tracker terminates after the Defense Center generates the compliance event, even though the 5-minute period has not expired. If RNA detects the BitTorrent TCP service on Host 1 at any time after the flow tracker terminates, it will create a new flow tracker.

RNA processes Flow Tracker 2 in the following stages:

1. RNA starts tracking flows at the 7-second marker when RNA detects the BitTorrent service on Host 2. Note that Flow 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. RNA detects that Host 2 transmitted data, totaling 2MB, that matches the signature for Flow Tracker 2:
   - 1MB to Host A, at the 10-second marker
   - 1MB to Host B, at the 20-second marker
3. Flow Tracker 2 expires after the 5 minutes, at the 307-second marker. RNA has not detected any more data transmissions that match the appropriate signature.
4. RNA stops tracking flows and does not generate a compliance event.

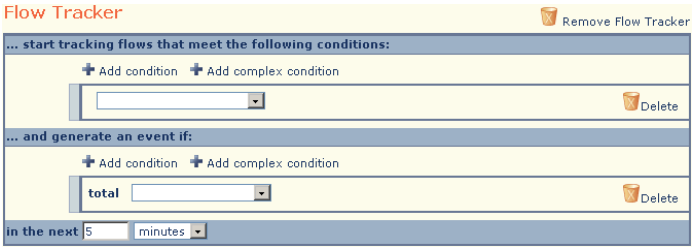
Keysight, EX-1004
Page 1242

To add a flow tracker:

Access: P&R Admin/Admin

1. On the Create Rule Page, click **Add Flow Tracker**.

   The Flow Tracker section appears.

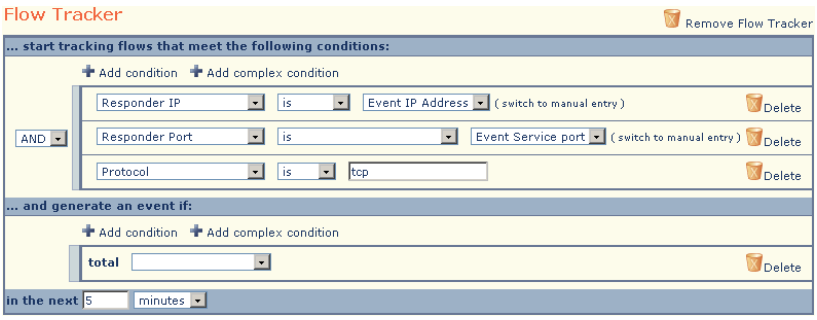


TIP! To remove a flow tracker click **Remove Flow Tracker**.

2. Specify which flows you want to track by setting flow tracker criteria.

   You can set flow tracker criteria by creating a single, simple condition, or you can create more elaborate constructs by combining and nesting conditions.

   For example, the following graphic shows three conditions that specify that RNA track flows that occur on the port involved in the policy violation and that involve the host that triggered the policy violation.



   See Understanding Rule-Building Mechanics on page 1251 for information on how to use the web interface to build conditions. The syntax you can use to build flow tracker conditions is described in Syntax for Flow Trackers on page 1244.

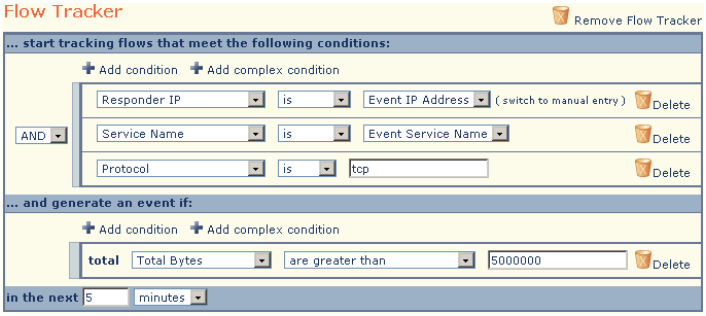
Keysight, EX-1004
Page 1243

3. Based on the flows you decided to track in step 2, describe when you want to generate a compliance 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 compliance event.

   For example, the following graphic shows a flow tracker that generates an event if more than about 5MB of data are transmitted by the host in the five minutes following the policy violation.



   See Understanding Rule-Building Mechanics on page 1251 for information on how to use the web interface to build conditions. The syntax you can use to build flow tracker conditions is described in Syntax for Flow Tracker Events on page 1247.

4. Optionally, continue with the procedure in the next section, Adding Snooze and Inactive Periods, to add inactive and snooze periods.

Syntax for Flow Trackers

Platform: DC

Component: RNA

The Syntax for Flow Trackers table describes how to build a flow tracker condition that specifies the kind of flows you want to track.

You should keep in mind that flows detected by 3D Sensors with RNA and flows exported by NetFlow-enabled devices contain different information. For example, flows detected by 3D Sensors with RNA do not contain TCP flag information. Therefore, if you want to specify that a flow event have a certain TCP flag to trigger a compliance rule, none of the flows detected by 3D Sensors with RNA will trigger the rule. As another example, NetFlow records do not contain information about which host in the flow is the initiator and which is the responder. When RNA processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether

Keysight, EX-1004
Page 1244

those ports are well-known. For more information, see How is NetFlow Data Different from RNA Flow Data? on page 1142.

Syntax for Flow Trackers

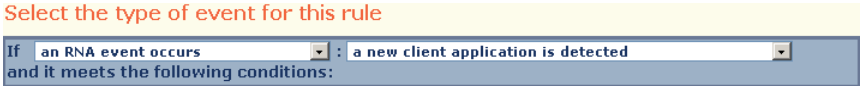
If you specify...	Select an operator, then...
Flow Duration	Type the flow duration, in seconds.
Initiator IP, Responder IP, or Initiator/Responder IP	Either type a single IP address or use CIDR notation to specify a range of IP addresses. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.
Initiator Port or Responder Port	Type the port number.
Protocol	Type the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Detection Engine	Select the detection engine that detected the flow (for flows detected by 3D Sensors with RNA), or processed the flow (for flow data exported by a NetFlow-enabled device).
Flow Type	Select whether you want to track flows based on how they were detected: by a 3D Sensor with RNA (**RNA**) or exported by a NetFlow-enabled device (**NetFlow**).
NetFlow Device	Select the IP address of the NetFlow-enabled device that exported the flows you want to track. If you did not add any NetFlow-enabled devices to your deployment (using the system settings), the NetFlow Device drop-down list is blank.
TCP Flags	Select the TCP flag that flows must contain in order to track them. **IMPORTANT!** Only flows exported by NetFlow-enabled devices contain TCP flag data.
Service Name	Select one or more services from the list of available services.
Application	Select one or more applications from the list of available applications.

Keysight, EX-1004
Page 1245

Syntax for Flow Trackers (Continued)

If you specify...	Select an operator, then...
Application Type	Select one or more application types from the list of available application types. Application types include web browsers, email clients, instant messengers, and so on.
Application Version	Type the version of the application.
Initiator Bytes, Responder Bytes, or Total Bytes	Type one of: • the number of bytes transmitted by the hosts on the monitored network segment (**Initiator Bytes**) • the number of bytes received by the hosts on the monitored network segment (**Responder Bytes**) • the number of bytes both transmitted and received by the hosts on the monitored network segment (**Total Bytes**)
Initiator Packets, Responder Packets, or Total Packets	Type one of: • the number of packets transmitted by the hosts on the monitored network segment (**Initiator Packets**) • the number of packets received by the hosts on the monitored network segment (**Responder Packets**) • the number of packets both transmitted and received by the hosts on the monitored network segment (**Total Packets**)
Payload	Select one or more payloads from the list of available payloads.
Payload Type	Select one or more payload types from the list of available payload types.

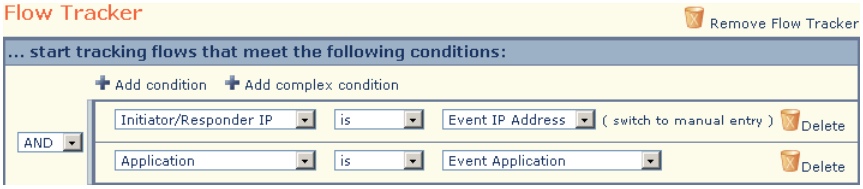
Note that you can often use event data when constructing a flow tracker. For example, assume your compliance rule triggers when RNA detects a new client application on one of your monitored hosts; that is, the rule triggers when an RNA event whose base event type is **a new client application is detected** is generated.



Keysight, EX-1004
Page 1246

Further assume that when you detect this new application, you want to track flows involving the new application on the host where it was detected. Because RNA "knows" the IP address of the host and the client application name, you can build a simple flow tracker that tracks those flows.

In fact, when you add a flow tracker to this type of compliance rule, the flow tracker is populated with those default constraints; that is, the **Initiator/Responder IP** is set to the **Event IP Address** and the **Application** is set to the **Event Application**.



TIP! To specify that the flow tracker track flows 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 Flow Tracker Events

Platform: DC

Component: RNA

The Syntax for Flow Tracker Events table describes how to how to build a flow tracker condition that specifies when you want to generate a compliance event based on the flows you are tracking.

Syntax for Flow Tracker Events

If you specify...	Select an operator, then...
Number of Flows	Type the total number of flows detected on the monitored network segment.
Total Bytes, Initiator Bytes, or Responder Bytes	Type one of: • the total bytes transmitted on the monitored network segment (**Total Bytes**) • the number of bytes transmitted by the hosts on the monitored network segment (**Initiator Bytes**) • the number of bytes received by the hosts on the monitored network segment (**Responder Bytes**)

Keysight, EX-1004
Page 1247

Syntax for Flow Tracker Events (Continued)

If you specify...	Select an operator, then...
Total Packets, Initiator Packets, or Responder Packets	Type one of: • the total packets transmitted on the monitored network segment (**Total Packets**) • the number of packets transmitted by the hosts on the monitored network segment (**Initiator Packets**) • the number of packets received by the hosts on the monitored network segment (**Responder Packets**)
Unique Initiators or Unique Responders	Type one of: • the number of unique hosts that initiated sessions that were detected on the monitored network segment (**Unique Initiators**) • the number of unique hosts that responded to flow transactions that were detected on the monitored network segment (**Unique Responders**)

Adding a User Qualification

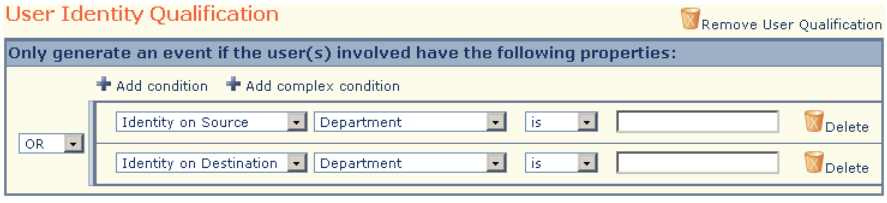
Platform: DC

Component: RUA

If you are using an flow, intrusion, RNA, or host input event to trigger your compliance rule, you can constrain the rule based on the identity of a user involved in the event. This constraint is called a *user identity qualification*.

IMPORTANT! You **cannot** add a user identity qualification to a compliance rule that triggers on a traffic profile change or on the detection of a RUA event.

For example, as shown in the following graphic, you could constrain a compliance rule so that it triggers only when the identity of the source or destination user is one from the sales department.



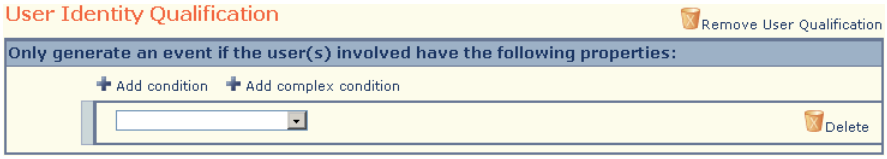
Keysight, EX-1004
Page 1248

To add a user identity qualification:

Access: P&R Admin/Admin

1. On the Create Rule Page, click **Add User Qualification**.

   The User Identity Qualification section appears.



   TIP! To remove a User Identity profile qualification, click **Remove User Qualification**.

2. Build the user 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 1251 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 Identity Qualifications on page 1249.

3. Optionally, continue with the procedures in the following sections to add a flow tracker and inactive and snooze periods.
   - Adding a Flow Tracker on page 1237
   - Adding Snooze and Inactive Periods on page 1250

Syntax for User Identity Qualifications

Platform: DC

Component: RUA

When you build a user qualification condition, you must first select the identity you want to use to constrain your compliance 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 flow 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 an RNA 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.

Keysight, EX-1004
Page 1249

Note that the Defense Center obtains certain information about RUA users, including first and last names, department, telephone number, and email address, from an optional Defense Center-LDAP server connection. This information may not be available for all RUA users in the database.

Syntax for User Qualifications

If you specify...	Select an operator, then...
Username	Type the username of the RUA user you want to use to constrain the compliance rule.
Authentication Protocol	Select an authentication protocol (or user type) from the list of available protocols. This is the protocol that was used to detect the RUA user.
First Name	Type the first name of the RUA user you want to use to constrain the compliance rule.
Last Name	Type the last name of the RUA user you want to use to constrain the compliance rule.
Department	Type the department of the RUA user you want to use to constrain the compliance rule.
Phone	Type the telephone number of the RUA user you want to use to constrain the compliance rule.
Email	Type the email address of the RUA user you want to use to constrain the compliance rule.

Adding Snooze and Inactive Periods

Platform: DC

Component: None

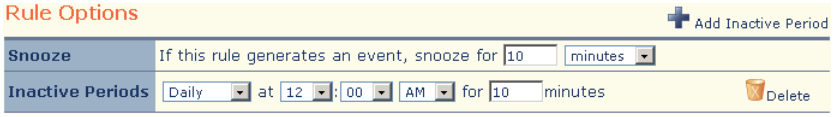
You can configure *snooze periods* in compliance rules. When a compliance rule triggers, a snooze period instructs the Defense Center to stop firing that rule for a specified interval, even if the rule is violated again during the interval. When the snooze period has elapsed, the rule can trigger again (and start a new snooze period).

For example, you may have a host on your network that should never generate traffic. A simple compliance rule that triggers whenever RNA detects a flow involving that host can create multiple compliance events in a short period of time, depending on the network traffic to and from the host. To limit the number of compliance events exposing your policy violation, you can add a snooze period so that the Defense Center generates a compliance event only for the first flow (within a time period that you specify) that RNA detects involving that host.

Keysight, EX-1004
Page 1250

You can also set up inactive periods in compliance rules. During inactive periods, the compliance 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 compliance 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 compliance rule that is configured with a snooze period and an inactive period.



To add a snooze period:

Access: P&R Admin/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: P&R Admin/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 compliance rule.

> **TIP!** To delete an inactive period, click **Delete** 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 Compliance Policies on page 1214 to save the rule.

Understanding Rule-Building Mechanics

Platform: DC

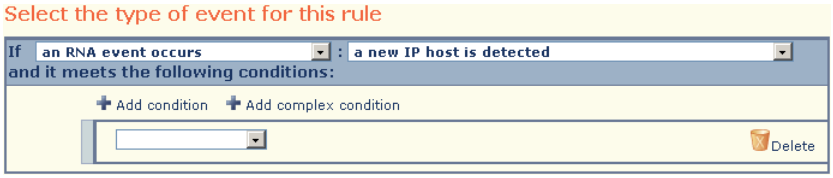
Component: None

You build compliance rules, flow trackers, user qualifications, and host profile qualifications by specifying the conditions under which they trigger. You can

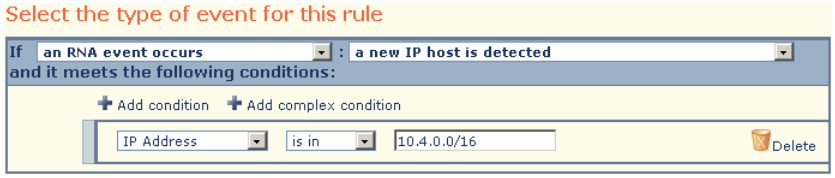
Keysight, EX-1004
Page 1251

create simple conditions, or you can create more elaborate constructs by combining and nesting conditions.

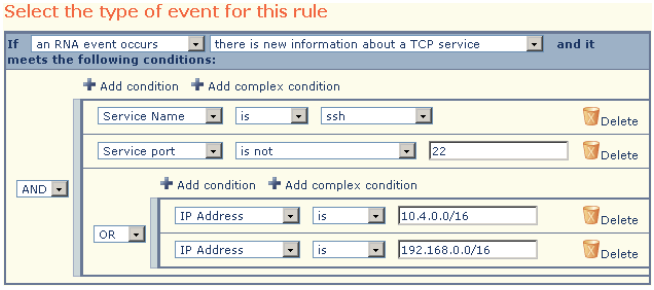
For example, if you want to generate a compliance event every time a new host is detected, you can create a very simple rule with no conditions, as shown in the following graphic.



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.



But the following rule, which detects ssh activity on a nonstandard 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 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 1253
- Adding and Linking Conditions on page 1255
- Using Multiple Values in a Condition on page 1258

Keysight, EX-1004
Page 1252

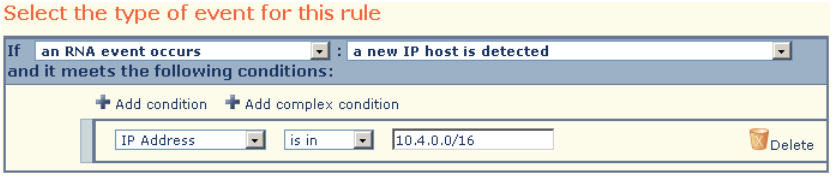
Building a Single Condition

Platform: DC

Component: None

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 compliance 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`.



To build the compliance rule trigger criteria in the example above:

Access: P&R Admin/Admin

1. Begin building a compliance rule.

   For more information, see Creating Rules for Compliance Policies on page 1214.

2. On the Create Rule page, under **Select the type of event for this rule**, select **an RNA 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.

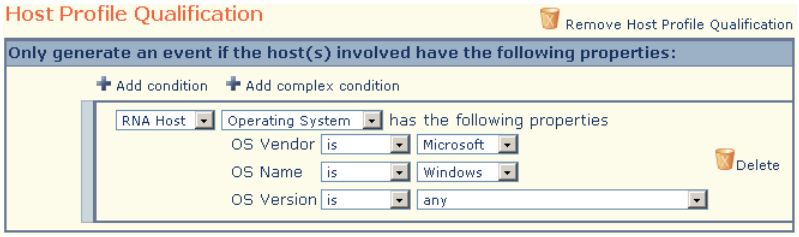
4. Select **is in** from the operator drop-down list that appears.

   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 CIDR notation. For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

5. Type `10.4.0.0/16` in the text field.

Keysight, EX-1004
Page 1253

In contrast, the following host profile qualification is more complex; it constrains a compliance rule such that the rule triggers only if the host involved in the RNA event on which the rule is based is running a version of Microsoft Windows.



To build the host profile qualification in the example above:

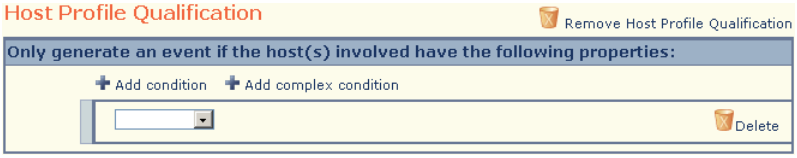
Access: P&R Admin/Admin

1. Build a compliance rule that triggers on an RNA event.

   For more information, see Creating Rules for Compliance Policies on page 1214.

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 compliance rule.

   Because this host profile qualification is part of a compliance rule based on an RNA event, the only available category is **RNA 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 compliance rule triggers, a host profile qualification, or a flow tracker. Within compliance rule triggers, the categories further depend on what kind of event is the basis for your compliance rule.

Keysight, EX-1004
Page 1254

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 1258.

For more information on the syntax for building compliance rule trigger criteria, see:

- Syntax for Flow Events on page 1218
- Syntax for Intrusion Events on page 1221
- Syntax for RNA Events on page 1223
- Syntax for RUA Events on page 1226
- Syntax for Host Input Events on page 1227
- Syntax for Traffic Profile Changes on page 1229

For more information on the syntax for building host profile qualifications, user qualifications, and flow trackers, see:

- Syntax for Host Profile Qualifications on page 1234
- Syntax for User Identity Qualifications on page 1249
- Syntax for Flow Trackers on page 1244
- Syntax for Flow Tracker Events on page 1247

Adding and Linking Conditions

Platform: DC

Component: None

You can create simple compliance rule triggers, flow trackers, and host profile qualifications, user qualifications, or you can create more elaborate constructs by combining and nesting conditions.

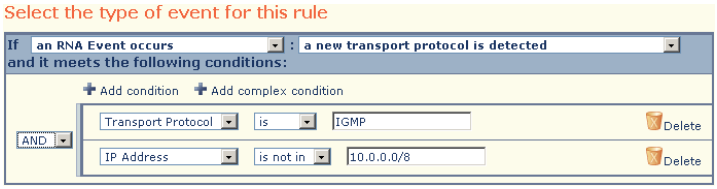
When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together.

- 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.

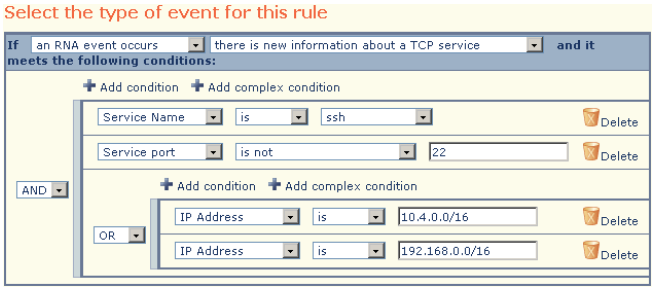
For example, the following compliance rule trigger criteria contains two conditions, linked by **AND**. This means that the rule triggers if both conditions are

Keysight, EX-1004
Page 1255

true, that is, if a host with an IP address that is not in the 10.x.x.x subnet transmits an IGMP message.



In contrast, the following rule, which detects ssh activity on a nonstandard 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.



This rule triggers if the ssh service is detected on a nonstandard port; the first two conditions demand that the service 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	Service name is ssh
B	Service port is not 22
C	IP Address is in 10.4.0.0/8
D	IP Address is in 196.168.0.0/16

Keysight, EX-1004
Page 1256

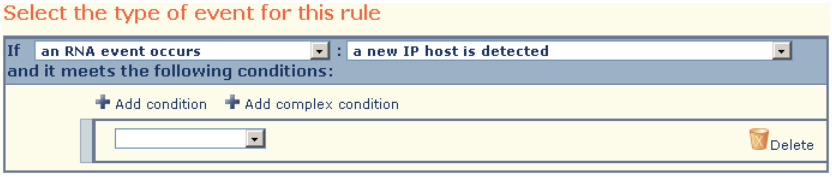
To add a single condition:

Access: P&R Admin/Admin

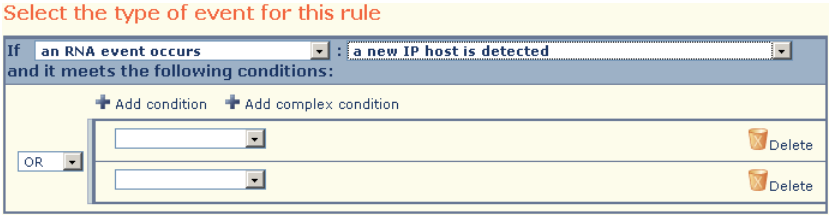
- To add a single condition, click **Add condition** above the current condition.

  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:



  The result is:



Keysight, EX-1004
Page 1257

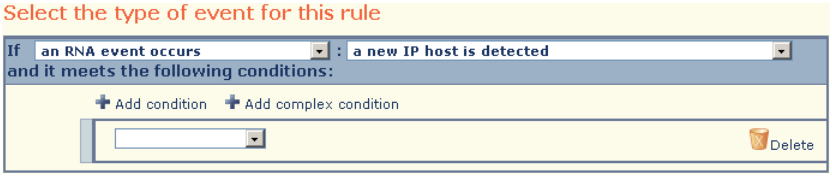
To add a complex condition:

Access: P&R Admin/Admin

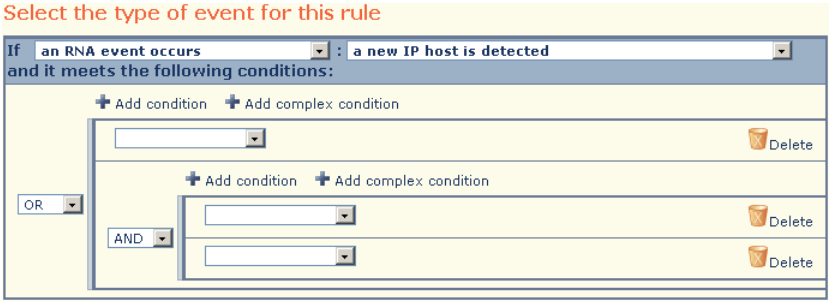
▶ Click **Add complex condition** above the current condition.

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: P&R Admin/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

Platform: DC

Component: None

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.

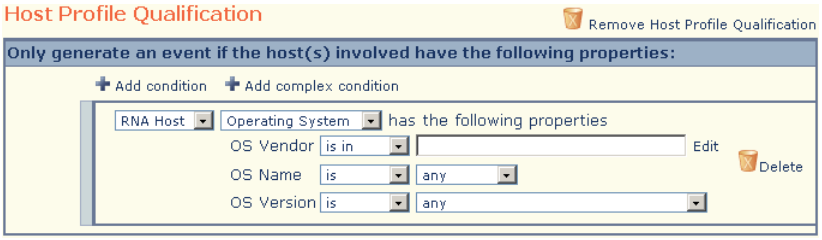
Keysight, EX-1004
Page 1258

To include multiple values in one condition:

Access: P&R Admin/Admin

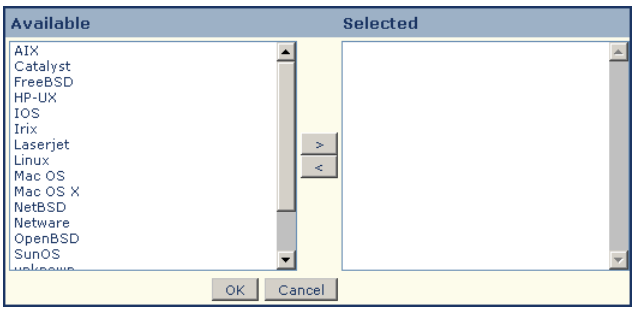
1. Build a condition, choosing **is in** or **is not in** as the operator.

   The drop-down list changes to a text field.



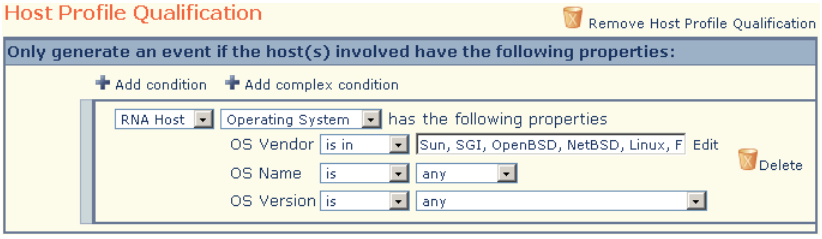
2. Click anywhere in the text field or on the **Edit** link.

   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.

   The Create Rule page appears again. Your selections appear in the value field of your condition.



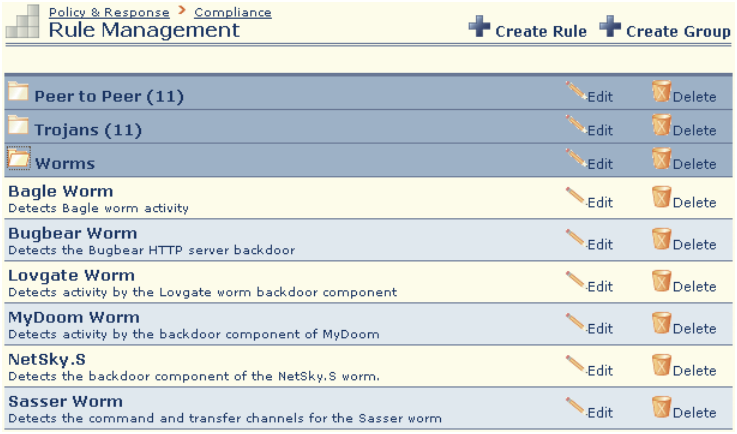
Keysight, EX-1004
Page 1259

Managing Rules for Compliance Policies

Platform: DC

Component: None

Use the Rule Management page to manage compliance rules used within compliance policies.



You can create, modify, and delete rules. You can also create rule groups to help you organize compliance rules. For more information on modifying rules, deleting rules, and creating rule groups, see:

- Modifying a Rule on page 1260
- Deleting a Rule on page 1261
- Creating a Rule Group on page 1261
- Deleting a Rule Group on page 1262

For more information on creating rules, see Creating Rules for Compliance Policies on page 1214.

Modifying a Rule

Platform: DC

Component: None

Use the following procedure to modify an existing compliance rule.

To modify an existing rule:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Rule Management**.

   The Rule Management page appears.

Keysight, EX-1004
Page 1260

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 **Edit**.

   The Create Rule page appears.
4. Make modifications as necessary and click **Update Rule**.

   The rule is updated.

Deleting a Rule

Platform: DC

Component: None

You cannot delete compliance rules that you are using in one or more 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 Modifying a Compliance Policy on page 1270.

To delete an existing rule:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Rule Management**.

   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 **Delete**.

   The rule is deleted.

Creating a Rule Group

Platform: DC

Component: None

Create rule groups to help you organize compliance 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 compliance rules; you cannot assign a group of rules to a compliance 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 Compliance Policies on page 1214
- Modifying a Rule on page 1260.

For information on deleting a rule group, see Deleting a Rule Group on page 1262.

Keysight, EX-1004
Page 1261

To create a rule group:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Rule Management**.
   The Rule Management page appears.
2. Click **Create Group**.
   The Groups page appears.



3. In the **Group Name** field, type a name for the group.
4. Click **Add Group**.
   The group is added.

Deleting a Rule Group

Platform: DC

Component: None

When you delete a rule group, rules that were in the group are **not** deleted. Rather, they merely become ungrouped.

To delete a rule group:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Rule Management**.
   The Rule Management page appears.
2. Next to the rule group you want to delete, click **Delete**.
3. Confirm that you want to delete the rule group.
   The group is deleted and rules that were in that group become ungrouped.

Creating Compliance Policies

Platform: DC

Component: None

After you create compliance rules or compliance white lists (or both), and, optionally, alert responses and remediations, you can use them to build compliance policies.

When your network traffic meets the criteria specified in a compliance rule or white list in an active policy, the Defense Center generates either a compliance event or white list event. It also launches any responses you assigned to the rule

Keysight, EX-1004
Page 1262

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 compliance rules, compliance white lists, and responses you can use to build a compliance policy, see the following sections:

- Creating Rules for Compliance Policies on page 1214
- Creating Compliance White Lists on page 1167
- Configuring Responses for Compliance Policies on page 1281.

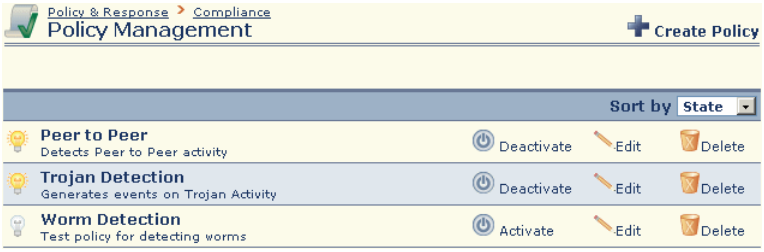
TIP! Optionally, you can create a skeleton policy and modify it later to add rules and responses.

To create a compliance policy:

Access: P&R Admin/Admin

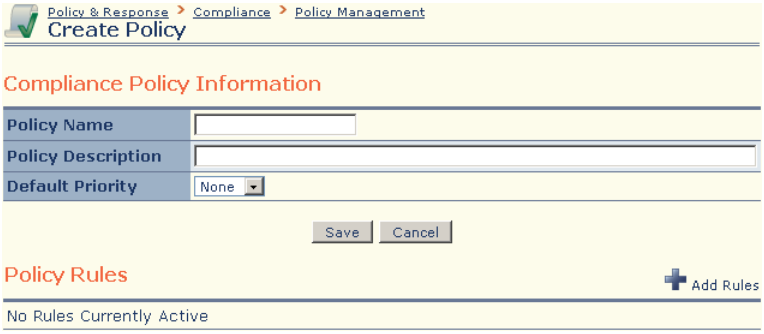
1. Select **Policy & Response > Compliance > Policy Management**.

   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 1264.

Keysight, EX-1004
Page 1263

4. Add one or more rules or white lists to the compliance policy.

   See Adding Rules and White Lists to a Compliance Policy on page 1265.

5. Optionally, set rule and white list priorities.

   See Setting Rule and White List Priorities on page 1266.

6. Optionally, add responses to the rules or white lists you added.

   Adding Responses to Rules and White Lists on page 1267.

7. Click **Save**.

   The policy is saved.

IMPORTANT! You must activate the policy before it can generate compliance and white list events and launch responses to policy violations. For more information, see Managing Compliance Policies on page 1269.

Providing Basic Policy Information

Platform: DC

Component: None

You must give each policy an identifying name. Optionally, you can add a short description to the policy.



You can also assign a user-defined priority to your policy. If your compliance policy is violated, the resultant compliance 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 Compliance Policy on page 1265.

To provide basic policy information:

Access: P&R Admin/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.

Keysight, EX-1004
Page 1264

3. From the **Default Priority** 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 Compliance Policy on page 1265.

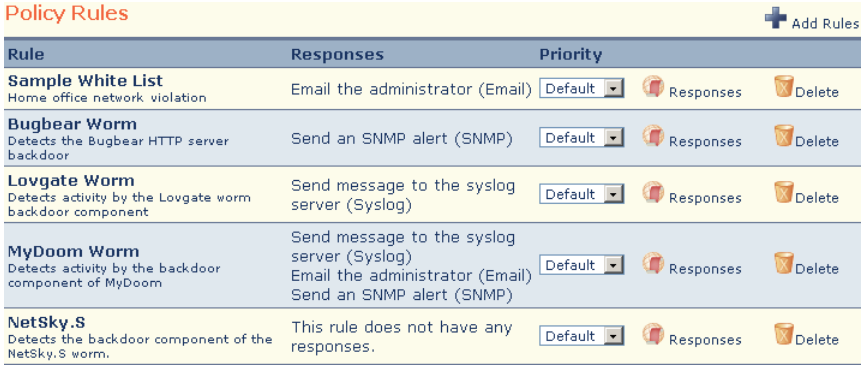
Adding Rules and White Lists to a Compliance Policy

Platform: DC

Component: None

A compliance policy contains one or more compliance rules or white lists. When any rule or white list in a policy is violated, RNA 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 compliance policy composed of a compliance white list and a set of compliance rules, configured with a variety of responses.



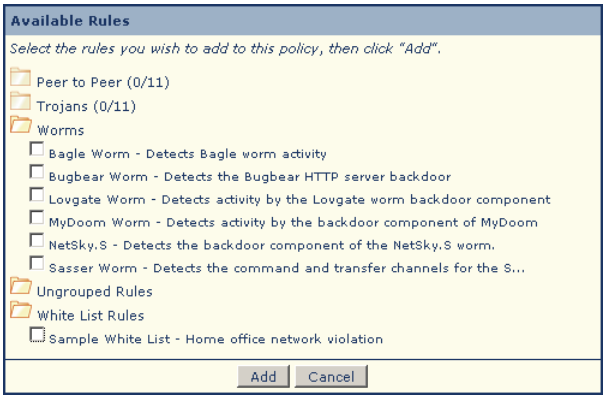
Keysight, EX-1004
Page 1265

To add rules or white lists to a compliance policy:

Access: P&R Admin/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 1266.

Setting Rule and White List Priorities

Platform: DC

Component: None

You can assign a user-defined priority to each compliance rule or compliance white list in your compliance 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.

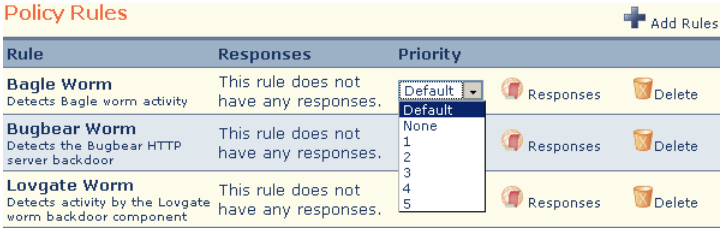
For example, consider a policy where the policy itself has a priority of 1 and its rules or white lists are set with the default priority, with the exception of one rule given a priority of 3. If the priority 3 rule triggers, the resulting compliance event shows 3 as its priority value. If other rules or white lists in the policy trigger, the resulting events show 1 as their priority values, retained from the policy's priority.

Keysight, EX-1004
Page 1266

To set rule or white list priorities:

Access: P&R Admin/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 1267.

Adding Responses to Rules and White Lists

Platform: DC

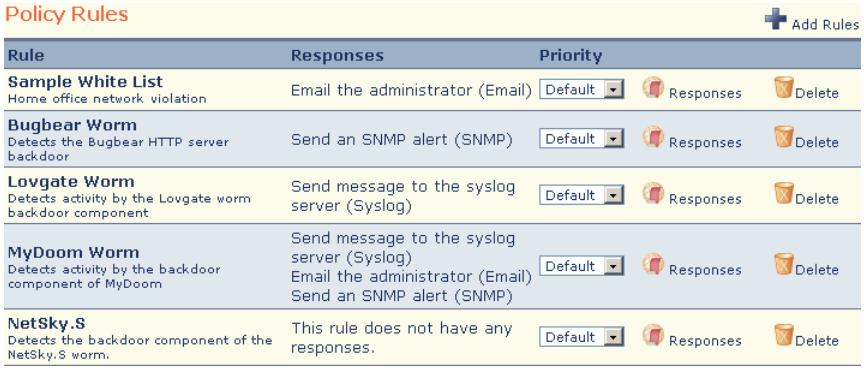
Component: RNA

Within a compliance 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, RNA 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 Responses for Compliance Policies on page 1281.

IMPORTANT! Do **not** assign an Nmap remediation as a response to a compliance rule that triggers on a traffic profile change. The remediation will not launch.

Keysight, EX-1004
Page 1267

The following graphic shows a compliance policy composed of a compliance white list and a set of compliance rules, configured with a variety of responses.

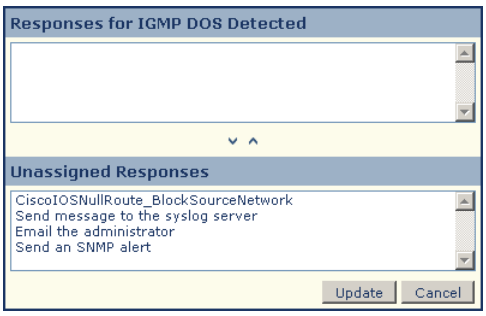


To add responses to rules and white lists:

Access: P&R Admin/Admin

1. On the Create Policy page, next to a rule or white list where you want to add responses, click **Responses**.

   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.

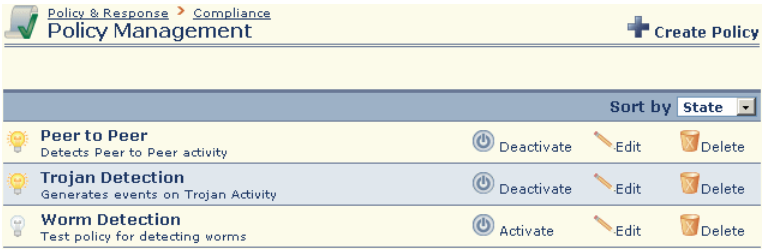
Keysight, EX-1004
Page 1268

Managing Compliance Policies

Platform: DC

Component: None

You manage compliance policies on the Policy Management page. You can create, modify, sort, activate, deactivate, and delete policies.



A policy's status appears with its name. If the light bulb icon next to the policy name is lit, the policy is active. If it is dark, the policy is inactive. If you want the policy to generate compliance 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 compliance 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 compliance policy (if you deactivated it) or add the white list to another active compliance policy (if you deleted the white list from a compliance policy or deleted the policy itself). Note that the re-evaluation 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 RNA as a Compliance Tool on page 1156.

For more information on managing compliance policies, see:

- Activating and Deactivating Compliance Policies on page 1270
- Modifying a Compliance Policy on page 1270
- Deleting a Compliance Policy on page 1271

Keysight, EX-1004
Page 1269

For information on creating new policies, see Creating Compliance Policies on page 1262.

Activating and Deactivating Compliance Policies

Platform: DC

Component: None

Use the following procedure to either activate or deactivate a compliance policy.

To activate or deactivate a policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Policy Management**.

   The Policy Management page appears.

2. You have two options:
   - To **activate** a policy so that it will generate events and launch associated responses when triggered, click **Activate** next to the policy.
   - To **deactivate** a policy so that it will **not** generate events or launch associated responses when triggered, click **Deactivate** next to the policy.

Modifying a Compliance Policy

Platform: DC

Component: None

Use the following procedure to modify a compliance policy.

To modify a policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Policy Management**.

   The Policy Management page appears.

2. Click **Edit** next to the policy.

   The Create Policy page appears. See Creating Compliance Policies on page 1262 for information on the various configurations you can change. To remove a rule or white list from a compliance policy, on the Create Policy page, click **Delete** 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.

Keysight, EX-1004
Page 1270

Deleting a Compliance Policy

Platform: DC

Component: None

Use the following procedure to delete a compliance policy.

To delete a policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > Compliance > Policy Management**.
   The Policy Management page appears.
2. Click **Delete** next to the policy you want to delete.
   The policy is deleted.

Working with Compliance Events

Platform: DC/MDC

Component: None

When a compliance rule within an active compliance policy triggers, the Defense Center generates a compliance event and logs it to the database. You can search, view, and delete compliance events.

IMPORTANT! When a compliance white list within an active policy triggers, the Defense Center generates a white list event. For more information, see Working with White List Events on page 1197.

TIP! For information on configuring the number of compliance and white list events saved in the database, see Configuring Database Event Limits on page 1747.

For more information, see the following sections:

- Viewing Compliance Events on page 1272
- Understanding the Compliance Events Table on page 1274
- Searching for Compliance Events on page 1276

Keysight, EX-1004
Page 1271

Viewing Compliance Events

Platform: DC/MDC

Component: None

You can view a table of compliance events, and then manipulate the event view depending on the information you are looking for.

The page you see when you access compliance events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of compliance 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 1618.

The Compliance Event Actions table below describes some of the specific actions you can perform on an compliance events workflow page.

Compliance 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 1495.
sort and constrain events on the current workflow page	find more information in Sorting Drill-down Workflow Pages on page 1612.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1613.
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 1593.
learn more about the columns that appear	find more information in Understanding the Compliance Events Table on page 1274.
modify the time and date range for displayed events	find more information in see Setting Event Time Constraints on page 1597. 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.

Keysight, EX-1004
Page 1272

Compliance Event Actions (Continued)

To...	You can...
drill down to the next page in the workflow, constraining on a specific value	use one of the following methods: • 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
delete compliance events from the system	use one of the following methods: • To delete some events, select the checkboxes 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. See Purging the RNA and RUA Databases on page 2014 for information on deleting all RNA events from the database.
navigate to other event views to view associated events	find more information in Navigating between Workflows on page 1614.
temporarily use a different workflow	click **Workflows**. For more information, see Selecting Workflows on page 1590.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for compliance events	click **Search**. For more information, see Searching for Compliance Events on page 1276.

Keysight, EX-1004
Page 1273

To view compliance events:

Access: Any Analyst/Admin

- Select **Analysis & Reporting > Compliance > Compliance Events**.

  The first page of the default compliance events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

TIP! If you are using a custom workflow that does not include the table view of compliance events, from the **Workflows** menu on the toolbar, select **Compliance Events**.

Understanding the Compliance Events Table

Platform: DC/MDC

Component: None

When a compliance rule triggers, the Defense Center generates a compliance event. The fields in the compliance events table are described in the Compliance Event Fields table.

Compliance Event Fields

Field	Description
Time	The date and time that the compliance event was generated.
Impact Flag	The impact flag assigned to the compliance event based on the correlation between IPS data, RNA data, and vulnerability information. For more information, see Using Impact Flags to Evaluate Events on page 307.

Keysight, EX-1004
Page 1274

Compliance Event Fields (Continued)

Field	Description
Inline Result	One of: • a black down arrow, indicating that IPS dropped the packet that triggered the intrusion rule • a gray down arrow, indicating that IPS would have dropped the packet in an inline 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 Note that IPS 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 435, Setting Drop Behavior in an Inline Deployment on page 358, and Tap Mode on page 227.
Source IP	The IP address of the source host in the event that triggered the policy violation.
Destination IP	The IP address of the destination host in the event that triggered the compliance rule.
Source User	The name of the user logged in to the source host in the event that triggered the policy violation
Destination User	The name of the user logged in to the destination host in the event that triggered the policy violation
Source Port	The source port associated with the event that triggered the compliance rule.
Destination Port	The destination port associated with the event that triggered the compliance rule.
Description	The description of the compliance event. The information in the description depends on how the rule was triggered. For example, if the rule was triggered by an operating system information update event, the new operating system name and confidence level appears. If a new TCP service information event triggered the rule, the port, service name, confidence level, and any service subtypes appear. If the rule was triggered by a new TCP service event, the service port appears.
Policy	The name of the policy that was violated.
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.

Keysight, EX-1004
Page 1275

Compliance Event Fields (Continued)

Field	Description
Src Host Criticality	The user-assigned host criticality of the source host involved in the compliance event: None, Low, Medium, or High. Note that only compliance events generated by rules based on RNA events, host input events, or flow events contain a source host criticality. For more information on host criticality, see Assigning a Host Criticality Value to a Host on page 986.
Dst Host Criticality	The user-assigned host criticality of the destination host involved in the compliance event: None, Low, Medium, or High. Note that only compliance events generated by rules based on RNA events or flow events contain a destination host criticality. For more information on host criticality, see Assigning a Host Criticality Value to a Host on page 986.
Detection Engine	The name of the detection engine that generated the event that triggered the compliance rule.
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 compliance events table, see the following:

- Viewing Compliance Events on page 1272
- Searching for Compliance Events on page 1276

Searching for Compliance Events

Platform: DC/MDC

Component: None

You can search for specific compliance events. You may want to create searches customized for your network environment, then save them to re-use later. The

Keysight, EX-1004
Page 1276

search criteria you can use are described in the Compliance Event Search Criteria table.

Compliance Event Search Criteria

Field	Search Criteria Rules
Policy	Enter all or part of the name of the policy you want to search for.
Rule	Enter all or parts of the name of the rule you want to search for.
Description	Enter all or part of the compliance event description. The information in the description depends on how the rule was triggered. For example, if the rule was triggered by an operating system information update event, the new operating system name and confidence level appears. If a new TCP service information event triggered the rule, the port, service name, confidence level, and any service subtypes appear. If the rule was triggered by a new TCP service event, the service port appears.
Priority	Specify the priority of the compliance event, which is determined by the priority of the compliance rule whose triggering generated the event, or of the compliance policy that was violated. Enter none for no priority. For information on setting compliance rule and policy priorities, see Providing Basic Policy Information on page 1264 and Setting Rule and White List Priorities on page 1266.
Source IP	Specify the IP address of the source host in the event that triggered the compliance rule that generated the compliance event. See Specifying IP Addresses in Searches on page 1557 for detailed information about how to enter IP addresses.
Destination IP	Specify the destination IP address of the host or hosts whose compliance events you want to view. Note that only compliance events generated by rules based on intrusion events and flow events contain a destination IP.See Specifying IP Addresses in Searches on page 1557 for detailed information about how to enter IP addresses.
Source User	Specify the name of the user logged in to the source host in the event that triggered the policy violation
Destination User	Specify the name of the user logged in to the destination host in the event that triggered the policy violation
Source Port	Specify the source port associated with the event that triggered the compliance rule. See Specifying Ports in Searches on page 1558 for information about the syntax for specifying ports.
Destination Port	Specify the destination port associated with the event that triggered the compliance rule. See Specifying Ports in Searches on page 1558 for information about the syntax for specifying ports.

Keysight, EX-1004
Page 1277

Compliance Event Search Criteria (Continued)

Field	Search Criteria Rules
Source/Destination IP	Specify the source or destination IP address of the host or hosts whose compliance events you want to view.
Impact Flag	For Defense Centers, specify the impact flag assigned to the compliance event based on the correlation between IPS and RNA data. Valid values are `red`, `orange`, `yellow`, `blue`, and `gray`. For more information, see Using Impact Flags to Evaluate Events on page 307.
Inline Result	Type either: • `dropped`, to specify whether the packet was 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 IPS 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 435, Setting Drop Behavior in an Inline Deployment on page 358, and Tap Mode on page 227.
Src Host Criticality	Specify the host criticality of the source host involved in the compliance event: `None`, `Low`, `Medium`, or `High`. Note that only compliance events generated by rules based on RNA events contain a source host criticality. For more information on host criticality, see Assigning a Host Criticality Value to a Host on page 986.
Dst Host Criticality	Specify the host criticality of the destination host involved in the compliance event: `None`, `Low`, `Medium`, or `High`. Note that only compliance events generated by rules based on RNA events contain a destination host criticality. For more information on host criticality, see Assigning a Host Criticality Value to a Host on page 986.
Detection Engine	Specify the name of the detection engine, detection engine group, sensor, or sensor group that generated the event that caused the Defense Center to generate the compliance events you want to view. For more information, see Specifying Detection Engine/Appliance Names on page 1559.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

Keysight, EX-1004
Page 1278

To search for compliance events:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Searches > Compliance Events**.

   The Compliance Events search page appears.



   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 created automatically when you save the search.

3. Enter your search criteria in the appropriate fields, as described in the Compliance Event Search Criteria table. If you enter multiple criteria, the search returns only the records that match all the criteria.

Keysight, EX-1004
Page 1279

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 compliance events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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.

Keysight, EX-1004
Page 1280

Chapter 29

Configuring Responses for Compliance Policies

You can configure the Defense Center to launch responses to network discovery events (also called RNA events), to intrusion events that have specific impact flag values, and to compliance policy violations.

The most basic kind of response you can launch is an alert. Alerts notify you, in one of three ways, of an RNA event, impact flag, or compliance policy violation. You can receive alerts via email, an simple network management protocol (SNMP) trap server, or a syslog server.

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 compliance 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 also 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 compliance 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 compliance policy is blocked. For more information about writing your own remediation modules, refer to the *Sourcefire Remediation API Guide*.

You can group alerts and remediations so that a compliance policy violation triggers all of the responses within the group.

Keysight, EX-1004
Page 1281

You must use a Defense Center to create alerts and assign them to RNA events or impact flags, or to configure and use remediations with RNA events. You can also use a Master Defense Center to create alerts and assign them to intrusion events with specific impact flags.

For more information, see:

- Configuring Alerts on page 1282
- Assigning Notification Alerts to RNA Events on page 1293
- Configuring External Alerting on Impact Flags on page 1295
- Configuring Remediations on page 1296
- Working with Remediation Status Events on page 1343
- Configuring Response Groups on page 1351

Configuring Alerts

Platform: DC/MDC

Component: RNA

You can create three types of alerts: syslog, email, and SNMP. RNA events, intrusion events that have specific impact flag values, and compliance rules can use these alerts. The information you receive in an alert depends on the type of event that triggered the alert (or, in the case of alerts launched from within compliance policies, the type of event that triggered the compliance policy violation). The Alert Information table describes the information you receive in alerts.

Keysight, EX-1004
Page 1282

If you configure an alert as a response to a compliance rule that contains a flow tracker, the alert information you receive is the same as that for alerts on traffic profile changes, even if the compliance rule is based on a different kind of event.

Alert Information

Event Type	Alert Information
RNA or host input event	• the event type • the name of the detection engine that generated the event • the date and time the alert was generated • the IP address of the host involved in the event • a description of the event (for example, if you want to be alerted when RNA detects a new network protocol, the alert includes the name of the protocol) For more information on RNA and host input event types, see Understanding RNA Network Discovery Event Types on page 1005 and Understanding RNA Host Input Event Types on page 1010.
RUA event	• the date and time the alert was generated • the event type • the user associated with the event • the IP address of the host involved in the event • a description of the event • the name of the detection engine that generated the event For more information on RUA events, see Working with RUA Events on page 1500.
flow event	none
traffic profile change	• the date and time the alert was generated • the number of connections open when the compliance rule triggered • the top 5 initiators and the number of connections open for each initiator when the compliance rule triggered • the top 5 responders and the number of connections open for each responder when the compliance rule triggered

Keysight, EX-1004
Page 1283

Alert Information (Continued)

Event Type	Alert Information
white list event	• the date and time the alert was generated • the IP address of the host involved in the event • the port involved in the event, if any • a description of the event
intrusion event	• the rule name • the impact flag value of the event • the name of the detection engine that generated the event • the date and time the alert was generated • a description of the event

Before you can assign alerts to RNA events or compliance policies, you must create them on the Alerts page.



When you create an alert, its type and status appear with its name. The light bulb icon next to the alert indicates whether the alert is active. If you want to launch an alert in response to an event or to a compliance rule, you must activate it first. If the light bulb icon is lit, the alert is active. If it is dark, the alert is inactive. For more information, see Activating and Deactivating Alerts on page 1293.

The term `In Use` indicates an alert that is used in a compliance policy or assigned to an RNA event or impact flag, while `Not Used` indicates an alert that is not in use. Also, you can sort alerts by state (active versus inactive) or alphabetically by name using the **Sort by** drop-down list.

For more information, see:

- Creating Alerts on page 1285
- Modifying Alerts on page 1292
- Deleting Alerts on page 1292
- Activating and Deactivating Alerts on page 1293

Keysight, EX-1004
Page 1284

After you create an alert, launch it in response to an RNA event, an impact flag, or a rule within a compliance policy. For more information, see:

- Assigning Notification Alerts to RNA Events on page 1293
- Configuring External Alerting on Impact Flags on page 1295
- Adding Responses to Rules and White Lists on page 1267

Creating Alerts

Platform: DC/MDC

Component: RNA

You can create three types of alerts to assign to events, impact flags, and rules within compliance policies. For more information, see:

- Creating Syslog Alerts on page 1285
- Creating Email Alerts on page 1289
- Creating SNMP Alerts on page 1290

Creating Syslog Alerts

Platform: DC/MDC

Component: RNA

You can configure the Defense Center to send messages to an external syslog server.

As part of the alert you can specify the severity and facility associated with the syslog messages to ensure that they are processed properly by the remote 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 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 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.

You can select any type of facility when creating a syslog alert, but you should 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

Keysight, EX-1004
Page 1285

system (if you are logging syslog messages to a UNIX- or Linux-based system) should indicate which facilities are saved to which log files on the server.

WARNING! The computer you configure to receive the syslog messages must be set up to accept remote messages. Otherwise, the Defense Center may send syslog messages to the host, but they will not be accepted.

The Available Syslog Facilities table lists the facilities you can select.

Available Syslog Facilities

Facility	Description
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
USER	A message generated by a user-level process.
MAIL	A message generated by a mail system.
DAEMON	A message generated by a system daemon.
AUTH	A message associated with security and authorization.
SYSLOG	A message generated by the syslog daemon.
LPR	A message generated by the printing subsystem.
NEWS	A message generated by the network news subsystem.
UUCP	A message generated by the UUCP subsystem.
CRON	A message generated by the clock daemon.
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
FTP	A message generated by the FTP daemon.
NTP	A message generated by the NTP daemon.
AUDIT	A message generated by the audit subsystem.
ALERT	An alert message.

Keysight, EX-1004
Page 1286

Available Syslog Facilities (Continued)

Facility	Description
CLOCK	A message generated by the clock daemon. **IMPORTANT!** Some systems use the CLOCK facility and others use CRON. You should be familiar with the configuration of the system where you want to send syslog messages.
LOCAL0-LOCAL7	A message generated by an internal process.

You can select from the following standard syslog severity levels:

Syslog Severity Levels

Level	Description
EMERG	A panic condition broadcast to all users.
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
ERROR	An error condition.
WARN	Warning messages.
NOTICE	Conditions that are not error conditions, but require attention.
INFO	Informational messages.
DEBUG	Messages that contain debugging information.

To create a syslog alert:

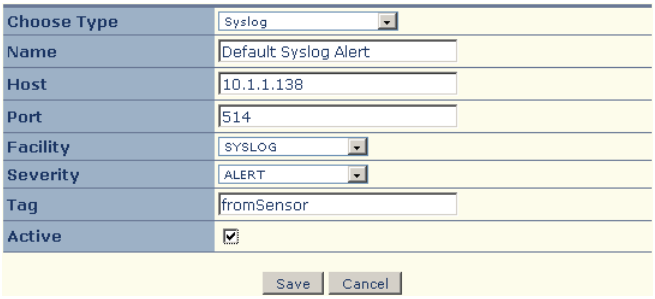
Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.
   The Alerts page appears.
2. Click **Create Alert**.
   The Create Alert page appears.

Keysight, EX-1004
Page 1287

3. From the **Choose Type** list in the Configure Alerts section, select **Syslog**.

   The Create Alert page expands.



4. In the **Name** field, type the name you want to use to identify the saved response.
5. In the **Host** field, type the IP address or host name of the server that will receive the syslog alerts.

   TIP! If you want to specify multiple syslog servers, separate each IP address with a comma.

6. In the **Port** field, type the port the server uses for syslog messages. By default, this value is 514.
7. From the **Facility** list, select the facility you would like to use for syslog alerts.

   See the Available Syslog Facilities table on page 1286 for a list of all available facilities.
8. From the **Severity** list, select a severity.

   See the Syslog Severity Levels table on page 1287 for a list and description of all available severities.
9. In the **Tag** field, type the tag name (using alphanumeric characters only) that you want to appear with the syslog message.

   For example, if you wanted all messages sent to syslog to be preceded with `FromDC`, type `FromDC` in the field.

   IMPORTANT! You can use only alphanumeric characters in tag names. You **cannot** use spaces and underscores.

10. Select **Active** to activate the alert so you can use it within a compliance policy or as a response to an RNA event or an impact flag.
11. Click **Save**.

    The syslog alert is saved.

Keysight, EX-1004
Page 1288

Creating Email Alerts

Platform: DC/MDC

Component: RNA

You can configure the Defense Center to send an alert via email in response to an RNA event, impact flag, or a compliance policy violation.

WARNING! Before you can send email alerts, you must enable DNS and the Defense Center must be able to reverse resolve its own IP address.

To create an email alert:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.

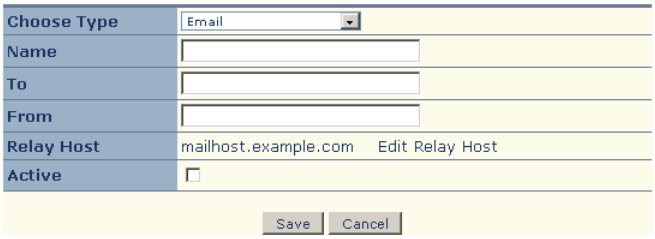
   The Alerts page appears.

2. Click **Create Alert**.

   The Create Alert page appears.

3. From the **Choose Type** list in the Configure Alerts section, select **Email**.

   The Create Alert page expands.



4. In the **Name** field, type the name you want to use to identify the email alert.
5. In the **To** field, type the email address (or multiple addresses separated by commas) that you want to receive notifications.
6. In the **From** field, type the name that will appear as the sender of the notification email.
7. In the **Relay Host** field, verify that the name or IP address of the mail server is the one that you want to use to transmit the alert.

   The relay host is configured as part of the system policy. If you need to change the mail server, or if you have not yet configured a relay host, click **Edit Relay Host**. The System Policy page appears in a pop-up window. For detailed information about editing the system policy to configure a relay host, see Configuring a Mail Relay Host and Notification Address on page 1751. Note that you must apply the system policy after you edit it for your changes to take effect.

Keysight, EX-1004
Page 1289

8. Select **Active** to activate the alert so that you can use it within a compliance policy or as a response to an RNA event or an impact flag.
9. Click **Save**.

   The email alert is saved.

Creating SNMP Alerts

Platform: DC/MDC

Component: RNA

You can configure the Defense Center to send an alert to an SNMP trap server in response to an RNA event, impact flag, or compliance policy violation. You can configure SNMP alerts using SNMPv1, SNMPv2, or SNMPv3.

When you use SNMPv3, RNA uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message. Currently, this Engine ID value is the hexadecimal version of the Defense Center's IP address with `01` at the end of the string. For example, if the Defense Center sending the SNMP alert has an IP address of `172.168.1.50`, the Engine ID is `0xAC10013201` or if the Defense Center had an IP address of `10.1.1.77`, `0x0a01014D01` is used as the Engine ID.

TIP! 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:

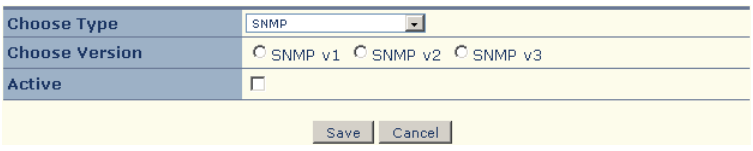
Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.

   The Alerts page appears.
2. Click **Create Alert**.

   The Create Alert page appears.
3. From the **Choose Type** list in the Configure Alerts section, select **SNMP**.

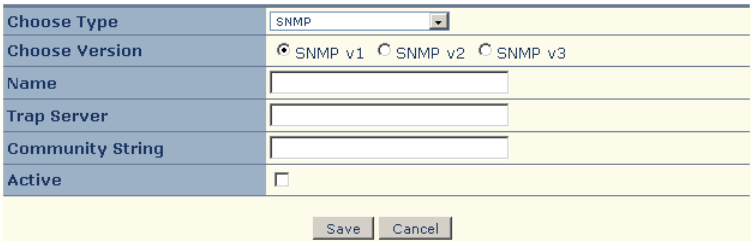
   The Create Alert page expands.



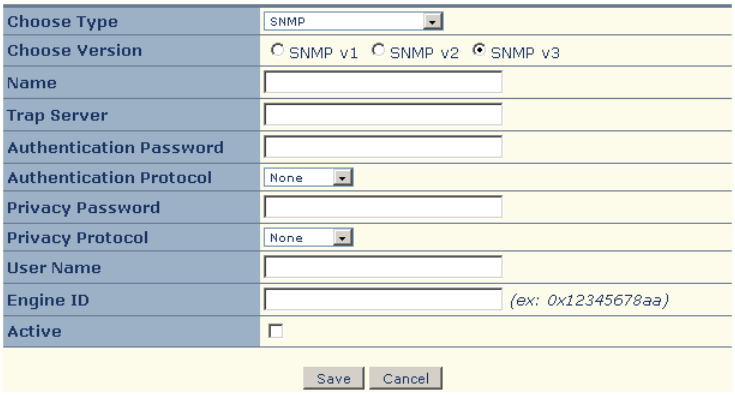
Keysight, EX-1004
Page 1290

4. From the **Choose Version** section, select the SNMP version you want to use.

   If you selected SNMP v1 or SNMP v2, their options appear.



   If you selected SNMP v3, the SNMP v3 options appear.



5. In the **Name** field, type the name that you want to use to identify the SNMP response.
6. In the **Trap Server** field, type the name or IP address of the SNMP trap server.
7. If you selected SNMP v1 or SNMP v2, type the SNMP community name in the **Community String** field and skip to step 14.
8. If you selected SNMP v3, in the **Authentication Password** field, type the password required for authentication with the SNMP server.
9. From the **Authentication Protocol** drop-down list, select the protocol you want to use for authentication.
10. In the **Privacy Password** field, type the SNMP privacy key required by the SNMP server.
11. From the **Privacy Protocol** list, select **None** to use no privacy protocol or **DES** to use Data Encryption Standard as the privacy protocol.
12. In the **User Name** field, type the name of the user that you want to authenticate with the SNMP server.
13. In the **Engine ID** field, type the identifier for the SNMP engine.

Keysight, EX-1004
Page 1291

14. Select **Active** to activate the alert so that you can use it within a compliance policy or as a response to an RNA event or an impact flag.
15. Click **Save**.

    The SNMP alert is saved.

Modifying Alerts

Platform: DC/MDC

Component: RNA

Use the following procedure to modify an alert.

To modify an alert:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.

   The Alerts page appears.
2. Click **Edit** next to the alert you want to modify.

   The Create Alert page appears.
3. Make changes as needed.
4. You have two options.
   - Click **Save** to modify the existing alert. If the alert is active and in use, the changes you made take effect immediately.
   - Click **Save as New** to create a new alert with the information you specified. Note that to create a new alert, you must change the name.

Deleting Alerts

Platform: DC/MDC

Component: RNA

You can delete any alert as long as it is not in use, either within a compliance policy or in response to an RNA event or impact flag.

To delete an alert:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.

   The Alerts page appears.
2. On the Alerts page, click **Delete** next to the alert you want to delete.

   The alert is deleted.

Keysight, EX-1004
Page 1292

Activating and Deactivating Alerts

Platform: DC/MDC

Component: RNA

If you want to temporarily disable an alert without deleting it, you can deactivate it. This leaves the alert on the system, but it will not be launched in response to an RNA event, an impact flag, or a compliance policy violation. To enable the alert again, you must activate it.

IMPORTANT! Even you deactivate an alert, it is still considered in use and therefore cannot be deleted unless you remove it in all of the places it is used.

To deactivate an alert:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.
   The Alerts page appears.
2. Click **Deactivate** next to the alert you want to deactivate.

To activate an alert:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Alerts**.
   The Alerts page appears.
2. Click **Activate** next to the alert you want to activate.

Assigning Notification Alerts to RNA Events

Platform: DC

Component: RNA

After you create alerts as described in Configuring Alerts on page 1282, you can configure the Defense Center to send the alerts when it generates specific RNA events.

IMPORTANT! You cannot assign response groups to RNA events. You can use response groups only within compliance policies.

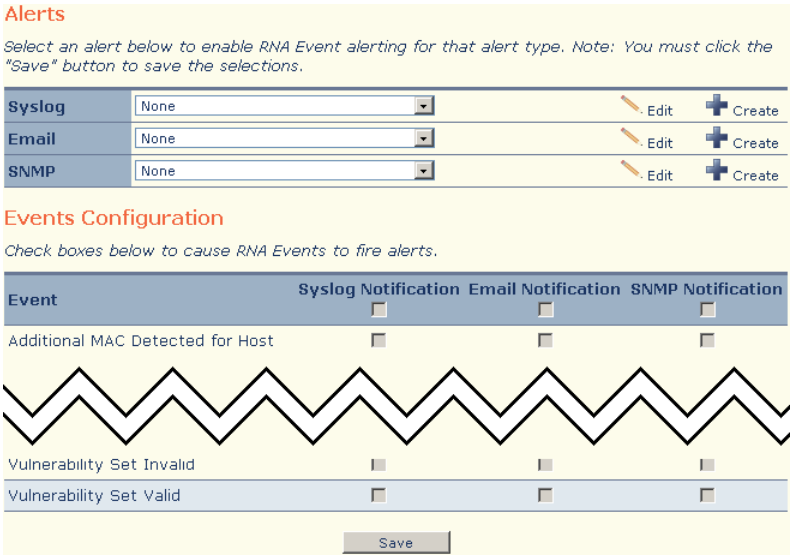
Keysight, EX-1004
Page 1293

To assign alerts to events:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > RNA Event Alerts**.

   The RNA Event Alerts page appears.



2. In the Alerts section, select the alert you want to use for each alert type.

   TIP! You can create alerts by clicking **Create** next to the type of alert you want to create. For more information, see Creating Alerts on page 1285.

3. Next to each event type, select the check boxes that best match the response you want to occur when the event is generated. You can select one or more of the following:

   - **Syslog Notification** to launch the syslog alert you chose.
   - **Email Notification** to launch the email alert you chose.
   - **SNMP Notification** to launch the SNMP alert you chose.

   To select all events in a column, select the check box that appears in the column header.

   TIP! For more information about each event type, see Understanding RNA Network Discovery Event Types on page 1005 and Understanding RNA Host Input Event Types on page 1010.

4. Click **Save**.

Keysight, EX-1004
Page 1294

Configuring External Alerting on Impact Flags

Platform: DC/MDC

Component: RNA + IPS

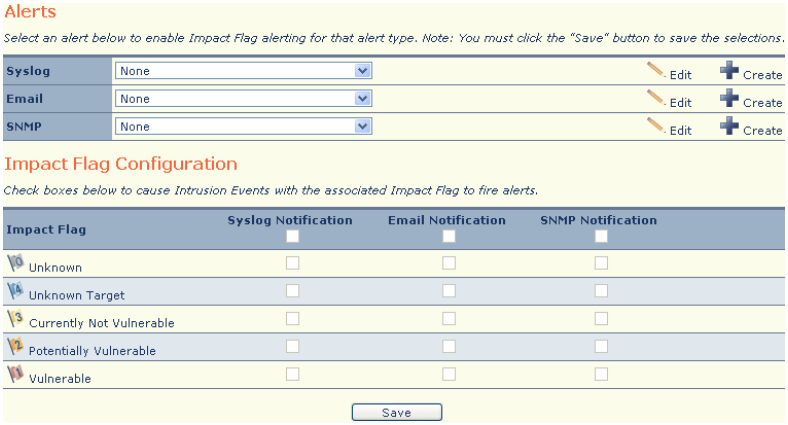
If you are managing sensors that are licensed for IPS and RNA, you can configure the system to notify you of events that have specific impact flag values. This allows you to receive notification of events based on how directly they impact your network.

To configure external alerting for events based on impact:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Impact Flag Alerts**.

   The Impact Flag Alerts page appears.



2. In the Alerts section, select the alert configuration you want to use for each alert type.

   TIP! You can create alerts by clicking **Create** next to the type of alert you want to create. For more information, see Creating Alerts on page 1285.

3. In the Impact Configuration section, select the check box for each type of notification you want to receive for each impact flag. You can select the following impact flags:
   - Unknown (impact flag 0 - gray)
   - Unknown Target (impact flag 4 - blue)
   - Currently Not Vulnerable (impact flag 3 - yellow)
   - Potentially Vulnerable (impact flag 2 - orange)
   - Vulnerable (impact flag 1 - red)

Keysight, EX-1004
Page 1295

4. Click **Save**.

Configuring Remediations

Platform: DC

Component: RNA

In addition to alerts, which are simple notifications of an event or a compliance policy violation, you can also configure responses called *remediations*. Remediations are programs that the Defense Center runs when a compliance 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 Check Point OPSEC SAM module, which, if you are running Check Point Firewall-1® NG FP3 or higher, allows you to dynamically block traffic sent to or sent from an IP address or network that violates a compliance policy.

  See Configuring Remediations for Check Point Firewalls on page 1298 for more information.
- 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 compliance policy.

  See Configuring Remediations for Cisco IOS Routers on page 1319 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 compliance policy.

  See Configuring Remediations for Cisco PIX Firewalls on page 1327 for more information.
- The Nmap Scanning module, which allows you to actively scan specific targets to determine operating systems and services running on those hosts.

  See Configuring Nmap Remediations on page 1333 for more information.
- The Set Attribute Value module, which allows you to set a host attribute on a host where a compliance event occurs.

  See Configuring Set Attribute Remediations on page 1340.

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.

Keysight, EX-1004
Page 1296

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 compliance 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 1343.

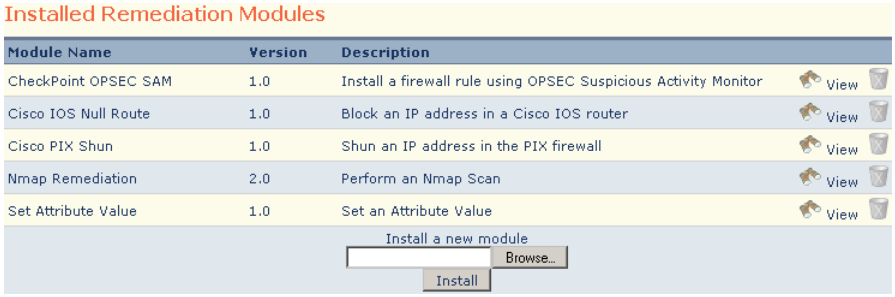
In addition to the default modules that Sourcefire provides, you can write custom remediation modules that perform other specific tasks when policy violations trigger. Refer to the *Sourcefire Remediation API Guide* for more information about writing your own remediation modules and installing them on 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: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Modules**.

   The Modules page appears.



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: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Modules**.

   The Modules page appears.

Keysight, EX-1004
Page 1297

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 Check Point Firewalls

Platform: DC

Component: RNA

Sourcefire provides a Check Point OPSEC SAM remediation module that allows you to install filters on Check Point firewalls when a compliance policy is violated. The Check Point OPSEC SAM remediation module supports Check Point Firewall-1 NG FP3 and higher.

IMPORTANT! A destination-based remediation will only work if you configure it to launch when a compliance rule that is based on a flow data event or intrusion event triggers. RNA network discovery events only transmit source hosts.

To create remediations for Check Point, use the following steps:

Access: P&R Admin/Admin

1. Create an OPSEC application on each Check Point firewall you want to connect to the Defense Center.

   See Creating the OPSEC Application on page 1299 for the procedures.
2. On the Defense Center, add an OPSEC instance for each Check Point firewall.

   See Adding an OPSEC Instance on page 1302 for the procedures.
3. Create specific remediations for each instance, based on the type of response you want to elicit on the firewall when compliance policies are violated.

   Each available remediation type is described in the following sections:
   - OPSEC Block To/From Destination IP/Network Remediations on page 1306 explains how to block all traffic sent to or received from the destination host or network in a compliance event.
   - OPSEC Block To Destination Service Remediations on page 1310 explains how to block service traffic to the destination host in a compliance event. The blocked service is determined by the protocol and port in the event.

Keysight, EX-1004
Page 1298

- OPSEC Block To/From Source IP/Network on page 1313 explains how to block all traffic sent to or received from the source host or network in a compliance event.
- OPSEC Block To Source Service on page 1316 explains how to block service traffic to the source host in a compliance event. The blocked service is determined by the protocol and port in the event.

4. Begin assigning OPSEC remediations to specific compliance policy rules.

TIP! You can clear configured responses from all configured Check Point firewalls at any time. To do this, select **Policy & Response > Remediations > Modules**. Click the **View** link that corresponds to the Check Point OPSEC SAM module. On the module details page that appears, click **Clear SAM Responses** to clear all responses. Note that this clears **all** responses for **all** configured Check Point firewall instances.

Creating the OPSEC Application

Platform: DC

Component: RNA

Before adding instances to the Check Point OPSEC SAM module and creating remediations, you should create an application for each Defense Center that will send remediations to the Check Point firewall.

IMPORTANT! The default port number for an OPSEC SAM server is 18183. For more information on modifying the default port for an OPSEC SAM server, see the Check Point documentation.

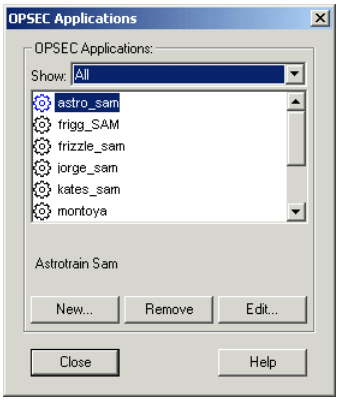
Keysight, EX-1004
Page 1299

To create and configure an OPSEC SAM application connection to the Defense Center:

Access: P&R Admin/Admin

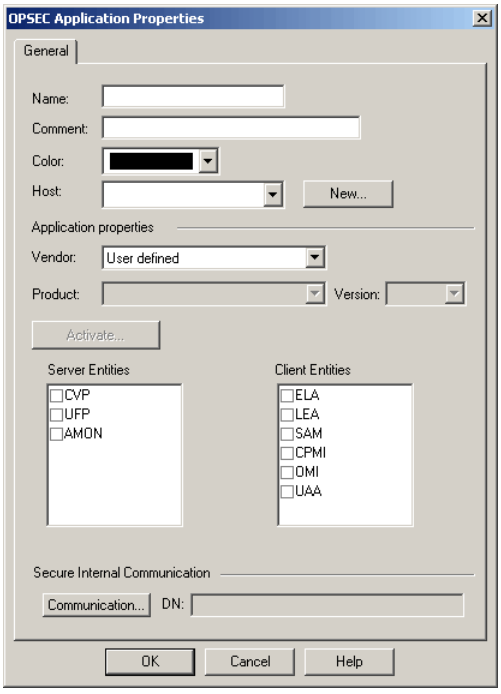
1. Open the Check Point SmartDashboard™ and, from the **Manage** menu, select **OPSEC Applications**.

   The OPSEC Applications dialog box appears.



2. Click **New** and, from the list that appears, select **OPSEC Application**.

   The OPSEC Application Properties page appears.



Keysight, EX-1004
Page 1300

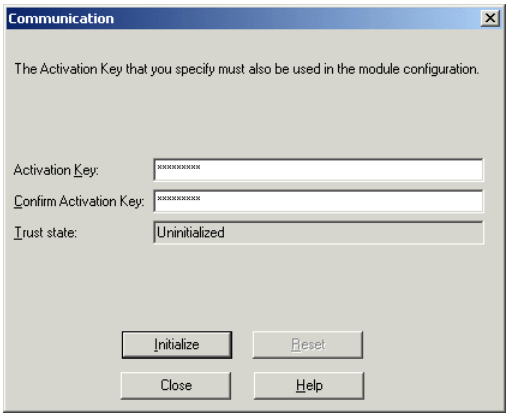
3. Fill in the fields as follows:
   - In the **Name** field, enter the name for the new OPSEC application that you will use to connect to the Defense Center. Spaces and special characters are not allowed in the application name. For example, you could use *pr_sam* as the application name.

     Note the name that you specify. You will need to enter it later when adding a remediation instance on the Defense Center.
   - Optionally, enter a comment in the **Comment** text box and select a color from the **Color** list.

     The Comment and Color features are only used by Check Point management software and have no effect on the connection with the Defense Center.
   - In the **Host** list, select the Defense Center that the firewall will communicate with.

     If the Defense Center does not appear in the list, click **New**, enter the Defense Center's full server name and domain (for example, `DCenter.example.com`) and IP address and click **OK**.
   - Use the default value, **User Defined**, for **Vendor**.
   - Make sure all check boxes are disabled under **Server Entities**.
   - In the **Client Entities** box, select the **SAM** check box to indicate that this application will implement Suspicious Activity Monitoring.
4. Click **Communication**.

   The Communication dialog box appears.



   The Communication dialog box allows you to specify an activation key to be used to ensure a secure transfer of authentication certificates between the Defense Center and the firewall.
5. In the **Activation Key** and **Confirm Activation Key** text boxes, enter a password.

   Note the activation key that you specify. You will need to enter it later when adding a remediation instance on the Defense Center.

Keysight, EX-1004
Page 1301

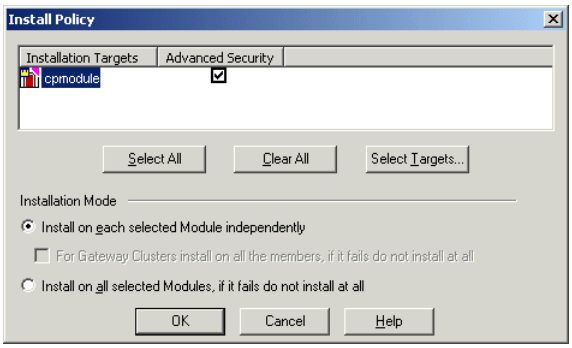
6. Click **Initialize** to create authentication certificates for the application and then click **Close**.

   The Communication dialog box closes and the OPSEC Application Properties page appears again.

7. Click **OK** to close the OPSEC Application Properties page and then click **Close** to close the OPSEC Applications dialog box.

8. From the **Policy** menu on the Check Point SmartDashboard, select **Install**.

   The Install Policy dialog box appears.



9. Select the module, enable **Install on each selected Module independently**, and click **OK**.

   The new policy is installed on the firewall.

10. Click **Close** to close the Installation Process dialog box.

    Your firewall is configured for OPSEC SAM communication with the Defense Center.

    If you have additional Check Point firewalls where you want the Defense Center to send remediations, repeat this procedure for each firewall. Otherwise, continue with Adding an OPSEC Instance to add an instance to the Defense Center.

Adding an OPSEC Instance

Platform: DC

Component: RNA

After you have configured an OPSEC application on the Check Point firewall, you can add an instance to the Defense Center. If you have multiple Check Point firewalls you want to send remediations to, you must create separate applications on the firewall (as described in Creating the OPSEC Application on page 1299) **and** add separate instances for each firewall to the Defense Center.

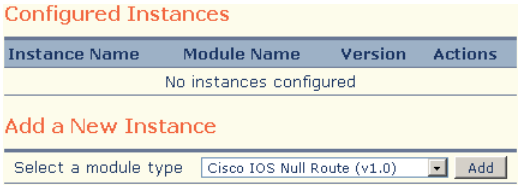
Keysight, EX-1004
Page 1302

To add a Check Point FW-1 instance:

Access: P&R Admin/Admin

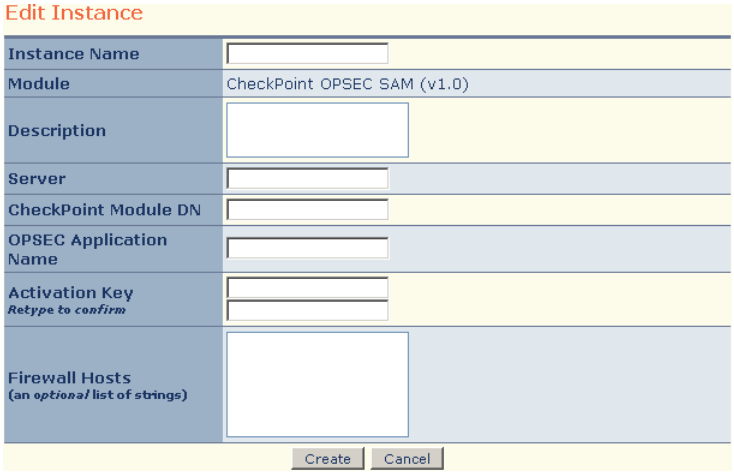
1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.



2. From the **Add a New Instance** list, select **Check Point OPSEC SAM** and click **Add**.

   The Instance Detail page appears.



3. In the **Instance Name** field, enter 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 to more than one Check Point firewall, you will have multiple instances, so you may want to choose a name such as `CPFW_01` and `CPFW_02`.

4. In the **Server** field, enter the server name (if DNS is enabled) or the IP address of the firewall.

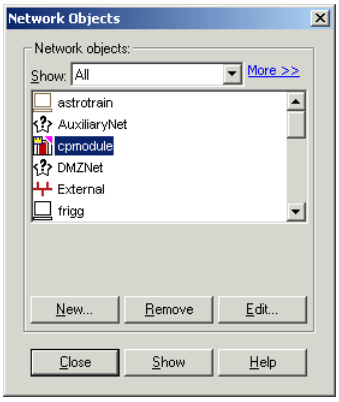
Keysight, EX-1004
Page 1303

5. Locate your Check Point module distinguished name.
   - If you do not know the Check Point module name, continue with step 6 for instructions on locating it.
   - If you already know the Check Point module name, skip to step 10 to enter it.

WARNING! The Check Point module distinguished name is **not** the same as the Check Point application distinguished name that was created in Creating the OPSEC Application on page 1299. The following steps provide detailed information about accessing the module distinguished name.

6. Open the Check Point SmartDashboard.
7. From the **Manage** menu, select **Network Objects**.

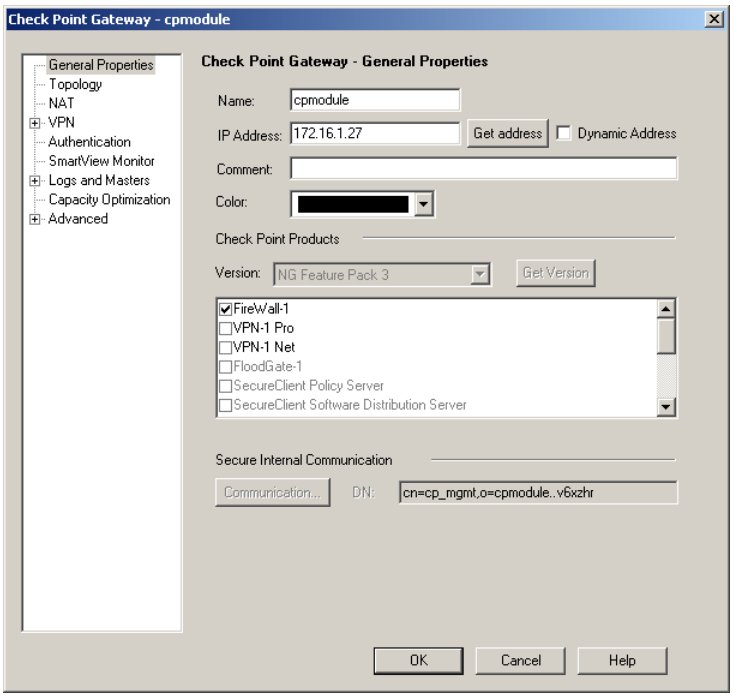
   The Network Objects dialog box appears.



Keysight, EX-1004
Page 1304

8. Select the Check Point module (typically **cpmodule**), and click **Edit**.

   The General Properties page for the module appears.



9. Under Secure Internal Communication, select and copy the module's distinguished name.

   In this example, the Check Point module DN name is `cn=cp_mgmt,o=cpmodule..v6xzhr`. This will be different for your Check Point server, but the format should be similar.

10. In the **Check Point Module DN** field, enter the Check Point firewall server module distinguished name (also known as the Secure Internal Communication, or SIC, name).

WARNING! The Check Point module DN name is **not** the same as the distinguished name for the application. If you are unsure of the module name, follow steps 6-9 of this procedure to locate it.

11. In the **OPSEC Application Name** field, enter the name of the application you created in Creating the OPSEC Application on page 1299.

12. In the **Activation Key** fields, enter the activation key you used when creating authentication credentials in Creating the OPSEC Application on page 1299.

Keysight, EX-1004
Page 1305

13. Optionally, in the **Firewall Hosts** field, enter any specific firewall object defined on the firewall server that you want to use as a target for the remediation.

    All, Localhost, and Gateways are provided by default for all Check Point firewall remediations. You should only populate this field if there are additional firewall objects you need to use as a remediation target.

    IMPORTANT! If you use the **All** option to set up OPSEC SAM responses to use more than one firewall host, and one of the firewall hosts fails, the failure message you receive does not indicate which host failed. As a workaround, you can review the firewall logs of the management firewall server to determine which specific responses failed and which succeeded.

14. 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 a compliance policy. See the following sections for more information:

    - OPSEC Block To/From Destination IP/Network Remediations on page 1306
    - OPSEC Block To Destination Service Remediations on page 1310
    - OPSEC Block To/From Source IP/Network on page 1313
    - OPSEC Block To Source Service on page 1316

OPSEC Block To/From Destination IP/Network Remediations

Platform: DC

Component: RNA

The Block To/From Destination IP/Network remediation configures your Check Point firewall to filter the destination IP or network that triggered a compliance rule.

IMPORTANT! Do **not** use this remediation as a response to a compliance rule that is based on an RNA event; RNA events only transmit a source host and not a destination host. You can use this remediation in response to compliance rules that are based on flow data events or intrusion events.

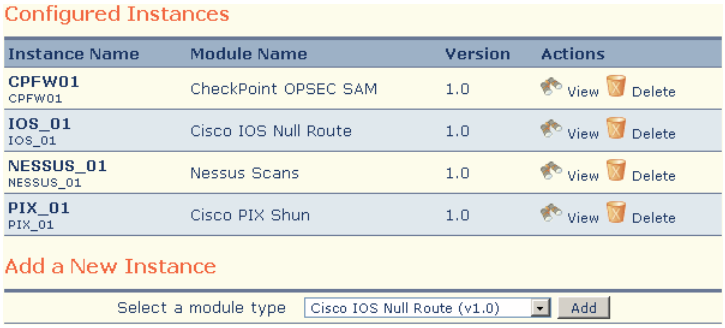
Keysight, EX-1004
Page 1306

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

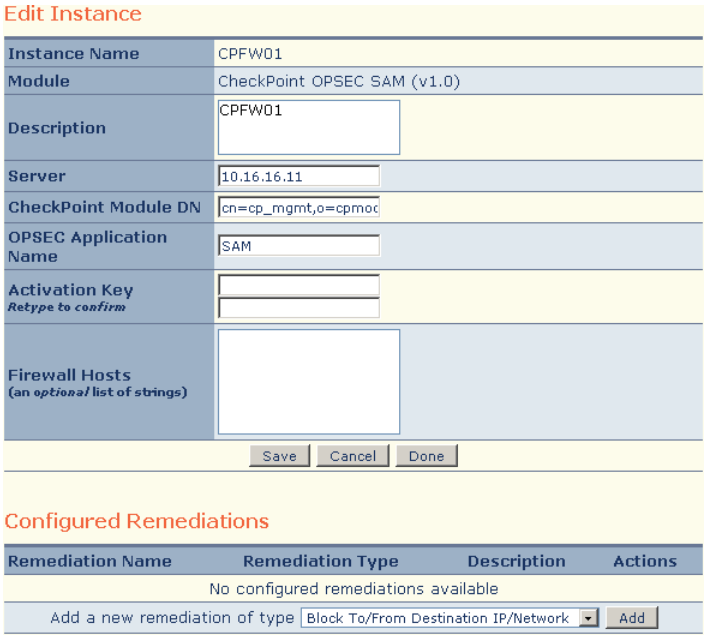
   The Instances page appears.



2. Next to the instance where you want to add the remediation, click **View**.

   If you have not yet added an instance, see Adding an OPSEC Instance on page 1302.

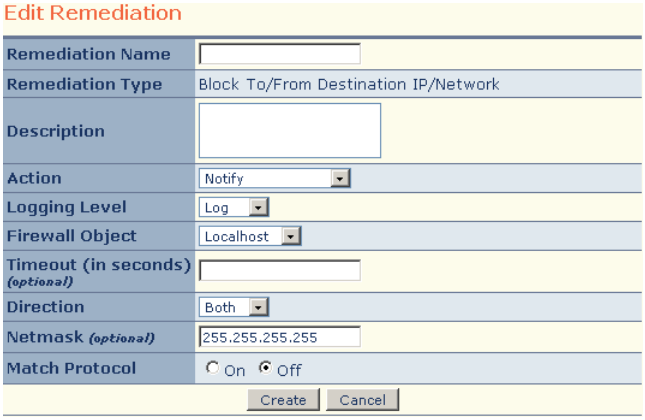
   The Instance Detail page appears.



Keysight, EX-1004
Page 1307

3. In the Configured Remediations section, select **Block To/From Destination IP/Network** and click **Add**.

   The Remediation Edit page appears.



4. In the **Remediation Name** field, enter a name for the remediation, using no spaces or special characters.
5. Optionally, in the **Description** field, enter a description for the remediation.
6. From the **Action** list, select the action you would like the firewall rule to perform. You can select from the following options:

Action	Description
Notify	All connection attempts are logged as specified in the Logging Level list. Traffic is not impeded.
Inhibit	All connection attempts are rejected. Existing connections will continue to function.
Inhibit and Close	All connection attempts are rejected. Existing connections are closed.
Drop	All connection attempts are dropped. Existing connections will continue to function.
Drop and Close	All connection attempts are dropped. Existing connections are closed.

Keysight, EX-1004
Page 1308

7. From the **Logging Level** list, select one of the following log types:

Log Type	Description
None	Does not log SAM responses to this rule.
Log	Performs detailed logging, but does not store the event that caused the SAM response.
Alert	Keeps more detailed logs along with the event that caused the SAM response.

8. From the **Firewall Object** list, select the firewall object that should receive the remediation.

   You can select from the predefined values of All, Gateways, or Localhost. You can also select from additional firewall objects that you added using the **Firewall Hosts** field when creating the instance. See Adding an OPSEC Instance on page 1302 for details about adding firewall objects to instances.

   IMPORTANT! If you use the **All** option to set up OPSEC SAM responses to use more than one firewall host, and one of the firewall hosts fails, the failure message you receive does not indicate which host failed. As a workaround, you can review the firewall logs of the management firewall server to determine which specific responses failed and which succeeded.

9. From the **Direction** list, select the traffic direction you want to filter from the firewall:
   - Select **To** to filter traffic sent to the host or network.
   - Select **From** to filter traffic sent from the host or network.
   - Select **Both** to filter traffic in both directions.

10. Optionally, in the **Timeout** field, enter the number of seconds that the specified action continues after it is initiated by the firewall.

    The maximum value is 2,147,483 seconds. Specifying a value of 0 (the default) causes the action to continue until it is manually disabled by a firewall administrator.

Keysight, EX-1004
Page 1309

11. Optionally, in the **Netmask** field, enter a subnet mask or use CIDR notation to describe the network you want to block. The default is 255.255.255.255 (matches a single IP address).

    For example, to block 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 30 addresses that include the triggering IP address, 255.255.255.224 or 27 as the netmask. In this case, if the IP address 10.1.1.15 triggered 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 255.255.255.255, or enter 32.

12. Specify whether the **Match Protocol** should be enabled or disabled:
    - Select **On** to include the protocol in the firewall configuration. For example, if a UDP packet triggers the remediation, only UDP traffic from the triggering host is filtered.
    - Select **Off** (the default) to filter all traffic from the triggering host, regardless of protocol.

13. Click **Create**, then click **Done**.

    The remediation is added.

OPSEC Block To Destination Service Remediations

Platform: DC

Component: RNA

The Block To Destination Service remediation configures your Check Point firewall to filter traffic sent to a service (protocol and port) running on the destination host. This is the destination host described in the compliance event.

IMPORTANT! Do **not** use this remediation as a response to a compliance rule that is based on an RNA event; RNA events only transmit a source host and not a destination host. You can use this remediation in response to compliance rules that are based on flow data events or intrusion events.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

2. Next to the instance you want to view, click **View**.

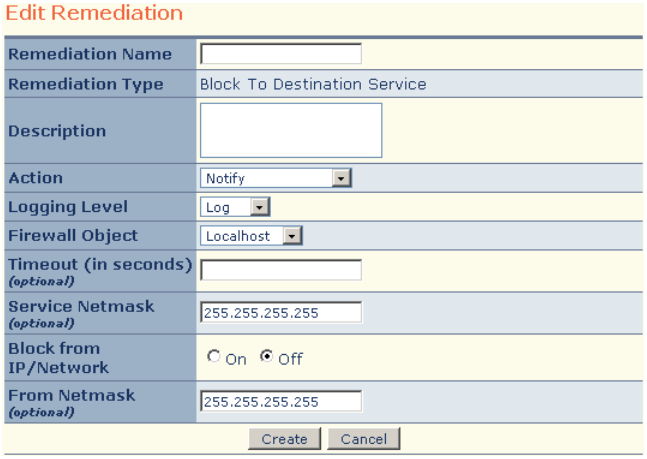
   If you have not yet added an instance, see Adding an OPSEC Instance on page 1302.

   The Instance Detail page appears.

Keysight, EX-1004
Page 1310

3. In the Configured Remediations section, select **Block To Destination Service** and click **Add**.

   The Remediation Edit page appears.



4. In the **Remediation Name** field, enter a name for the remediation, using no spaces or special characters.
5. Optionally, in the **Description** field, enter a description for the remediation.
6. From the **Action** list, select the action you would like the firewall rule to perform. You can select from the following options:

Action	Description
Notify	All connection attempts are logged as specified in the Logging Level list. Traffic is not impeded.
Inhibit	All connection attempts are rejected. Existing connections will continue to function.
Inhibit and Close	All connection attempts are rejected. Existing connections are closed.
Drop	All connection attempts are dropped. Existing connections will continue to function.
Drop and Close	All connection attempts are dropped. Existing connections are closed.

Keysight, EX-1004
Page 1311

7. From the **Logging Level** list, select one of the following log types:

Log Type	Description
None	Does not log SAM responses to this rule.
Log	Performs detailed logging, but does not store the event that caused the SAM response.
Alert	Keeps more detailed logs along with the event that caused the SAM response.

8. From the **Firewall Object** list, select the firewall object that should receive the remediation.

   You can select from the predefined values All, Gateways, or Localhost. You can also select from additional firewall objects that you added using the **Firewall Hosts** field when creating the instance. See Adding an OPSEC Instance on page 1302 for details about adding firewall objects to instances.

   IMPORTANT! If you use the **All** option to set up OPSEC SAM responses to use more than one firewall host, and one of the firewall hosts fails, the failure message you receive does not indicate which host failed. As a workaround, you can review the firewall logs of the management firewall server to determine which specific responses failed and which succeeded.

9. Optionally, in the **Timeout** field, enter the number of seconds that the specified action continues after it is initiated by the firewall.

   The maximum value is 2,147,483 seconds. Specifying a value of 0 (the default) causes the action to continue until it is manually disabled by a firewall administrator.

10. Optionally, in the **Service Netmask** field, enter a subnet mask or use CIDR notation to describe the network where you want to filter service access. The default is 255.255.255.255 (filters connections to the service on the IP address that triggered the compliance policy violation).

    For example, to filter access to the service on any host in the Class C network of the IP address that triggered the rule, use 255.255.255.0 or 24 as the netmask. To filter access to the service on a block of 30 addresses around the IP address that triggered the policy violation, use 255.255.255.224 or 27.

Keysight, EX-1004
Page 1312

11. Specify whether the **Block from IP/Network** option is enabled or disabled:
    - Select **On** to filter from the source IP address or network in the event to the service running on the destination IP address.
    - Select **Off** (the default) to indicate that traffic from the source IP address or network to the destination IP address or network is not filtered.

    WARNING! If you want to use this remediation in response to a rule based on an RNA event (which do not have source IP addresses), disable **Block from IP/Network**.

12. Optionally, in the **From Netmask** field, enter a subnet mask or use CIDR notation to describe the source network that you want to filter access for. The default is 255.255.255.255 (matches a single IP address).

    For example, to disallow the entire Class C network that the source IP address belongs to from contacting a service on the destination host, use 255.255.255.0 or 24 as the netmask.

    As another example, to block 30 addresses that include the source IP address from accessing the destination host's service, specify 255.255.255.224 or 27 as the netmask. In this case, if the IP address 10.1.1.15 was the source IP address in the event that triggered the remediation, all IP addresses between 10.1.1.1 and 10.1.1.30 are unable to access the service on the destination host. To block only the source IP address, leave the field blank, enter 255.255.255.255, or enter 32.

13. Click **Create**, then click **Done**.

    The remediation is added.

OPSEC Block To/From Source IP/Network

Platform: DC

Component: RNA

The Block To/From Source IP/Network remediation configures your Check Point firewall to filter the source IP or network that triggered the event. If the compliance policy violation was caused by a flow data event or an intrusion event, this is the source IP address in the event. For RNA events, this remediation triggers on the IP address associated with the event.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

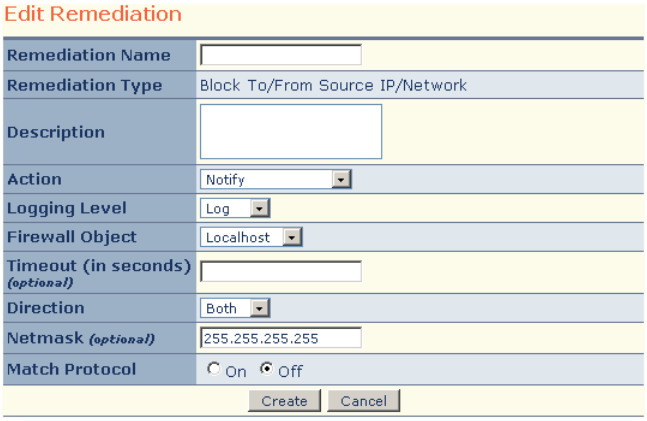
Keysight, EX-1004
Page 1313

2. Next to the instance where you want to add the remediation, click **View**. If you have not yet added an instance, see Adding an OPSEC Instance on page 1302.

   The Instance Detail page appears.

3. In the Configured Remediations section, select **Block To/From Source IP/Network** and click **Add**.

   The Remediation Edit page appears.



4. In the **Remediation Name** field, enter a name for the remediation, using no spaces or special characters.
5. Optionally, in the **Description** field, enter a description for the remediation.
6. From the **Action** list, select the action you would like the firewall rule to perform. You can select from the following options:

Action	Description
Notify	All connection attempts are logged as specified in the Logging Level list. Traffic is not impeded.
Inhibit	All connection attempts are rejected. Existing connections will continue to function.
Inhibit and Close	All connection attempts are rejected. Existing connections are closed.
Drop	All connection attempts are dropped. Existing connections will continue to function.
Drop and Close	All connection attempts are dropped. Existing connections are closed.

Keysight, EX-1004
Page 1314

7. From the **Logging Level** list, select one of the following log types:

Log Type	Description
None	Does not log SAM responses to this rule.
Log	Performs detailed logging, but does not store the event that caused the SAM response.
Alert	Keeps more detailed logs along with the event that caused the SAM response.

8. From the **Firewall Object** list, select the firewall object that should receive the remediation.

   You can select from the predefined values of All, Gateways, or Localhost. You can also select from additional firewall objects that you added using the **Firewall Hosts** field when creating the instance. See Adding an OPSEC Instance on page 1302 for details about adding firewall objects to instances.

   IMPORTANT! If you use the **All** option to set up OPSEC SAM responses to use more than one firewall host, and one of the firewall hosts fails, the failure message you receive does not indicate which host failed. As a workaround, you can review the firewall logs of the management firewall server to determine which specific responses failed and which succeeded.

9. Optionally, in the **Timeout** field, enter the number of seconds that the specified action continues after it is initiated by the firewall.

   The maximum value is 2,147,483 seconds. Specifying a value of 0 (the default) causes the action to continue until it is manually disabled by a firewall administrator.

10. From the **Direction** list, select the traffic direction you want to filter from the firewall:
    - Select **To** to filter traffic sent to the source host or network.
    - Select **From** to filter traffic sent from the source host or network.
    - Select **Both** to filter traffic in both directions.

Keysight, EX-1004
Page 1315

11. Optionally, in the **Netmask** field, enter a subnet mask or use CIDR notation to describe the network you want to block. The default is 255.255.255.255 (matches a single IP address).

    For example, to block an entire Class C network when a single host triggers a rule (this is not recommended), use 255.255.255.0 or 24 as the netmask.

    As another example, to block 30 addresses that include the triggering source IP address, specify 27 or 255.255.255.224 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 255.255.255.255, or enter 32.

12. Specify whether the **Match Protocol** should be enabled or disabled:
    - Select **On** to include the protocol specified in the triggering event in the firewall configuration. For example, if a UDP packet triggers the remediation, only UDP traffic from the triggering host is filtered.
    - Select **Off** (the default) to filter all traffic from the triggering host, regardless of protocol.

13. Click **Create**, then click **Done**.

    The remediation is added.

OPSEC Block To Source Service

Platform: DC

Component: RNA

The Block To Source Service remediation configures your Check Point firewall to filter traffic sent to a service (protocol and port) running on the source host described in the event that triggered the compliance policy violation. If the compliance policy violation was caused by a flow data event or an intrusion event, this is the source IP address in the event. For RNA events, this remediation triggers on the IP address associated with the event.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

2. Next to the instance where you want to add the remediation, click **View**.

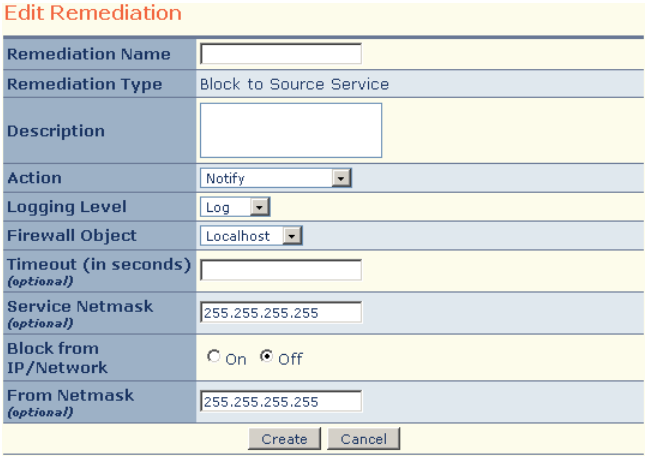
   If you have not yet added an instance, see Adding an OPSEC Instance on page 1302.

   The Instance Detail page appears.

Keysight, EX-1004
Page 1316

3. In the Configured Remediations section, select **Block To Source Service** and click **Add**.

   The Remediation Edit page appears.



4. In the **Remediation Name** field, enter a name for the remediation, using no spaces or special characters.
5. Optionally, in the **Description** field, enter a description for the remediation.
6. From the **Action** list, select the action you would like the firewall rule to perform. You can select from the following options:

Action	Description
Notify	All connection attempts are logged as specified in the Logging Level list. Traffic is not impeded.
Inhibit	All connection attempts are rejected. Existing connections will continue to function.
Inhibit and Close	All connection attempts are rejected. Existing connections are closed.
Drop	All connection attempts are dropped. Existing connections will continue to function.
Drop and Close	All connection attempts are dropped. Existing connections are closed.

Keysight, EX-1004
Page 1317

7. From the **Logging Level** list, select one of the following log types:

Log Type	Description
None	Does not log SAM responses to this rule.
Log	Performs detailed logging, but does not store the event that caused the SAM response.
Alert	Keeps more detailed logs along with the event that caused the SAM response.

8. From the **Firewall Object** list, select the firewall object that should receive the remediation.

   You can select from the predefined values of All, Gateways, or Localhost. You can also select from additional firewall objects that you added using the **Firewall Hosts** field when creating the instance. See Adding an OPSEC Instance on page 1302 for details about adding firewall objects to instances.

   IMPORTANT! If you use the **All** option to set up OPSEC SAM responses to use more than one firewall host, and one of the firewall hosts fails, the failure message you receive does not indicate which host failed. As a workaround, you can review the firewall logs of the management firewall server to determine which specific responses failed and which succeeded.

9. Optionally, in the **Timeout** field, enter the number of seconds that the specified action continues after it is initiated by the firewall.

   The maximum value is 2,147,483 seconds. Specifying a value of 0 (the default) causes the action to continue until it is manually disabled by a firewall administrator.

10. Optionally, in the **Service Netmask** field, enter a subnet mask or use CIDR notation to describe the network where you want to filter service access. The default is 255.255.255.255 (filters connections to the service on the IP address that triggered the compliance policy violation).

    For example, to filter access to the service on any host in the Class C network of the IP address that triggered the rule, use 255.255.255.0 or 24 as the netmask. To filter access to the service on a block of 30 addresses around the IP address that triggered the policy violation, use 255.255.255.224 or 27.

Keysight, EX-1004
Page 1318

11. Specify whether the **Block from IP/Network** option is enabled or disabled:
    - Select **On** to filter from the destination IP address or network in the event to the service running on the source IP address.
    - Select **Off** (the default) to indicate that traffic from the destination IP address or network to the source IP address or network is not filtered.

    WARNING! If you want to use this remediation in response to a rule based on an RNA event (which does not have destination IP addresses), disable **Block from IP/Network**.

12. Optionally, in the **From Netmask** field, enter a subnet mask or use CIDR notation to describe the destination network that you want to filter access for. The default is 255.255.255.255 (matches a single IP address).

    For example, to disallow the entire Class C network that the destination IP address belongs to from contacting a service on the source host, use 255.255.255.0 or 24 as the netmask.

    As another example, to block 30 addresses that include the destination IP address from accessing the source host's service, specify 255.255.255.224 or 27 as the netmask. In this case, if the IP address 10.1.1.15 was the destination IP address in the event that triggered the remediation, all IP addresses between 10.1.1.1 and 10.1.1.30 are unable to access the service on the source host. To block only the destination IP address, leave the field blank, enter 255.255.255.255, or enter 32.

13. Click **Create**, then click **Done**.

    The remediation is added.

Configuring Remediations for Cisco IOS Routers

Platform: DC

Component: RNA

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 compliance 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 compliance 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).

Keysight, EX-1004
Page 1319

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 compliance rule that is based on a flow data event or intrusion event triggers. RNA network 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: P&R Admin/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 1320 for the procedures.

3. Create specific remediations for each instance, based on the type of response you want to elicit on the router when compliance policies are violated.

   Each available remediation type is described in the following sections:

   - Cisco IOS Block Destination Remediations on page 1322
   - Cisco IOS Block Destination Network Remediations on page 1324
   - Cisco IOS Block Source Remediations on page 1325
   - Cisco IOS Block Source Network Remediations on page 1326

4. Begin assigning Cisco IOS remediations to specific compliance policy rules.

Adding a Cisco IOS Instance

Platform: DC

Component: RNA

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.

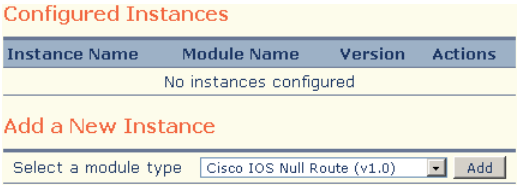
Keysight, EX-1004
Page 1320

To add a Cisco IOS instance:

Access: P&R Admin/Admin

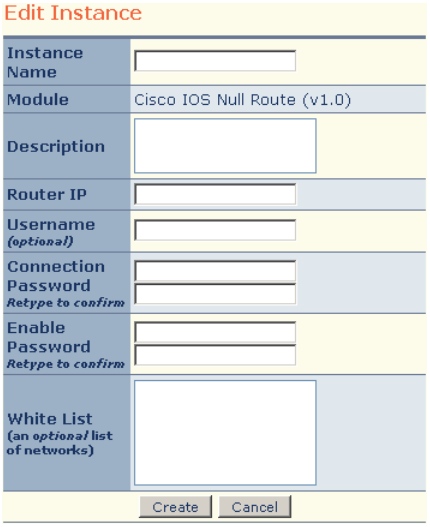
1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.



2. From the **Add a New Instance** list, select **Cisco IOS Null Route (v1.0)** and click **Add**.

   The Instance Detail page appears.



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.

Keysight, EX-1004
Page 1321

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 59.
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 compliance policies. See the following sections for more information:

   - Cisco IOS Block Destination Remediations on page 1322
   - Cisco IOS Block Destination Network Remediations on page 1324
   - Cisco IOS Block Source Remediations on page 1325
   - Cisco IOS Block Source Network Remediations on page 1326

Cisco IOS Block Destination Remediations

Platform: DC

Component: RNA

The Cisco IOS Block Destination remediation allows you to block traffic sent from the router to the destination host in a compliance event.

IMPORTANT! Do **not** use this remediation as a response to a compliance rule that is based on an RNA event; RNA events only transmit a source host and not a destination host. You can use this remediation in response to compliance rules that are based on flow data events or intrusion events.

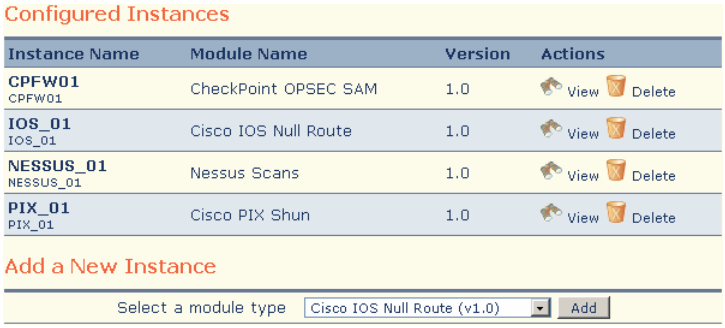
Keysight, EX-1004
Page 1322

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

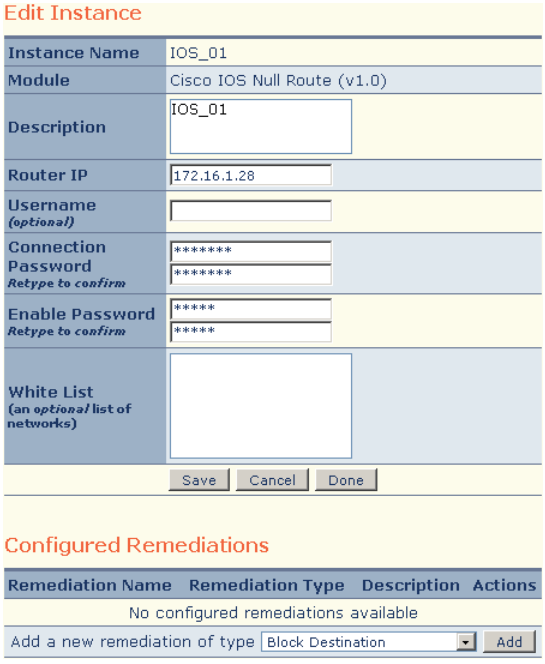
   The Instances page appears.



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 1320.

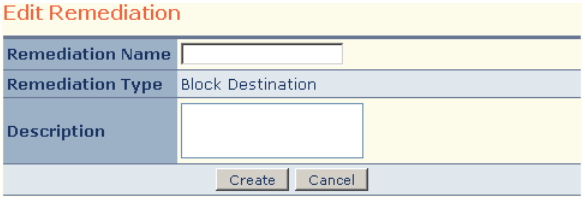
   The Instance Detail page appears.



Keysight, EX-1004
Page 1323

3. In the Configured Remediations section, select **Block Destination** and click **Add**.

   The Remediation Edit 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 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

Platform: DC

Component: RNA

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 compliance event.

IMPORTANT! Do **not** use this remediation as a response to a compliance rule that is based on an RNA event; RNA events only transmit a source host and not a destination host. You can use this remediation in response to compliance rules that are based on flow data events or intrusion events.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

2. Next to the instance where you want to add the remediation, click **View**.

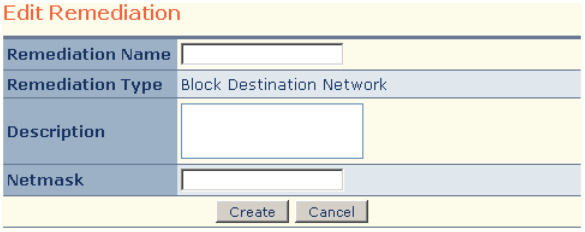
   If you have not yet added an instance, see Adding a Cisco IOS Instance on page 1320.

   The Instance Detail page appears.

Keysight, EX-1004
Page 1324

3. In the Configured Remediations section, select **Block Destination Network** and click **Add**.

   The Remediation Edit 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 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

Platform: DC

Component: RNA

The Cisco IOS Block Source remediation allows you to block any traffic sent from the router to the source host included in a compliance event that violates a compliance policy. The source host is the source IP address in the flow data event or intrusion event upon which the compliance rule is based, or the host IP address in an RNA event.

Keysight, EX-1004
Page 1325

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

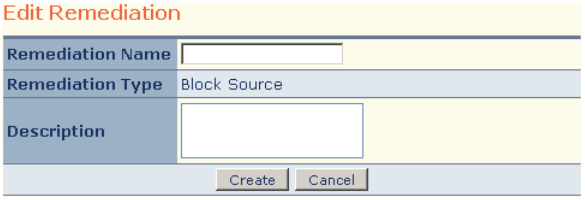
2. Next to the instance where you want to add the remediation, click **View**.

   If you have not yet added an instance, see Adding a Cisco IOS Instance on page 1320.

   The Instance Detail page appears.

3. In the Configured Remediations section, select **Block Source** and click **Add**.

   The Remediation Edit 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 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

Platform: DC

Component: RNA

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 compliance event. The source host is the source IP address in the flow data event or intrusion event upon which the compliance rule is based, or the host IP address in an RNA event.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

Keysight, EX-1004
Page 1326

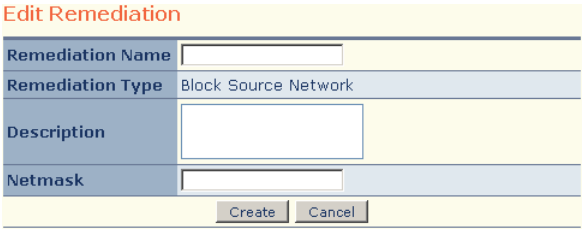
2. Next to the instance where you want to add the remediation, click **View**.

   If you have not yet added an instance, see Adding a Cisco IOS Instance on page 1320.

   The Instance Detail page appears.

3. In the Configured Remediations section, select **Block Source Network** and click **Add**.

   The Remediation Edit page appears.



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

Platform: DC

Component: RNA

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 compliance policy

Keysight, EX-1004
Page 1327

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 compliance rule that is based on a flow data event or intrusion event triggers. RNA network 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: P&R Admin/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 1329 for the procedures.

3. Create specific remediations for each instance, based on the type of response you want to elicit on the firewall when compliance policies are violated.

   The available remediation types are described in the following sections:

   - Cisco PIX Block Destination Remediations on page 1331
   - Cisco PIX Block Source Remediations on page 1333

4. Begin assigning Cisco PIX remediations to specific compliance policy rules.

Keysight, EX-1004
Page 1328

Adding a Cisco PIX Instance

Platform: DC

Component: RNA

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.

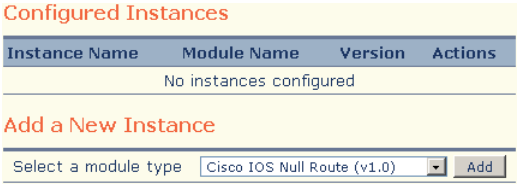
IMPORTANT! Sourcefire recommends that you use an SSH connection instead of a Telnet connection. Data transmitted using SSH is encrypted, making it much more secure than Telnet.

To add a Cisco PIX instance:

Access: P&R Admin/Admin

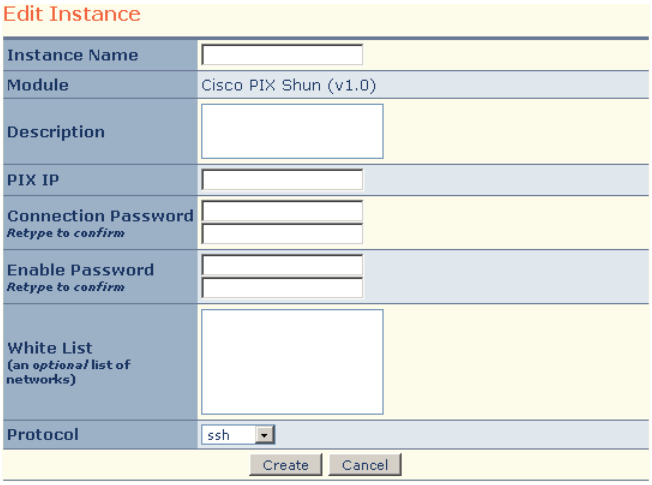
1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.



2. From the **Add a New Instance** list, select **Cisco PIX Shun** and click **Add**.

   The Instance Detail page appears.



Keysight, EX-1004
Page 1329

3. In the **Instance Name** field, enter 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. In the **PIX IP** field, enter the IP address of the Cisco PIX firewall you want to use for the remediation.

5. 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.

6. 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.

7. 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 59.

8. From the **Protocol** list, select the method you want to use to connect to the firewall.

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 in compliance policies. See the following sections for more information:

   - Cisco PIX Block Destination Remediations on page 1331
   - Cisco PIX Block Source Remediations on page 1333

Keysight, EX-1004
Page 1330

Cisco PIX Block Destination Remediations

Platform: DC

Component: RNA

The Cisco PIX Block Destination remediation allows you to block traffic sent from the destination host in a compliance event.

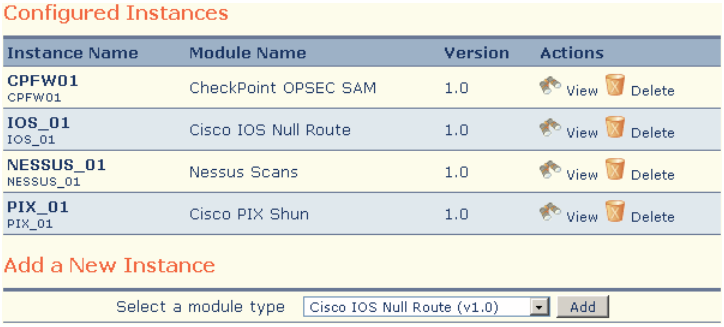
IMPORTANT! Do **not** use this remediation as a response to a compliance rule that is based on an RNA event; RNA events only transmit a source host and not a destination host. You can use this remediation in response to compliance rules that are based on flow data events or intrusion events.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

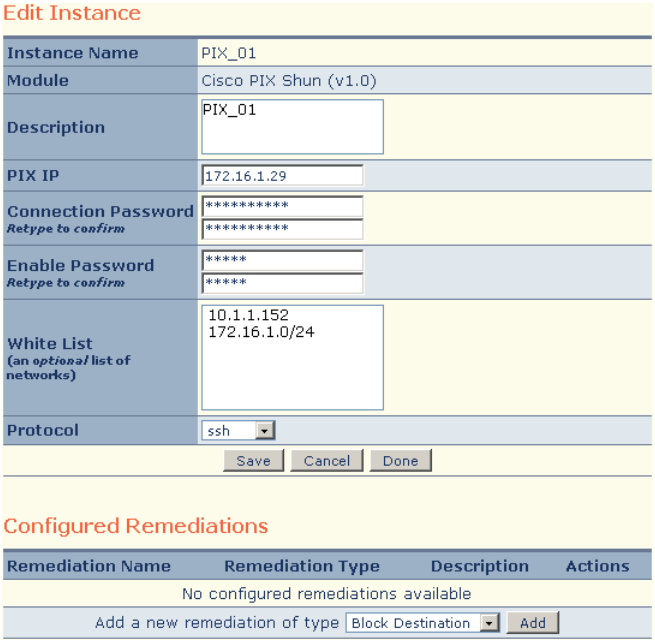


Keysight, EX-1004
Page 1331

2. Next to the instance where you want to add the remediation, click **View**.

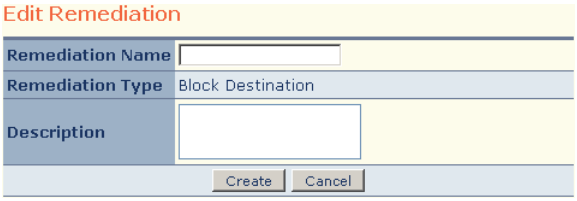
   If you have not yet added an instance, see Adding a Cisco PIX Instance on page 1329.

   The Instance Detail page appears.



3. In the Configured Remediations section, select **Block Destination** and click **Add**.

   The Remediation Edit 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_BlockDest`.

5. Optionally, in the **Description** field, enter a description of the remediation.

6. Click **Create**, then click **Done**.

   The remediation is added.

Keysight, EX-1004
Page 1332

Cisco PIX Block Source Remediations

Platform: DC

Component: RNA

The Cisco PIX Block Source remediation allows you to block any traffic sent from the source host included in the event that violates a compliance policy. The source host is the source IP address in the flow data event or intrusion event upon which the compliance rule is based, or the host IP address in an RNA event.

To add the remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

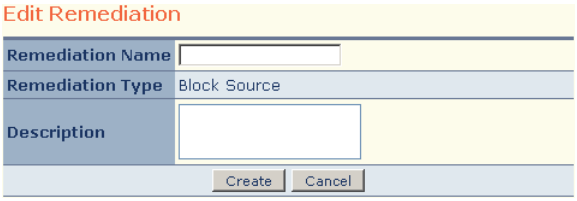
2. Next to the instance where you want to add the remediation, click **View**.

   If you have not yet added an instance, see Adding a Cisco PIX Instance on page 1329.

   The Instance Detail page appears.

3. In the Configured Remediations section, select **Block Source** and click **Add**.

   The Remediation Edit 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

Platform: DC

Component: RNA

You can respond to a compliance event by scanning the host where the triggering event occurred. You can choose to scan only the port from the event that triggered the compliance event.

Keysight, EX-1004
Page 1333

To set up Nmap scanning in response to a compliance 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 1334
- Nmap Scan Remediations on page 1335

Adding an Nmap Scan Instance

Platform: DC

Component: RNA

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 service information. You can set up scan instances for the local Nmap module on your Defense Center and for any sensors 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 sensor. 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:

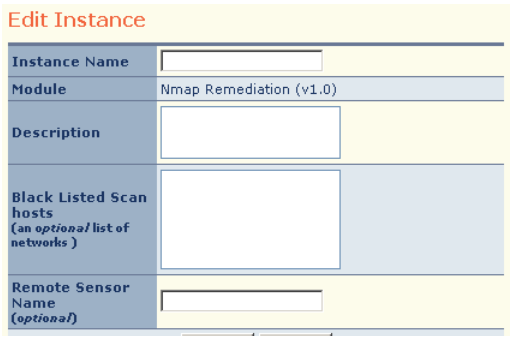
Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

   The Instances page appears.

2. Select **Nmap Remediation (v1.0)** from the **Add a module type** drop-down list and click **Add**.

   The Instance Detail 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 (-).

Keysight, EX-1004
Page 1334

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:
   - an exact IP address (for example, `192.168.1.101`)
   - 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 59.
6. Optionally, to run the scan from a remote sensor instead of the Defense Center, specify the name or IP address of the sensor in the **Remote Sensor Name** field.
7. Click **Create**.

   The scan instance is created.

Nmap Scan Remediations

Platform: DC

Component: RNA

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 compliance 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 RNA, NetFlow, and the host input feature can add hosts to the network map.

For more information on the specific settings in an Nmap remediation, see Understanding Nmap Remediations on page 1419.

Note that Nmap-supplied service and operating system data remains static until you run another Nmap scan. If you plan to scan a host for operating system and service data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and service data up-to-date. For more information, see Automating Nmap Scans on page 1722. 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.

Keysight, EX-1004
Page 1335

To create a Nmap remediation:

Access: P&R Admin/Admin

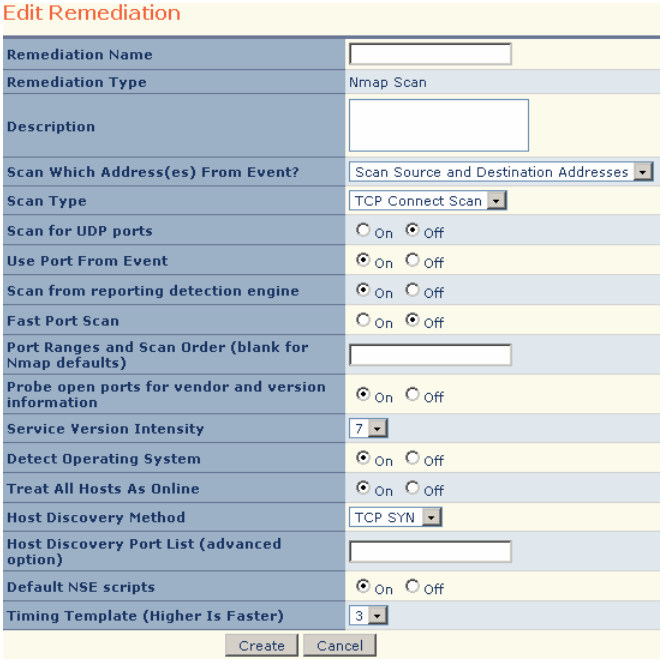
1. Select **Operations > Tools > Scanners**.

   The Scanners page appears.



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.

Keysight, EX-1004
Page 1336

5. If you plan to use this remediation in response to a compliance rule that triggers on an intrusion event, a flow event, or a RUA 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 compliance rule that triggers on an RNA 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 compliance 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 compliance policy violations, configure the **Use Port From Event** option:

Keysight, EX-1004
Page 1337

- Select **On** to scan the port in the compliance event, rather than the ports you specify in step 12.

  If you scan the port in the compliance 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 compliance 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 sensor 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.

    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 compliance policy violation, as described in step 8.
12. To probe open ports for service vendor and version information, configure **Probe open ports for vendor and version information:**
    - Select **On** to scan open ports on the host for service information to identify service vendors and versions.
    - Select **Off** to continue using RNA service 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.

Keysight, EX-1004
Page 1338

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 RNA 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 service, 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.

Keysight, EX-1004
Page 1339

20. Click **Save**, then click **Done**.

    The remediation is created.

Configuring Set Attribute Remediations

Platform: DC

Component: RNA

You can respond to a compliance 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 985 and Working with User-Defined Host Attributes on page 987.

To configure setting an attribute value in response to a compliance 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 1340
- Set Attribute Value Remediations on page 1341

Adding a Set Attribute Value Instance

Platform: DC

Component: RNA

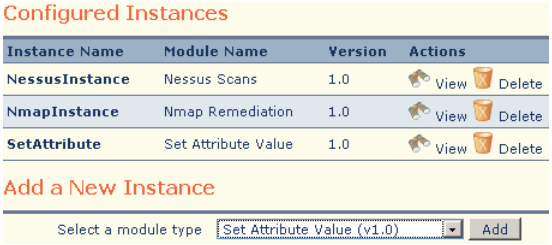
You can set up an instance to set attribute values in response to compliance rule violations.

To create a set attribute instance:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

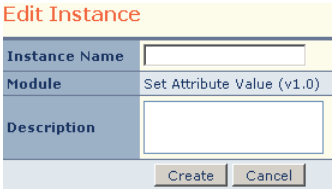
   The Instances page appears.



Keysight, EX-1004
Page 1340

2. Select **Set Attribute Value (v1.0)** from the **Add a module type** drop-down list and click **Add**.

   The Instance Detail 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

Platform: DC

Component: RNA

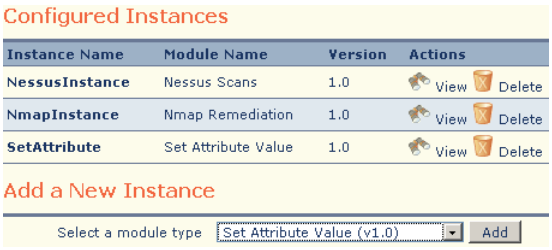
You can create a set attribute value remediation for each attribute value you want to be able to set in response to a compliance rule violation. If the attribute you want to set is a text attribute, you can set the remediation to use the description from the event as the attribute value.

To create a set attribute value remediation:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Remediations > Instances**.

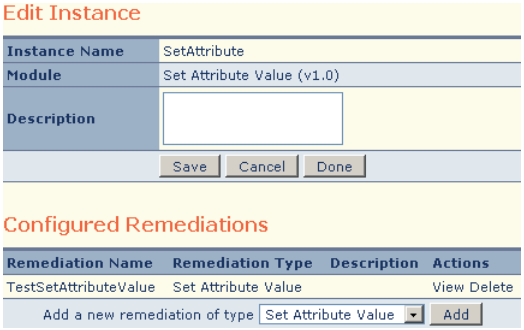
   The Instances page appears.



Keysight, EX-1004
Page 1341

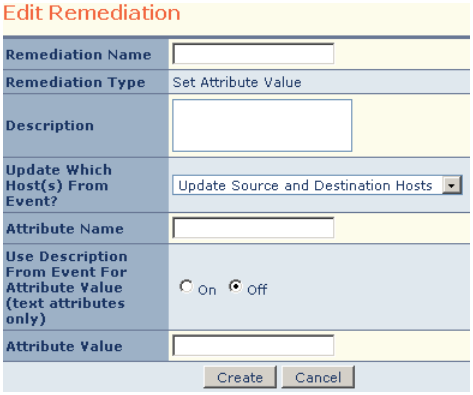
2. Click **View** next to the scan instance where you want to add a remediation.

   The Edit Instance page appears.



3. Select **Set Attribute Value** from the **Add a new remediation of type** drop-down list.

   The Remediation Edit page appears.



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.

Keysight, EX-1004
Page 1342

6. If you plan to use this remediation in response to a compliance rule that triggers on an intrusion event, RUA event, or a flow data 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 compliance rule that triggers on an RNA 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

Platform: DC

Component: RNA

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 1597
- Searching for Remediation Status Events on page 1349

Keysight, EX-1004
Page 1343

Viewing Remediation Status Events

Platform: DC

Component: RNA

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 1618.

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 1347.
modify the time and date range for displayed events	see Setting Event Time Constraints on page 1597. 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 1606 and Sorting Drill-down Workflow Pages on page 1612.
temporarily use different workflow	click **Workflows**. For more information, see Selecting Workflows on page 1590.
navigate to the compliance events view to see associated events	click **Compliance Events**. For more information, see Navigating between Workflows on page 1614.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.

Keysight, EX-1004
Page 1344

Remediation Status Actions (Continued)

To...	You can...
generate a report based on the data in the table view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
drill down to the next page in the workflow, constraining on a specific value	use one of the following methods: • 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
delete remediation status events from the system	use one of the following methods: • 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	click **Search**. For more information, see Searching for Remediation Status Events on page 1349.

Keysight, EX-1004
Page 1345

To view remediation status events:

Access: P&R Admin/Admin

- Select **Policy & Response > Responses > Remediations > Status**.

  The first page of the default remediations workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

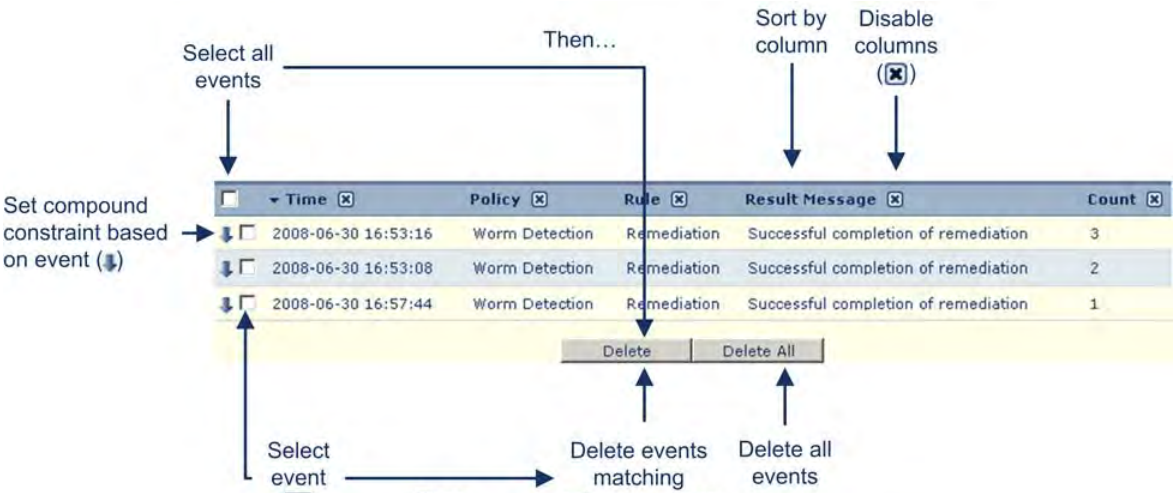
TIP! If you are using a custom workflow that does not include the table view of remediations, from the **Workflows** menu on the toolbar, select **Remediation Status**.

Working with Remediation Status Events

Platform: DC

Component: RNA

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.

Keysight, EX-1004
Page 1346

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 1606.
- Using Compound Constraints on page 1609.
- Sorting Drill-down Workflow Pages on page 1612.
- Understanding the Remediation Status Table on page 1347

Understanding the Remediation Status Table

Platform: DC

Component: RNA

You can configure the Defense Center to launch a variety of responses to policy violations and to RNA network 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 Responses for Compliance Policies on page 1281.

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 compliance policy that was violated and triggered the remediation.
Remediation Name	The name of the remediation that was launched.

Keysight, EX-1004
Page 1347

Remediation Status Fields (Continued)

Field	Description
Result Message	A message that describes what happened when the remediation was launched. Status messages include: • 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 service • Unable to gain required privileges on remote device or service • Timeout logging into remote device or service • Timeout executing remote commands or services • The remote device or service was unreachable • The remediation was attempted but failed • Failed to execute remediation program • Unknown/unexpected error **IMPORTANT!** If custom remediation modules are installed, you may see additional status messages that are implemented by the custom module.
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: P&R Admin/Admin

▶ Select **Policy & Response > Responses > Remediations > Status**.

The table view appears. For information on working with remediation status events, see Working with Remediation Status Events on page 1343.

TIP! If you are using a custom workflow that does not include the table view of remediation status events, click **Workflows**. On the Select Workflow page, click **Remediation Status**.

Keysight, EX-1004
Page 1348

Searching for Remediation Status Events

Platform: DC

Component: RNA

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 re-use 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	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: • 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 service • Unable to gain required privileges on remote device or service • Timeout logging into remote device or service • Timeout executing remote commands or services • The remote device or service was unreachable • The remediation was attempted but failed • Failed to execute remediation program • Unknown/unexpected error **IMPORTANT!** If you installed custom remediation modules, you may be able to enter additional status messages implemented by the custom module.
Time	Specify the date and time the Defense Center launched the remediation. See Specifying Time Constraints in Searches on page 1556 for the syntax for entering time.
Remediation Name	Enter the exact name of the remediation that was launched. This is the name you specified when you created the remediation.
Policy	Enter the name of the compliance policy that triggered the remediation.
Rule	Enter the name of the compliance rule that triggered the remediation.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

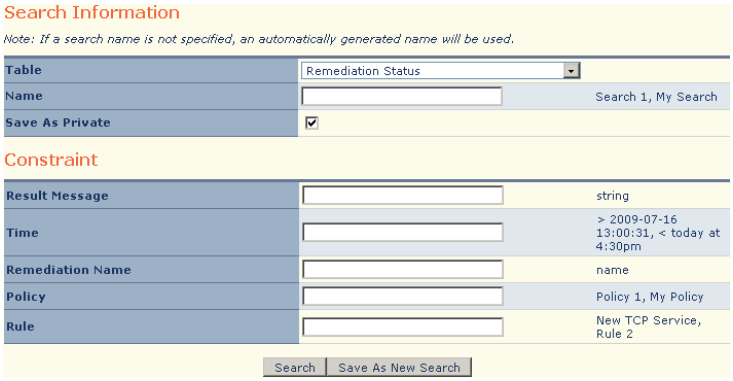
Keysight, EX-1004
Page 1349

To search for remediation status events:

Access: P&R Admin/Admin

1. Select **Analysis & Reporting > Searches > Remediation Status**.

   The Remediation Status search page appears.



> **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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

Keysight, EX-1004
Page 1350

- 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.

Configuring Response Groups

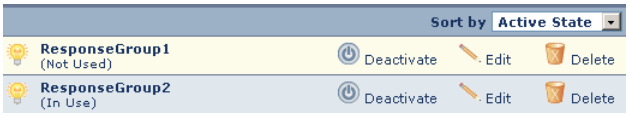
Platform: DC

Component: RNA

You can group alerts and remediations so that a policy violation triggers all of the responses within the group.

IMPORTANT! Response groups can only be used within compliance policies. You **cannot** assign them to launch in response to RNA events on the RNA Event Alerts page or impact flags on the Impact Flag Alerts page.

Before you can assign response groups to compliance policies, you must create the groups on the Groups page.



The light bulb icon next to the group indicates whether the group is active. If you want to assign a response group to a rule within a compliance policy, you must activate it. If the light bulb icon is lit, the group is active. If it is dark, the group is inactive. For more information, see Activating and Deactivating Response Groups on page 1353.

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 1352
- Modifying a Response Group on page 1353
- Deleting a Response Group on page 1353
- Activating and Deactivating Response Groups on page 1353

Keysight, EX-1004
Page 1351

Creating a Response Group

Platform: DC

Component: RNA

You can place individual alerts and remediations in response groups, which can then be assigned to rules within compliance 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:

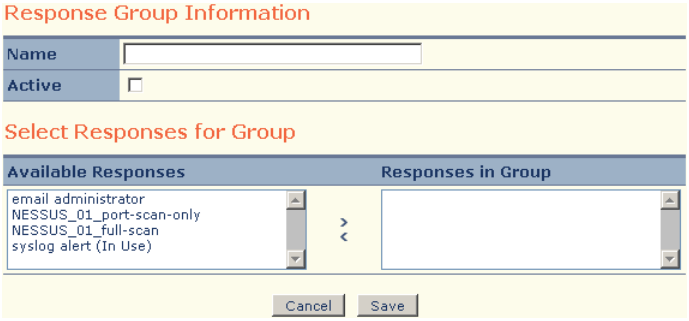
Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Groups**.

   The Groups page appears.

2. Click **Create Group**.

   The Groups page expands.



3. Under **Response Group Information**, in the **Name** field, enter a name for the new group.
4. Select **Active** to activate the group so that you can use it in response to a compliance policy violation.
5. Under **Select Responses for Group**, 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.

Keysight, EX-1004
Page 1352

Modifying a Response Group

Platform: DC

Component: RNA

Use the following procedure to modify a response group.

To modify a response group:

Access: P&R Admin/Admin

1. Select **Policy & Response > Responses > Groups**.

   The Groups page appears.

2. Click **Edit** next to the group you want to modify.

   The Groups page reloads with the information for the group you chose.

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

Platform: DC

Component: RNA

You can delete a response group as long as it is not in use within a compliance 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: P&R Admin/Admin

1. Select **Policy & Response > Responses > Groups**.

   The Groups page appears.

2. Click **Delete** next to the group you want to delete.

   The response group is deleted.

Activating and Deactivating Response Groups

Platform: DC

Component: RNA

If you want to temporarily disable a response group without deleting it, you can deactivate it. This leaves the group on the system but does not launch it when a

Keysight, EX-1004
Page 1353

policy to which the group is assigned is violated. To enable the group again, you must activate it.

IMPORTANT! Even if you deactivate a group, it is still considered in use and cannot be deleted unless you remove it from the compliance policies where it is applied.

To deactivate a response group:

Access: P&R Admin/Admin

- On the Groups page, click **Deactivate** next to the group you want to deactivate.

To activate a response group:

Access: P&R Admin/Admin

- On the Groups page, click **Activate** next to the group you want to activate.

Keysight, EX-1004
Page 1354

Chapter 30
Enhancing RNA Detection

The Sourcefire 3D System collects data from your network in a variety of ways, including IPS, RNA, scanners such as Nmap, and application, scanner, or user input through the *host input* feature, which uses an import file or an API interface to import data from a third-party application or allows users to set specific data for a host using the web interface. The collected information is most valuable to you when the Sourcefire 3D System can correlate the results of all possible sources of 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, RNA cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that RNA 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 Sourcefire 3D 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 service data does not automatically map to RNA vulnerability information. If you want to see vulnerabilities and perform impact correlation for hosts using third-party operating system and service data, you must map the vendor and version information from the third-party system to the vendor and version listed in RNA's vulnerability database (VDB). You also may want to maintain the host input data on an ongoing basis.

If RNA cannot identify services running on hosts on your network, you can create user-defined RNA detectors that allow RNA to identify the services based on a

Keysight, EX-1004
Page 1355

port or a pattern. You can also import, activate, and deactivate certain service and application detectors to further customize the service and application detection capability of the Sourcefire 3D System.

You can also replace RNA detection of operating system and service data using scan results from the Nmap active scanner or augment the RNA vulnerability lists with third-party vulnerabilities. RNA may reconcile data from multiple sources to determine the identity for a service. For more information on how RNA does this, see Understanding Current Identities on page 1360. For more information on active scanning, see Configuring Active Scanning on page 1418.

For more information, see the following sections:

- Assessing Your Detection Strategy on page 1356
- Enhancing Your Network Map on page 1358
- Using Custom Fingerprinting on page 1363
- Working with Service and Application Detectors on page 1381
- Importing Host Input Data on page 1406

Assessing Your Detection Strategy

Platform: DC

Component: RNA

Before you make any changes to RNA'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 Sensors Correctly Placed? on page 1356
- Do Unidentified Operating Systems Have a Unique TCP Stack? on page 1357
- Can RNA Identify All Services? on page 1358
- Have You Applied Patches that Fix Vulnerabilities? on page 1358
- Do You Want to Track Third-Party Vulnerabilities? on page 1358

Are Your Sensors Correctly Placed?

Platform: DC

Component: RNA

If network devices such as load balancers, proxy servers, or NAT devices reside between the 3D Sensor with RNA and the unidentified or misidentified host, place a sensor closer to the misidentified host rather than using custom fingerprinting. Sourcefire does not recommend using custom fingerprinting in this scenario.

Keysight, EX-1004
Page 1356

Do Unidentified Operating Systems Have a Unique TCP Stack?

Platform: DC

Component: RNA

If RNA 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 RNA 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 RNA by default (see Detected Operating Systems on page 2020 for a list of detected operating systems and kernels) 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 RNA cannot identify, you would benefit from creating a custom fingerprint, which allows RNA 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, RNA 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, etc.) 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 RNA 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 application uses to describe products to the RNA definitions for those products. For more information, see Managing Third-Party Product Mappings on page 1408.

RNA may reconcile data from multiple sources to determine the current identity for an operating system or service. For more information on how RNA does this, see Understanding Current Identities on page 1360.

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.

Keysight, EX-1004
Page 1357

However, note that active scan data and host input data may not be updated with the frequency of RNA data.

Can RNA Identify All Services?

Platform: DC

Component: RNA

If a host is correctly identified by RNA but has unidentified services, you can create a user-defined service detector to provide RNA with port and pattern matching information to help identify the service. For more information, see Creating a User-Defined Service Detector on page 1387.

Have You Applied Patches that Fix Vulnerabilities?

Platform: DC

Component: RNA

If RNA 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 RNA database. For more information, see Mapping Third-Party Product Fixes on page 1410.

Do You Want to Track Third-Party Vulnerabilities?

Platform: DC

Component: RNA

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 to vulnerability identifiers in the RNA 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 1412.

Enhancing Your Network Map

Platform: DC

Component: RNA

RNA 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 RNA decides which data to use for a service or operating system identity can help you decide how best to augment RNA's passive detection capabilities with active input sources.

Keysight, EX-1004
Page 1358

For more information, see the following topics:

- Understanding Passive Detection on page 1359
- Understanding Active Detection on page 1359
- Understanding Current Identities on page 1360
- Understanding Identity Conflicts on page 1362

Understanding Passive Detection

Platform: DC

Component: RNA

Passive detection is the detection of host operating system, client application, service, and payload information through analysis of traffic passively collected by RNA. RNA uses information in the VDB to help it identify your network assets.

If RNA cannot identify an operating system on a host, you can manually determine it and then create a custom server or client fingerprint to help RNA recognize that operating system on other hosts with similar operating system characteristics.

RNA uses all collected passive fingerprints for a host operating system to create a *derived fingerprint*. RNA 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 custom services on your network, you can augment RNA's service detection capabilities by creating custom service detectors that provide RNA with the information it needs to identify those services. NetFlow can also add passively detected service information to the network map.

Note that RNA does not use service and operating system data that it classified as *unknown* because it is unable to interpret the data. The sensor reports the identity to the Defense Center as `unknown` and the identity data is not used to derive fingerprints.

Understanding Active Detection

Platform: DC

Component: RNA

Active detection is addition, to the network map, of data collected by active sources, such as host operating system and service 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 services on hosts.

Keysight, EX-1004
Page 1359

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 service 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*.

RNA 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 an application whose results are imported through the command line, RNA retains both the identities from the Nmap results and the identities from the import application. Then RNA 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

Platform: DC

Component: RNA

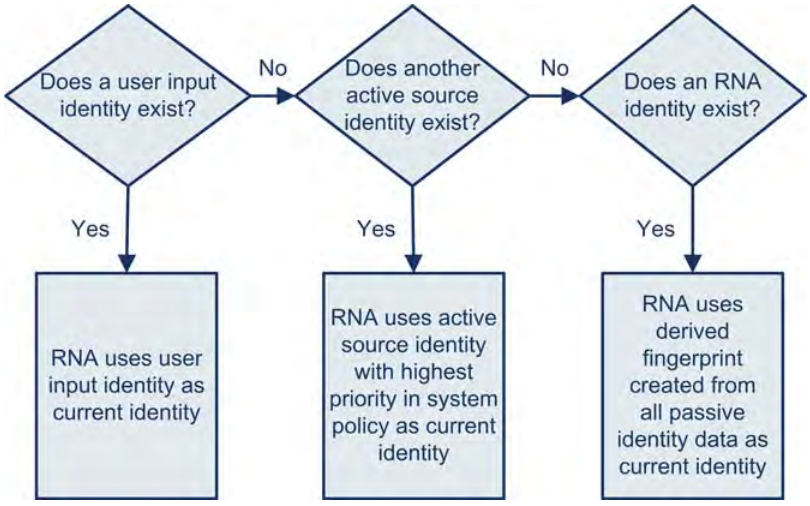
The *current identity* for a service or an operating system on a host is the identity that RNA finds most likely to be correct.

RNA uses the current identity for an operating system or service for the following purposes:

- to assign vulnerabilities to a host
- for impact assessment
- when evaluating compliance rules written against operating system identifications, host profile qualifications, and compliance white lists
- for display in the Hosts and Services table views in workflows
- for display in the host profile
- to calculate the operating system and service statistics on the RNA Statistics page

Keysight, EX-1004
Page 1360

RNA uses source priorities to determine which active identity should be used as the current identity for a service 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 RNA database may retain information from several sources for the operating system or for a particular service on a host.

RNA treats an operating system or service 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 system policy)
3. RNA
4. NetFlow

Note that a new higher priority service identity will not be override a current service 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 system policy or on your manual resolution, as described in Understanding Identity Conflicts on page 1362.

Keysight, EX-1004
Page 1361

Understanding Identity Conflicts

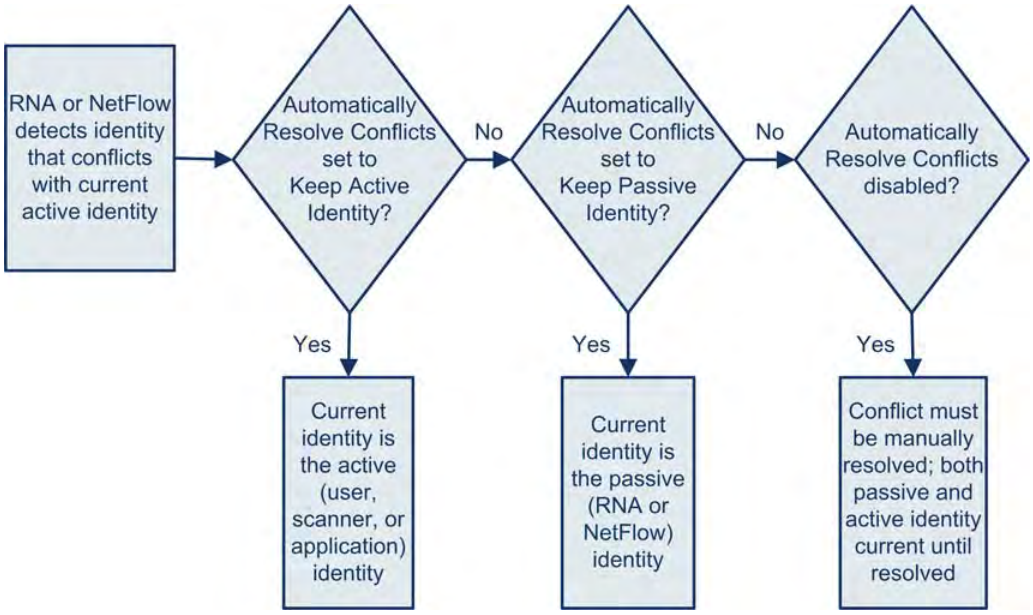
Platform: DC

Component: RNA

An *identity conflict* occurs when RNA 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, RNA 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 services on the host, RNA 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 (RNA) identity or always use the active identity. Unless you disable automatic resolution of identity conflicts in the system policy, identity conflicts are always automatically resolved.



A user with Administrator privileges can also configure RNA to generate an event when an identity conflict occurs. That user can then set up a compliance policy with a compliance rule that uses an Nmap scan as a compliance response. When an event occurs, Nmap scans the host to obtain updated host operating system and service data.

Keysight, EX-1004
Page 1362

Using Custom Fingerprinting

Platform: DC

Component: RNA

The Sourcefire 3D System includes operating system *fingerprints* that RNA uses to identify the operating system on each host it detects. However, sometimes RNA 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 RNA cannot match a host's operating system, it cannot identify the vulnerabilities for the host, because RNA derives the list of vulnerabilities for each host from its operating system fingerprint. For example, if RNA detects a host running Microsoft Windows, RNA 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, RNA cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that RNA 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 RNA 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, RNA uses the fingerprint to assign the custom operating system information you provide in the fingerprint. When RNA sees new traffic from a host that has already been detected and currently resides in the network map, RNA updates the host with the new fingerprint information. RNA 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 1356.

Keysight, EX-1004
Page 1363

You can create two types of fingerprints with RNA:

- Client fingerprints, which identify operating systems based on the SYN packet that the host sends when it connects to a TCP service running on another host on the network.

  See Fingerprinting Clients on page 1364 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 service.

  See Fingerprinting Servers on page 1370 for information about how to obtain a server fingerprint for a host.

After creating fingerprints, you must activate them before RNA can associate them with hosts. See Managing Fingerprints on page 1377 for more information.

IMPORTANT! If both a client and server fingerprint match the same host, the client fingerprint is used.

Fingerprinting Clients

Platform: DC

Component: RNA

Client fingerprints identify operating systems based on the SYN packet a host sends when it connects to a TCP service running on another host on the network.

If the Defense Center does not have direct contact with monitored hosts, you can specify a 3D Sensor 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 3D Sensor you use to obtain the fingerprint. (Sourcefire highly recommends that you directly connect the Defense Center or the 3D Sensor to the same subnet that the host is connected to.)
- The network interface (on the Defense Center or the 3D Sensor) 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.

Keysight, EX-1004
Page 1364

To obtain a client fingerprint for a host:

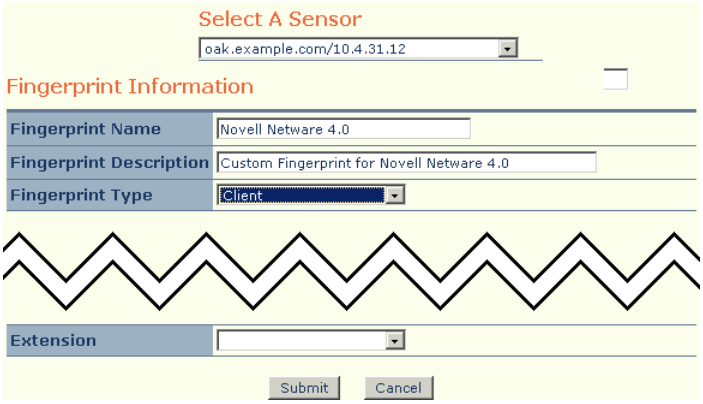
Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

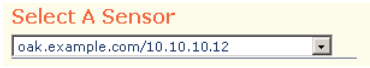
   The Custom Fingerprinting page appears.

2. Click **Create Custom Fingerprint**.

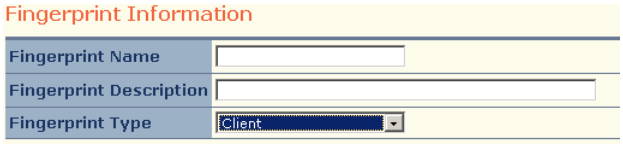
   The Create Custom Fingerprint page appears.



3. From the **Select A Sensor** list, select the Defense Center or the 3D Sensor that you want to use to collect the fingerprint.



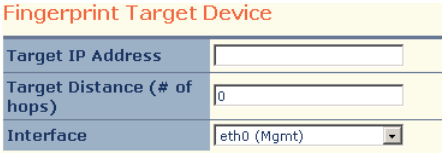
4. In the **Fingerprint Name** field, type an identifying name for the fingerprint.



5. In the **Fingerprint Description** field, type a description for the fingerprint.
6. From the **Fingerprint Type** list, select **Client**.

Keysight, EX-1004
Page 1365

7. In the **Target IP Address** field, type the IP address of the host you want to fingerprint.



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.

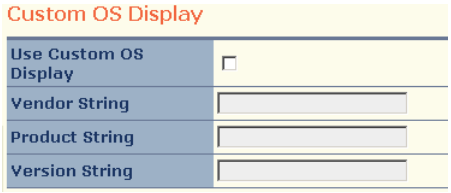
WARNING! This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by RNA.

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 3D Sensor 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 sensor, the sensor 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 sensor, refer to the *Installation Guide* for the specific model you are using to fingerprint.

Keysight, EX-1004
Page 1366

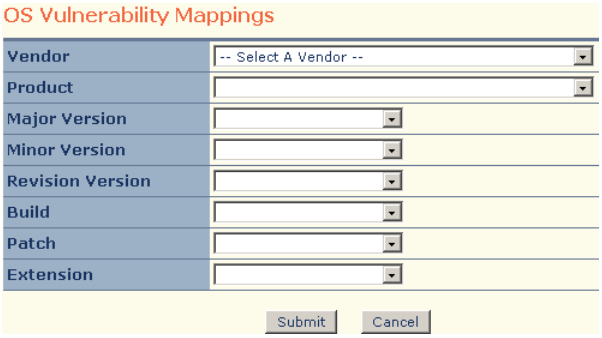
10. If you want to display custom information in the host profile for fingerprinted hosts (or if the host you want to fingerprint does not reside in the OS Vulnerability Mappings section), select **Use Custom OS Display** in the Custom OS Display section.



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.

11. In the OS Vulnerability Mappings section, select the operating system, product, and versions you want to use for vulnerability mapping.



The following lists the options, some of which may not be applicable:

- Vendor (required if you want to map vulnerabilities for the host or if you do not specify a custom operating system display)
- Product (required if you select a vendor)
- Major Version
- Minor Version
- Revision Version
- Build
- Patch
- Extension

Keysight, EX-1004
Page 1367

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.

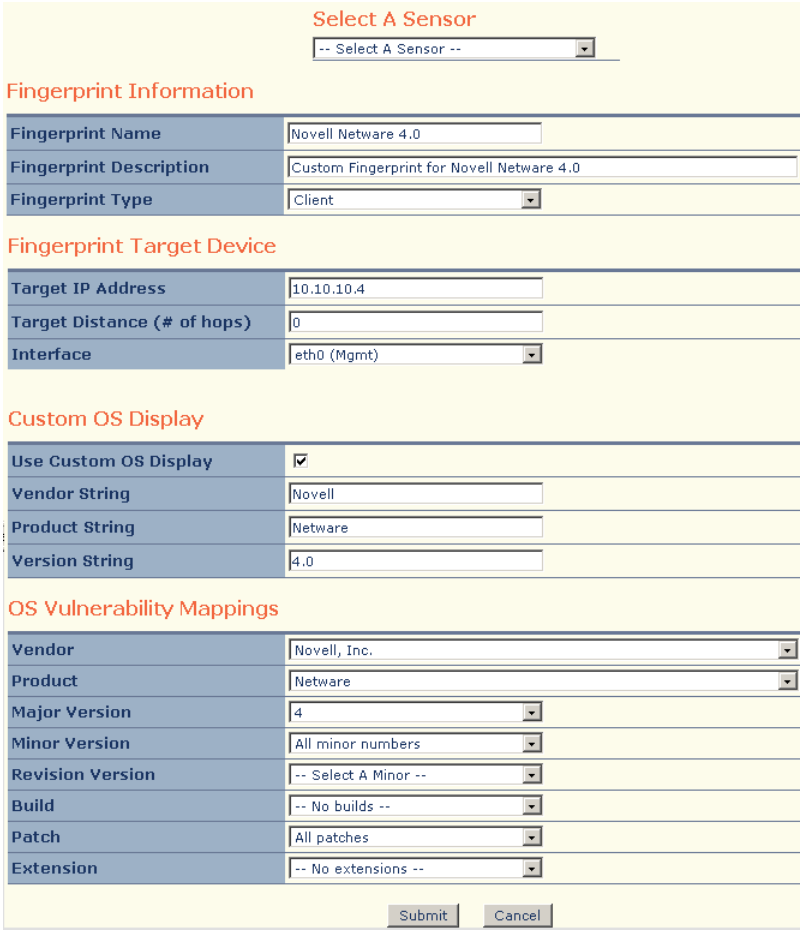
TIP! When creating a fingerprint, you assign a single vulnerability mapping for the fingerprint. After the fingerprint is created and activated, you can add additional vulnerability mappings for other versions of the operating system. See Editing an Active Fingerprint on page 1380 for more information.

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, RNA cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

The following graphic shows an example of a completed client fingerprint:

Keysight, EX-1004
Page 1368



12. Click **Submit**.

    The Custom Fingerprinting status page re-appears. The status page refreshes every ten seconds until it receives data from the host in question.

TIP! When you click **Submit**, 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`.

Keysight, EX-1004
Page 1369

13. Access the host you are trying to fingerprint and initiate a TCP connection to the Defense Center or managed 3D Sensor.

    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. The Custom Fingerprinting page should then reload with a "Ready" status.

IMPORTANT! To create an accurate fingerprint, traffic **must** be seen by the Defense Center or 3D Sensor collecting the fingerprint. If you are connected through a switch, traffic to a system other than the Defense Center or sensor may not be seen by RNA.

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 1377 for more information.

Fingerprinting Servers

Platform: DC

Component: RNA

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 service. 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 Defense Center or the 3D Sensor you use to obtain the fingerprint. Sourcefire highly recommends that you directly connect an unused interface on the Defense Center or the 3D Sensor to the same subnet that the host is connected to.
- The network interface (on the Defense Center or the 3D Sensor) 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 3D Sensor that is managed by the Defense Center and is closest to the host you intend to fingerprint when specifying server fingerprint properties.

Keysight, EX-1004
Page 1370

To obtain a server fingerprint for a host:

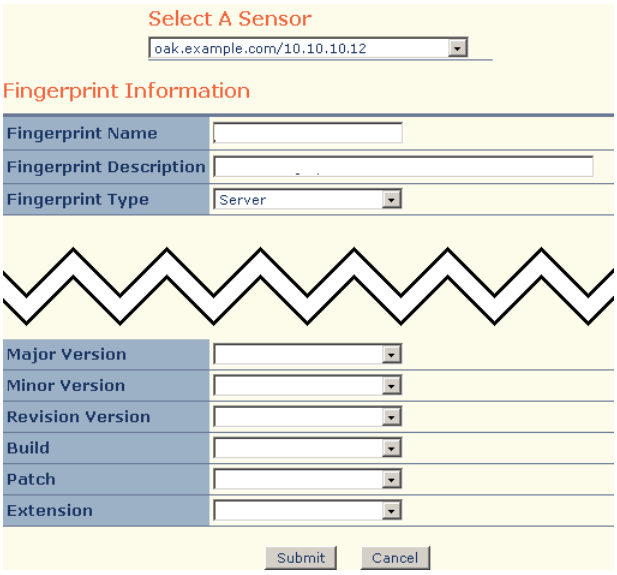
Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

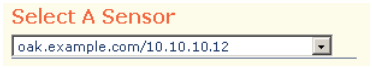
   The Custom Fingerprinting page appears.

2. Click **Create Custom Fingerprint**.

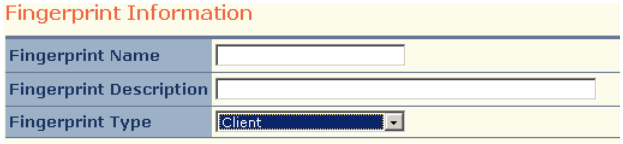
   The Create Custom Fingerprint page appears.



3. From the **Select A Sensor** list, select the Defense Center or the 3D Sensor that you want to use to collect the fingerprint.



4. In the **Fingerprint Name** field, type an identifying name for the fingerprint.

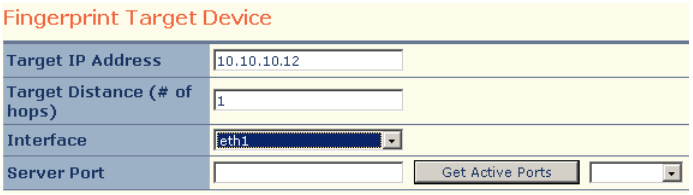


5. In the **Fingerprint Description** field, type a description for the fingerprint.

Keysight, EX-1004
Page 1371

6. From the **Fingerprint Type** list, select **Server**.

   Server fingerprinting options appear.



7. In the **Target IP Address** field, type the IP address of the host you want to fingerprint.
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 RNA.

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 3D Sensor 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 sensor, the sensor 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 sensor, refer to the *Installation Guide* for the specific model you are using to fingerprint.

10. Click **Get Active Ports**.

    If RNA 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).

Keysight, EX-1004
Page 1372

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.

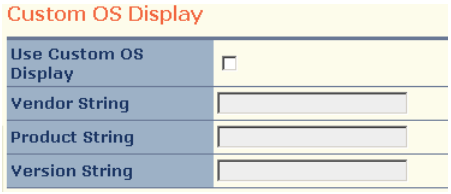
In addition, you should exclude that IP address from monitoring in your RNA detection policy while you create the fingerprint. Otherwise, the network map and RNA event views will be cluttered with inaccurate information about the host represented by that IP address. For more information, see What is an RNA Detection Policy? on page 889.

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.



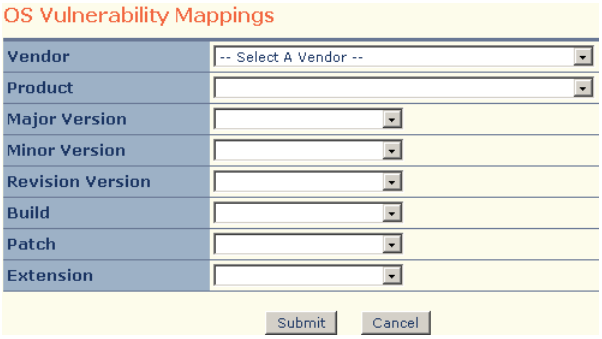
Keysight, EX-1004
Page 1373

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 from the following lists (if applicable):
    - **Vendor**
    - **Product**
    - **Major Version**
    - **Minor Version**
    - **Revision Version**
    - **Build**
    - **Patch**
    - **Extension**

    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.



Keysight, EX-1004
Page 1374

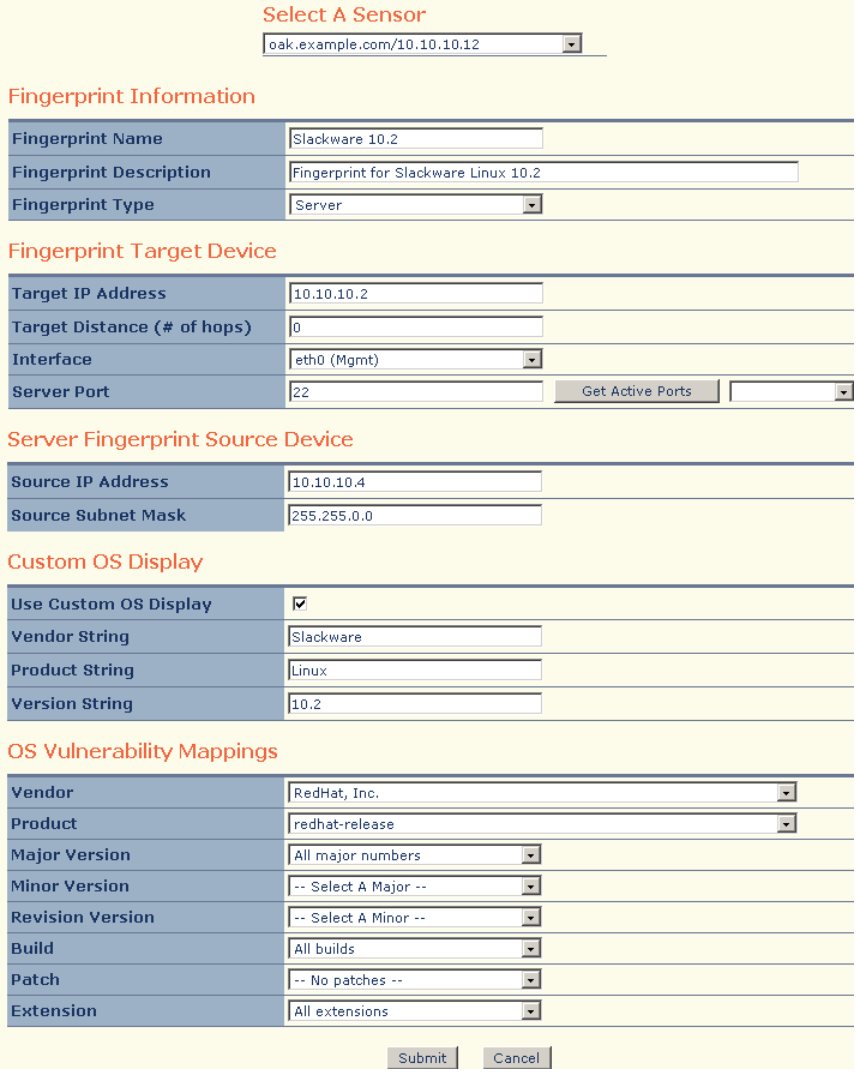
TIP! When creating a fingerprint, you assign a single vulnerability mapping for the fingerprint. After the fingerprint is created and activated, you can add additional vulnerability mappings for other versions of the operating system. See Editing an Active Fingerprint on page 1380 for more information.

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, RNA cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

The following graphic shows an example of a completed server fingerprint:

Keysight, EX-1004
Page 1375



17. Click **Submit**.
18. The Custom Fingerprinting status page appears. It reloads every ten seconds and should reload with a "Ready" status.

    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 **Edit** to access the Custom Fingerprinting page, and then click **Submit**.
19. After the fingerprint is created, activate it and, optionally, add vulnerability mappings. See Managing Fingerprints on page 1377 for more information.

Keysight, EX-1004
Page 1376

Managing Fingerprints

Platform: DC

Component: RNA

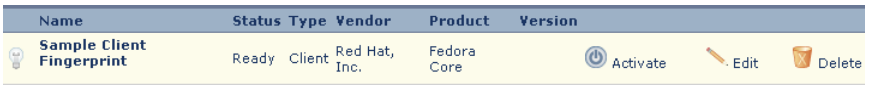
You can activate, deactivate, delete, view, and edit custom fingerprints. When creating a fingerprint, you assign a single vulnerability mapping for the fingerprint. For more information on creating a fingerprint, see Fingerprinting Clients on page 1364 and Fingerprinting Servers on page 1370. After 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: P&R Admin/Admin

▶ Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

The Custom Fingerprinting page appears.



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 1377
- Deactivating Fingerprints on page 1378
- Deleting Fingerprints on page 1378
- Editing Fingerprints on page 1379

Activating Fingerprints

Platform: DC

Component: RNA

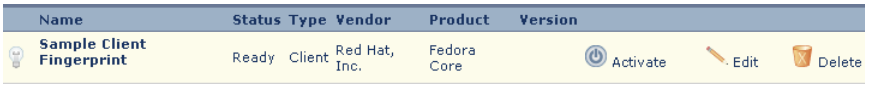
After creating a custom fingerprint, you must activate it before RNA can use it to identify hosts. After the new fingerprint is activated, RNA uses it to re-identify previously discovered hosts and discover new hosts.

To activate a fingerprint:

Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

   The Custom Fingerprinting page appears.



Keysight, EX-1004
Page 1377

2. Click **Activate** next to the fingerprint you want to activate.

IMPORTANT! The **Activate** option is only available if the fingerprint you created is valid. If **Activate** is not available, try creating the fingerprint again.

The Defense Center activates the fingerprint and propagates it to all managed 3D Sensors. The light bulb icon next to the fingerprint name changes to a lit light bulb icon.

Deactivating Fingerprints

Platform: DC

Component: RNA

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: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

   The Custom Fingerprinting page appears.



2. Click **Deactivate** next to the active fingerprint you want to deactivate.

   The Defense Center deactivates the fingerprint and propagates the deactivation to all managed 3D Sensors.

Deleting Fingerprints

Platform: DC

Component: RNA

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.

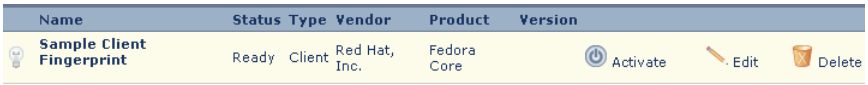
Keysight, EX-1004
Page 1378

To delete a fingerprint:

Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

   The Custom Fingerprinting page appears.



2. If the fingerprints you want to delete are active, click Deactivate next to each one to deactivate it.
3. Click **Delete** next to the fingerprint.
4. Click **OK** to confirm that you want to delete the fingerprint.

   The fingerprint is deleted.

Editing Fingerprints

Platform: DC

Component: RNA

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 re-send traffic to the Defense Center or managed 3D Sensor 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 1380
- Editing an Active Fingerprint on page 1380

Keysight, EX-1004
Page 1379

Editing an Inactive Fingerprint

Platform: DC

Component: RNA

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: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

   The Custom Fingerprinting page appears.

2. Click **Edit** 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 1364 for more information about the options you can configure.
   - If you are modifying a server fingerprint, see Fingerprinting Servers on page 1370 for more information about the options you can configure.

4. Click **Submit** to re-submit the fingerprint.

IMPORTANT! If you modified a client fingerprint, remember to send traffic from the host to the Defense Center or 3D Sensor gathering the fingerprint.

Editing an Active Fingerprint

Platform: DC

Component: RNA

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: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **Custom Fingerprinting**.

   The Custom Fingerprinting page appears.

2. Click **Edit** next to the fingerprint you want to edit.

   The Edit Custom Fingerprint Product Mappings page appears.

Keysight, EX-1004
Page 1380

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**.

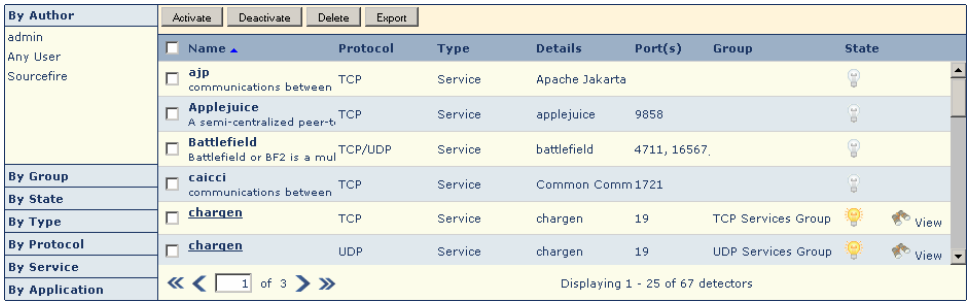
   The vulnerability mapping is added to the Pre-Defined OS Product Maps list.
6. Click **Save** to save your changes.

Working with Service and Application Detectors

Platform: DC

Component: RNA

When RNA analyzes IP traffic, it uses RNA detectors to identify the commonly used services and client applications running on hosts on your network. You use the RNA detectors page to customize the service and application detection capability of the Sourcefire 3D System.



The page provides information about each detector, including:

- the protocol (TCP, UDP, or both) of traffic that the detector inspects
- the group, if any, that the detector belongs to

  You can create new groups and move detectors between groups.
- the author of the detector

  Sourcefire-provided detectors have an author of **Sourcefire**. If you created a user-defined service detector, you are the author. You are also the author for any detector that you import or that you edit and save.
- whether the detector is a service or application detector

Keysight, EX-1004
Page 1381

- for service detectors, the service that the detector detects, the type of service detector (port or pattern), and the ports, if any, associated with the detector

  Port service detectors base service identification on the port that the service uses. *Pattern* service detectors base service identification on ASCII or hexadecimal patterns in the packet header. The Sourcefire 3D System considers detectors that identify service traffic by both port and pattern to be pattern detectors.
- for application detectors, the application that the detector detects as well as that application's type

  Application types categorize applications; for example, **database** and **instant messenger** are application types.
- the state (active or inactive) of the detector

  RNA uses only active detectors to analyze service traffic.

IMPORTANT! Although you can create and configure RNA detectors on a Defense Center that does not have an RNA host license installed, you must add the license before RNA can use the detectors; see Managing Feature Licenses on page 1787.

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

Keysight, EX-1004
Page 1382

delete some detectors but not others. This is because there are several different types of Sourcefire-provided detectors, as described in the following table.

RNA Detector Types

Detector Type	Description	Examples
internal	Service and application detectors that are delivered with updates to the Sourcefire 3D System. Internal detectors are always on; you cannot deactivate, delete, or otherwise configure them. **TIP!** To view a list of internal detectors, click **View RNA Internal Detectors**; see Viewing Internal RNA Detectors on page 1401 for more information.	aim (service) PuTTY (application)
add-on	Service and application detectors that are provided either via VDB updates or as individual importable detectors authored by Sourcefire Professional Services. Detectors provided through these delivery methods allow you to enhance RNA's detection capabilities with new pattern-based detectors without updating the Sourcefire 3D System as a whole. You can group, activate, and deactivate add-on detectors according to the needs of your organization. If you imported an individual detector, you can delete or export it. Finally, VDB updates may activate or deactivate add-on detectors. **IMPORTANT!** To improve performance, Sourcefire sometimes replaces add-on detectors with internal detectors delivered through updates to the Sourcefire 3D System.	Tivoli (service) TDS (application)
well-known port	Service detectors, provided by Sourcefire, that are based on detection of network traffic on well-known ports. You can group, activate, and deactivate well-known port detectors according to the needs of your organization. You can also view their properties and save copies that you can then edit and use. Note that by default, well-known port detectors are organized into two groups (**TCP Services Group** and **UDP Services Group**), based on their protocols.	finger
user-defined	Service detectors that you create, based on a port, pattern, or both. You have full control over user-defined detectors; you can group, enable, disable, edit, import, export and delete them.	any user-defined service detector

For a full list of the services and applications that the Sourcefire 3D System detects, see Detected Operating Systems and Services on page 2019. Keep in mind, however, that the detector list can change depending on the version of the Sourcefire 3D System and the VDB you have installed, as well as on any individual add-on detectors you may have imported. You should carefully read the release

Keysight, EX-1004
Page 1383

notes for each Sourcefire 3D System update as well as the advisories for each VDB update for information on updated RNA detectors.

For more information, see:

- Understanding RNA Service Detection on page 1384
- Creating a User-Defined Service Detector on page 1387
- Managing RNA Detectors on page 1394
- Working with RNA Detector Groups on page 1402

Understanding RNA Service Detection

Platform: DC

Component: RNA

When RNA inspects network traffic, it can identify service traffic in one of two ways:

- *Port* service detectors base service identification on the port that the traffic uses.
- *Pattern* service detectors base service identification on ASCII or hexadecimal patterns in the packet header.

Some service detectors use both port and pattern detection to increase the likelihood of correctly identifying traffic for a particular service. The following table lists the detection strategies for each of the types of RNA service detector.

RNA Service Detection Strategies

Detector Type	Service Detection Strategies
internal	port and pattern
add-on	port, pattern, or port and pattern
well-known port	port
user-defined	port, pattern, or port and pattern

When RNA detects service traffic, it first determines whether the service is running on a port configured in either:

- a Sourcefire-provided well-known port detector
- a user-defined service detector that uses a specific port as its only criterion

Keysight, EX-1004
Page 1384

If the service is running on one of those ports, RNA positively identifies the service.

IMPORTANT! Because you can create and activate user-defined port-based service detectors on ports used by Sourcefire-provided internal detectors, it is possible to override Sourcefire's detection capabilities. For example, if your user-defined service detector identifies all service traffic on port 22 as the `myservice` service, SSH traffic on port 22 will be misidentified as `myservice` traffic.

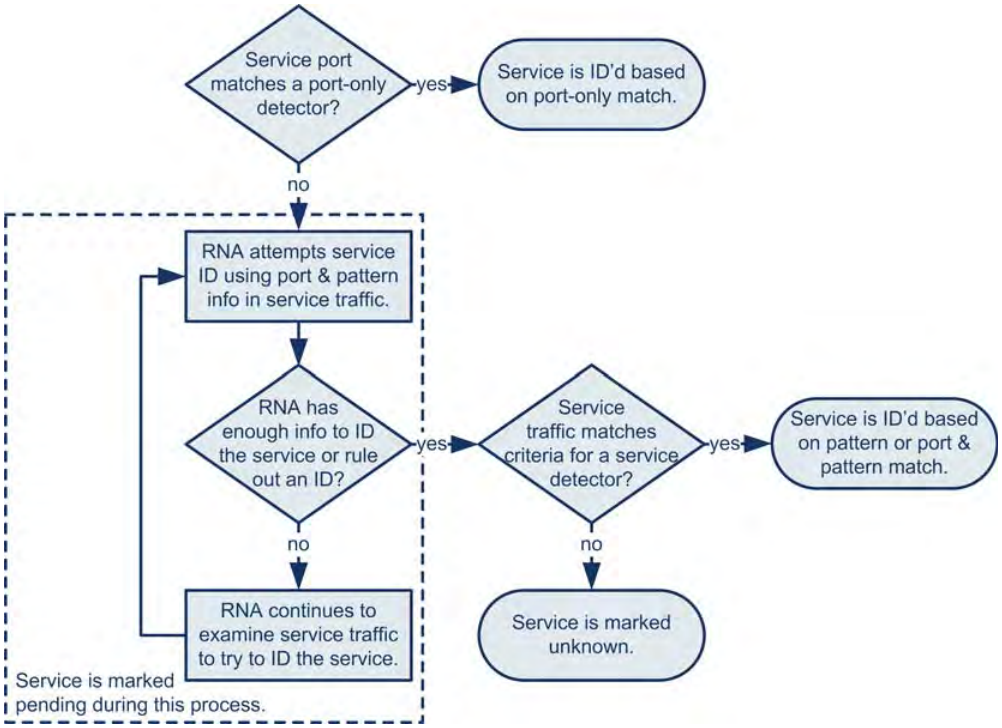
If the service is not running on one of those ports, RNA employs a more robust method to identify the service based on port and pattern matches. If two detectors both positively identify service traffic, the detector that employs the longer pattern match has precedence. Similarly, detectors with multiple pattern matches have precedence over single pattern matches.

Note that RNA identifies only those services running on hosts in your monitored networks, as defined in the RNA detection policy. For example, if an internal host accesses an FTP server on a remote site that you are not monitoring, RNA does not identify the service as FTP. On the other hand, if a remote or internal host accesses an FTP server on a host you are monitoring, RNA can positively identify the service. An exception occurs if RNA can identify the client application used in flows between a monitored host accessing a non-monitored server. In that case, RNA positively identifies the appropriate service that corresponds with the application in the flow, but does not add the service to the network map. For more information, see Implied Service Detection from Client Application Detection on page 2029.

The following diagram shows how the Sourcefire 3D System identifies service traffic. Note that the Sourcefire 3D System does not use this process for services

Keysight, EX-1004
Page 1385

added to the network map based on NetFlow or Nmap data. This process also does not apply to service identities you set using the host input feature.



Keysight, EX-1004
Page 1386

The following table outlines how the Sourcefire 3D System identifies detected services in the Defense Center web interface: the network map, host profiles, event views, and so on.

Sourcefire 3D System Identification of Services

Service Identification	Description
the service name	The Defense Center identifies a service with its name if the service was: • positively identified by RNA • added to the network map based on NetFlow data and there is a port-service correlation in `/etc/sf/services` • manually identified using the host input feature • identified by Nmap or another active source
`pending`	The Defense Center identifies a service as `pending` if RNA can neither positively nor negatively identify the service. Most often, RNA needs to collect and analyze more flow data (from which services are identified) before it can identify a pending service, but a service may also appear as pending for one of the following reasons: • the service was added to the network map based on NetFlow data and there is no port-service correlation in `/etc/sf/services` • you added the service to the network map using the host input feature, but did not specify a service name • service identities set by active sources have timed out (see Understanding Multiple Fingerprint Settings on page 1762), and RNA can neither positively nor negatively identify the service • for services identified in flows (see Understanding Flow Data Tables on page 1086), a monitored host accessed a non-monitored server and RNA cannot imply the service identity from the client application used in the flow
`unknown`	The Defense Center identifies a service as `unknown` if the service does not match any of RNA's detectors.

Creating a User-Defined Service Detector

Platform: DC

Component: RNA

If you use custom services on your network, you can create user-defined service detectors that provide RNA with the information it needs to identify those services. You can base service detection on the port or ports used by service traffic, patterns within the traffic, or on both ports and patterns.

For example, if you expect traffic for a custom service to use port 1180, you can create a service detector that detects traffic on that port. As another example, if

Keysight, EX-1004
Page 1387

you know that the header for any packet containing service traffic has a string of `ServiceName` in it, you can create a detector that registers the ASCII string of `ServiceName` as a pattern to match.

You must configure user-defined service detectors to use either a port or a pattern match; you cannot create a detector that uses neither. However, if you know both the port and a pattern for the service traffic, you can create a detector that uses both criteria; this increases the likelihood of correctly identifying traffic for that service.

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 individual add-on 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 1997.

To create a user-defined service detector:

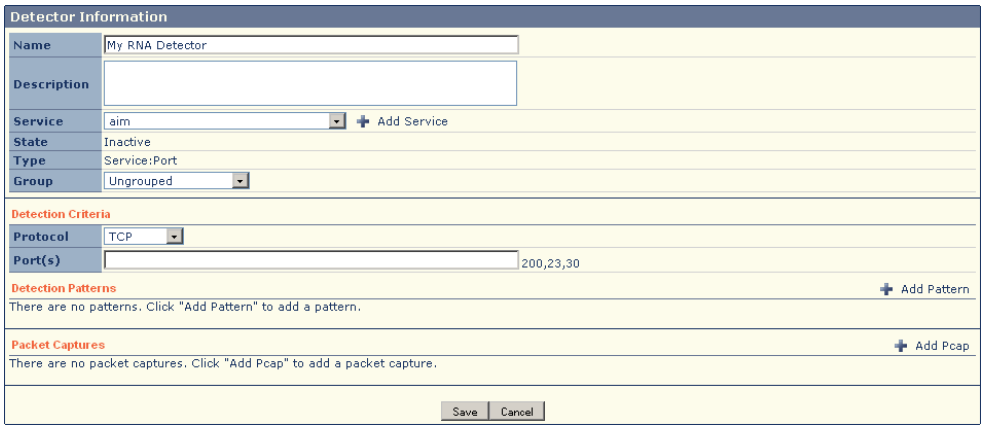
Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA Detectors page appears.

2. Click **Create Detector**.

   The Create RNA Detector page appears.



3. Provide basic detector information, such as the detector name, description, and group.

   See Providing Basic Service Detector Information on page 1389.

Keysight, EX-1004
Page 1388

4. Provide detection criteria, including the protocol of traffic the detector should inspect and the port that the traffic uses.

   See Specifying Detection Criteria for Service Detectors on page 1390.

5. Optionally, configure the detector to inspect traffic for matches to one or more patterns that occurs in traffic for that service.

   See Adding Detection Patterns to a Service Detector on page 1391.

6. Optionally, test the new detector against the contents of one or more PCAP files.

   See Testing a Service Detector Against Packet Captures on page 1392.

7. Click **Save**.

   The service detector is saved.

IMPORTANT! You must activate the detector before RNA can use it to analyze service traffic. For more information, see Activating and Deactivating RNA Detectors on page 1398.

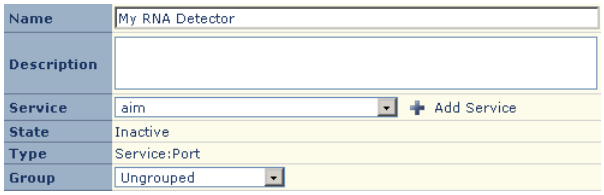
Providing Basic Service Detector Information

Platform: DC

Component: RNA

You must give each user-defined service detector a name, as well as identify the service you want to detect. Optionally, you can provide a brief description of the detector as well as add it to a group.

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 service traffic by port and pattern, the Sourcefire 3D System considers it a pattern detector.



If you are editing an existing detector, the Defense Center also displays the detector's author. If you created a user-defined service detector, you are the author. You are also the author for any detector that you import or that you edit and save.

Keysight, EX-1004
Page 1389

To provide basic service detector information:

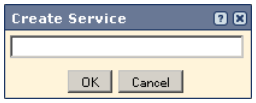
Access: P&R Admin/Admin

1. On the Create RNA Detector page, in the **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 service you want to detect. You have the following options:

   - If you are creating a detector for an existing service (for example, if you want to detect a particular service on a non-standard port), select the service from the **Service** drop-down list.
   - If you are creating a detector for a custom service, click **Add Service**. In the Create Service pop-up window that appears, type a name for the service and click **OK**.

3. Optionally, select a group for the detector from the **Group** drop-down list.

   By default, new detectors are ungrouped. For more information, see Working with RNA Detector Groups on page 1402.

4. Continue with the procedure in the next section, Specifying Detection Criteria for Service Detectors.

Specifying Detection Criteria for Service Detectors

Platform: DC

Component: RNA

When creating a user-defined service 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 a Service Detector on page 1391.

Keysight, EX-1004
Page 1390

To specify detection criteria for a service detector:

Access: P&R Admin/Admin

1. On the Create RNA 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 service traffic based on the port it uses, type the port 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 service detector to inspect traffic for matches to one or more patterns that occurs in traffic for that service, continue with the procedure in the next section, Adding Detection Patterns to a Service Detector.
   - If you want to test the new detector against the contents of one or more PCAP files, skip to Testing a Service Detector Against Packet Captures on page 1392.
   - If you are done creating the detector, click **Save**.

     The service detector is saved.

IMPORTANT! You must activate the detector before RNA can use it to analyze service traffic. For more information, see Activating and Deactivating RNA Detectors on page 1398.

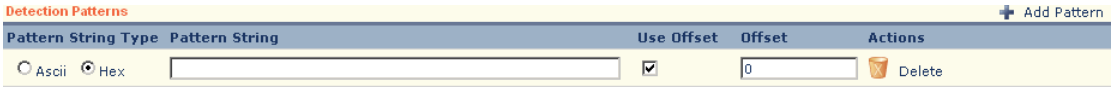
Adding Detection Patterns to a Service Detector

Platform: DC

Component: RNA

If you know that the header for any packet containing service traffic contains a particular pattern string, you can configure a user-defined service detector to search for that pattern.

Service 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 service traffic must match all of the patterns for the detector to positively identify the service.



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 Service Detectors on page 1390.

Keysight, EX-1004
Page 1391

To add a detection pattern to a service detector:

Access: P&R Admin/Admin

1. On the Create RNA Detector page, in the Detection Patterns section, click **Add Pattern**.
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 RNA should begin searching for the pattern; this is called the offset.

   To specify an offset, select the **Use Offset** check box, then 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 **Delete** 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 a Service Detector Against Packet Captures.
   - If you are done creating the detector, click **Save**.

     The service detector is saved.

IMPORTANT! You must activate the detector before RNA can use it to analyze service traffic. For more information, see Activating and Deactivating RNA Detectors on page 1398.

Testing a Service Detector Against Packet Captures

Platform: DC

Component: RNA

If you have a packet capture (PCAP) file that contains packets with traffic from the service you want to detect, you can test a user-defined service 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.

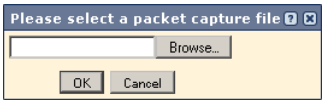
Keysight, EX-1004
Page 1392

To test a service detector against a PCAP file:

Access: P&R Admin/Admin

1. On the Create RNA Detector page, in the Packet Captures section, click **Add Pcap**.

   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 **Evaluate** 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 **Delete** next to the file you want to delete.

5. To save the detector, click **Save**.

IMPORTANT! You must activate the detector before RNA can use it to analyze service traffic. For more information, see Activating and Deactivating RNA Detectors on page 1398.

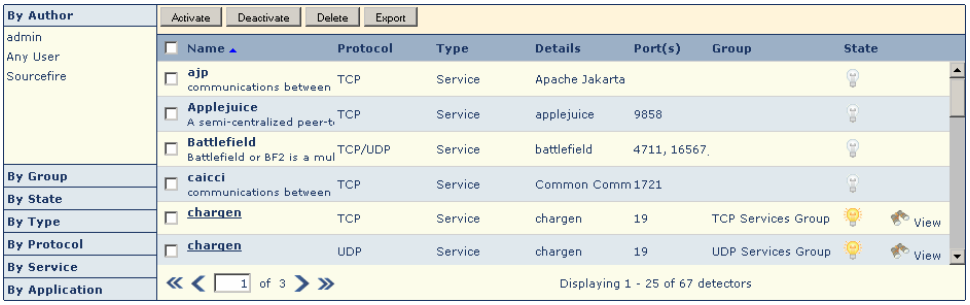
Keysight, EX-1004
Page 1393

Managing RNA Detectors

Platform: DC

Component: RNA

You view and manage RNA detectors on the RNA Detectors page.



From the RNA Detectors page, you can:

- sort, filter, and browse the list of detectors
- view a list of the Sourcefire-provided internal detectors
- view the properties of the Sourcefire-provided well-known port service detectors, and optionally save copies as new, user-defined detectors that you can modify
- create, modify, delete, and export user-defined service detectors
- delete and export any add-on detectors you individually imported
- create, modify, and delete detector groups
- activate and deactivate detectors

Note that you cannot modify internal or add-on detectors. In addition, you cannot delete Sourcefire-provided detectors, aside from any add-on detectors you individually imported.

For more information, see:

- Sorting the RNA Detector List on page 1395
- Filtering the RNA Detector List on page 1395
- Navigating to Other RNA Detector Pages on page 1397
- Activating and Deactivating RNA Detectors on page 1398
- Modifying Service Detectors on page 1399
- Deleting RNA Detectors on page 1400
- Viewing Internal RNA Detectors on page 1401

Keysight, EX-1004
Page 1394

Sorting the RNA Detector List

Platform: DC

Component: RNA

By default, the RNA 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 RNA detectors:

Access: P&R Admin/Admin

- On the RNA 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 RNA Detector List

Platform: DC

Component: RNA

You can filter the RNA detectors you display on the RNA Detectors page by a single criterion, or a combination of multiple criteria. The filter you construct is shown at the top of the page.



Keysight, EX-1004
Page 1395

The following table describes the filters you can apply to detectors.

RNA Detector Filter Groups

Filter Group	Description
By 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
By Group	Finds detectors according to which group they are in. You can also display ungrouped detectors. For more information, see Working with RNA Detector Groups on page 1402.
By State	Finds detectors according to their state, that is, **Active** or **Inactive**. For more information, see Activating and Deactivating RNA Detectors on page 1398.
By Type	Finds detectors according to the detector type: **Service** or **Application**. Service detectors have three subtypes you can use to further filter detectors: • **Port** service detectors include the Sourcefire-provided well-known port detectors, as well as any port-based user-defined service detectors. • **Pattern** service detectors include pattern-based or port-and-pattern-based user-defined service detectors. • **RNA** service detectors are the add-on service detectors. Application detectors only have one subtype: **RNA**; these are the add-on application detectors. For more information on detector types, see the RNA Detector Types table on page 1383.
By Protocol	Finds detectors according to which traffic protocol the detector inspects. Detectors can inspect TCP, UDP, or TCP and UDP traffic.
By Service	Finds detectors according to the service they detect. These services include the services detected by RNA by default (see the Detected Services table on page 2030), as well as any custom services you added when creating user-defined detectors.
By Application	Finds detectors according to the application type. For a list of application types that RNA detects, see Detected Client Applications on page 2037.

Keysight, EX-1004
Page 1396

Access: P&R Admin/Admin

The following table explains how to apply and remove filters to the RNA detector list.

Filtering RNA Detectors

To...	You can...
apply a filter	click the filter group you want to use to filter the detectors, then click the specific criterion you want to use. If the filter you are using has subfilters, you can click the subfilter to further filter the detectors.
remove a filter	click the delete icon (✖) next to the filter you want to remove.
remove all filters	click **Clear all** next to the list of filters applied to the detectors.

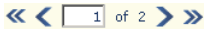
Navigating to Other RNA Detector Pages

Platform: DC

Component: RNA

Access: P&R Admin/Admin

The RNA 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.



Navigating RNA 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 double right arrow icon (>>).
jump to the first page	click the double left arrow icon (<<).

Keysight, EX-1004
Page 1397

Activating and Deactivating RNA Detectors

Platform: DC

Component: RNA

You must activate an RNA detector before you can use it to analyze network traffic.

Note that you can activate only one port-based service detector for each port. You can, however, activate multiple service detectors for a port if the detectors also use a pattern or patterns to determine service identity.

For example, the finger service detector is a well-known port service detector delivered with the Sourcefire 3D System; it identifies traffic on port 79 as finger traffic. This detector is active by default.

You can create a user-defined service detector for port 79, but you cannot activate it and use it to identify service traffic unless you either:

- deactivate the Sourcefire-provided finger detector
- add a pattern to your user-defined detector

Note that because the internal service detectors are port and pattern based, you can activate a user-defined port-based service detector that detects service traffic on the same port as one of the internal detectors. For example, the SSH internal detector uses port 22. You can create and activate a user-defined service detector on port 22; this allows you to supplement RNA's detection capability.

TIP! For improved performance, deactivate any add-on detectors that detect service or application traffic you are not interested in.

To activate or deactivate an RNA detector:

Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA 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 RNA Detectors on page 1394.

Keysight, EX-1004
Page 1398

3. You have the following options:
   - To **activate** a detector, so that RNA will use it when analyzing network traffic, click the activate icon ( ) next to the detector.
   - To **deactivate** a detector so that RNA will not use it when analyzing network traffic, click the deactivate icon ( ) next to the detector.

TIP! To activate or deactivate multiple detectors at once, select the check boxes next to the appropriate detectors, then click either **Activate** or **Deactivate** at the top of the page. You can also select all detectors in the current filtered view by selecting the check box at the top of the page.

Modifying Service Detectors

Platform: DC

Component: RNA

Use the following procedure to modify user-defined service detectors. Note that while you cannot modify any of the Sourcefire-provided detectors, you can view the properties of a well-known port service detector and save it as a new, user-defined service detector that you can modify and activate.

To modify a service detector:

Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA 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 RNA Detectors on page 1394.

3. You have the following options:
   - To modify a user-defined detector, click **Edit** next to the detector you want to modify.
   - To modify a copy of a Sourcefire-provided well-known port service detector, click **View** next to the detector you want to copy and modify.

   The Edit Service Detector page appears.

4. Make changes to the detector.

   See Creating a User-Defined Service Detector on page 1387 for information on the various configurations you can change.

Keysight, EX-1004
Page 1399

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.
   - If you are modifying a Sourcefire-provided well-known port detector, click **Save as New** to save the detector as a new, inactive user-defined detector.

IMPORTANT! RNA only uses active detectors to analyze service traffic. For more information, see Activating and Deactivating RNA Detectors on page 1398.

Deleting RNA Detectors

Platform: DC

Component: RNA

Use the following procedure to delete an RNA 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 RNA detectors, though you can deactivate many of them.

To delete an RNA detector:

Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA 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 RNA Detectors on page 1394.

TIP! To delete multiple detectors at once, select the check boxes next to the appropriate detectors, then click **Delete**. You can also select all detectors in the current filtered view by selecting the check box at the top of the page.

3. Click **OK** to confirm that you want to delete the detector.

   The detector is deleted.

Keysight, EX-1004
Page 1400

Viewing Internal RNA Detectors

Platform: DC

Component: RNA

The Sourcefire 3D System is delivered with many always-on internal RNA detectors that identify traffic for commonly used services and client applications. For your convenience, the Sourcefire 3D System provides a list of these internal detectors in the web interface.

Note that you can also see a list of the service and application detectors, including well-known port and add-on detectors, in Detected Operating Systems and Services on page 2019. Keep in mind, however, that the detector list can change depending on the version of the Sourcefire 3D System and the VDB you have installed, as well as on any individual add-on detectors you may have imported. 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 RNA detectors.

To view the internal detectors:

Access: P&R Admin/Admin

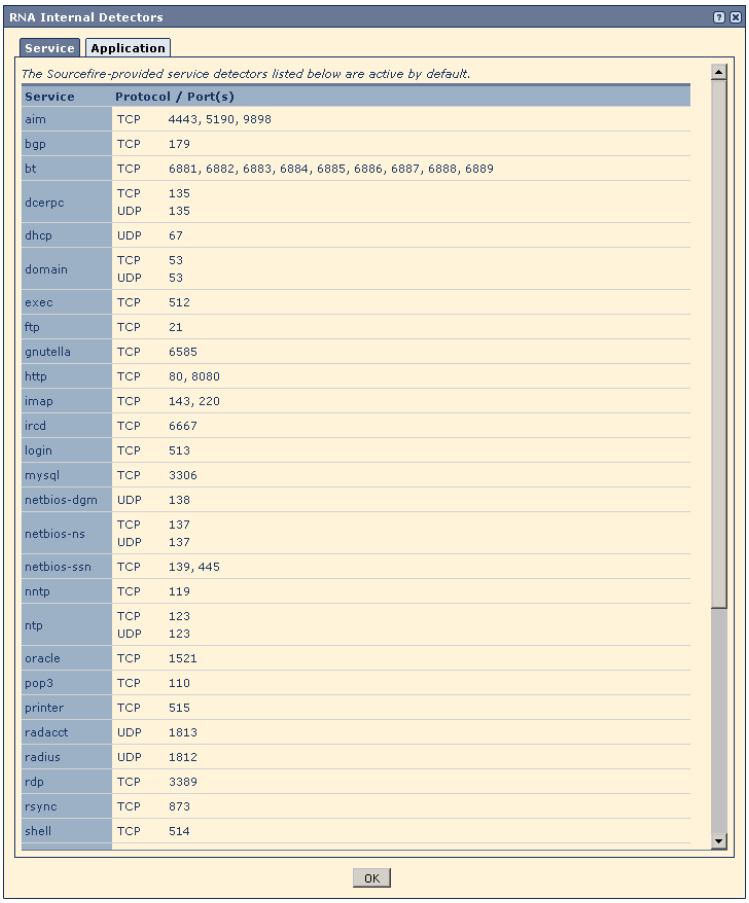
1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA Detectors page appears.

Keysight, EX-1004
Page 1401

2. Click **View RNA Internal Detectors**.

   The RNA Internal Detectors pop-up window appears, displaying the internal service detectors and their corresponding ports and protocols.



TIP! Click the **Application** tab to view the internal client application detectors and their corresponding types.

3. To close the pop-up window, click **OK**.

Working with RNA Detector Groups

Platform: DC

Component: RNA

RNA detector groups help you organize RNA detectors. The Sourcefire-provided well-known port service detectors are grouped by default according to the

Keysight, EX-1004
Page 1402

protocol they use (either TCP or UDP). You can also group user-defined and add-on detectors.

You can add an RNA detector to an existing group when you create the detector, modify an existing detector to add it to a different group, or specify which detectors belong to a group by modifying the group.

For more information, see:

- Creating RNA Detector Groups on page 1403
- Modifying RNA Detector Groups on page 1404
- Deleting RNA Detector Groups on page 1405

Creating RNA Detector Groups

Platform: DC

Component: RNA

Create RNA detector groups to help you organize RNA detectors. Note that a detector can only belong to one group at a time.

To create an RNA detector group:

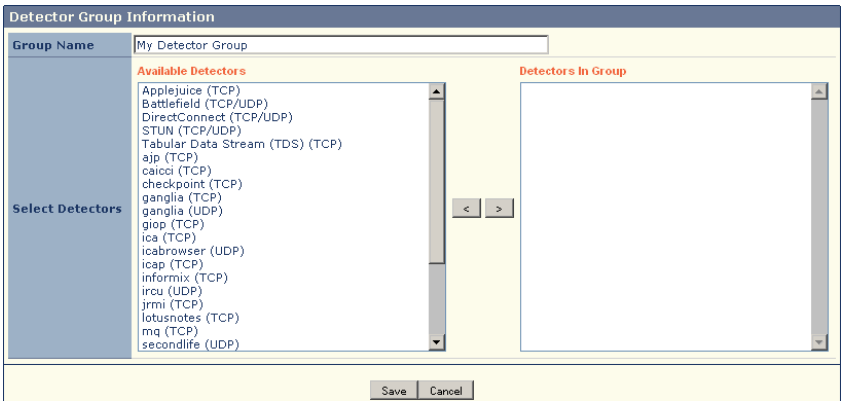
Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA Detectors page appears.

2. Click **Create Group**.

   The Create RNA Detector Group page appears.



3. In the **Group Name** field, type a name for the group.

Keysight, EX-1004
Page 1403

4. Optionally, add detectors to the group by selecting them from the **Available Detectors** list, then clicking the right arrow (**>**) button.

   TIP! Use Ctrl or Shift while clicking to select multiple detectors. You can also click and drag to select multiple adjacent detectors.

5. Click **Save**.

   The group is created.

Modifying RNA Detector Groups

Platform: DC

Component: RNA

Modify an RNA detector group to specify which detectors belong to it. Note that a detector can only belong to one group at a time. If you want to move a grouped service detector to another group, you must either:

- edit the detector as described in Modifying Service Detectors on page 1399
- use the following procedure to remove the detector from the original group, then again to add it to the new group

To modify an RNA detector group:

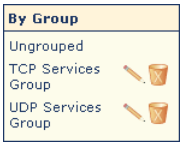
Access: P&R Admin/Admin

1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA Detectors page appears.

2. On the navigation panel on the left of the page, click **By Group**.

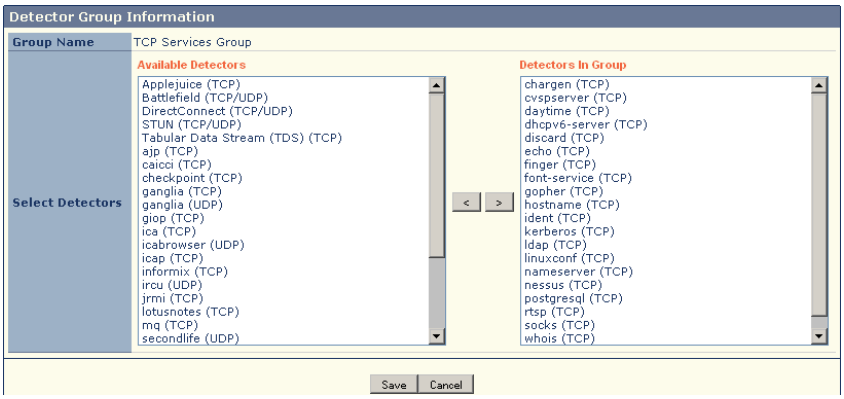
   A list of the existing RNA detector groups appears.



Keysight, EX-1004
Page 1404

3. Next to the group you want to edit, click the Edit icon ( ).

   The Edit RNA Detector Group page appears.



   TIP! You can also edit an RNA detector group by clicking its name on the RNA Detectors page.

4. You have the following options:
   - To move a detector out of the group, select it from the **Detectors In Group** list, then click the left arrow (**<**) button.
   - To move a detector into the group, select it from the **Available Detectors** list, then click the right arrow (**>**) button.

   TIP! Use Ctrl or Shift while clicking to select multiple detectors. You can also click and drag to select multiple adjacent detectors.

5. Click **Save**.

   The group is saved.

Deleting RNA Detector Groups

Platform: DC

Component: RNA

When you delete an RNA detector group, detectors that were in the group are **not** deleted. Rather, they merely become ungrouped.

To delete an RNA detector group:

Access: P&R Admin/Admin

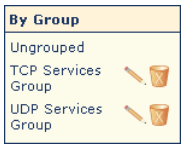
1. Select **Policy & Response** > **RNA** > **RNA Detectors**.

   The RNA Detectors page appears.

Keysight, EX-1004
Page 1405

2. On the navigation panel on the left of the page, click **By Group**.

   A list of the existing RNA detector groups appears.



3. Next to the group you want to delete, click the Delete icon ( ).
4. Click **OK** to confirm that you want to delete the group.

   The group is deleted and detectors that were in that group become ungrouped.

Importing Host Input Data

Platform: DC

Component: RNA

If your organization has the capability to write scripts or create command line import files to import network map data from third-party applications, you can import data to augment the information in the RNA network map. You can also use the host input feature by modifying operating system or service identities or deleting services, protocols, host attributes, or client applications using the web interface.

RNA may reconcile data from multiple sources to determine the current identity of an operating system or service. For more information on how RNA does this, see Understanding Current Identities on page 1360.

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 service definitions in the RNA database. For more information, see the following sections:

- Managing Third-Party Product Mappings on page 1408
- Mapping Third-Party Vulnerabilities on page 1412
- Managing Custom Product Mappings on page 1414

Keysight, EX-1004
Page 1406

Enabling Use of Third-Party Data

Platform: DC

Component: RNA

You can import network map data from third-party systems on your network. However, to enable features where IPS and RNA data are used together, such as RNA recommended rules 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 application that has more specific data on the operating systems or services running on hosts on your network, you can import that data using the host input feature. However, because third-party applications may name the products differently, you must map the third-party application vendor, product, and versions to the corresponding RNA product definition. After you map the products, you must enable RNA vulnerability mappings for impact assessment in the system policy to allow impact correlation. For versionless or vendorless services, you need to map vulnerabilities for the service types in the system policy. For more information, see Mapping Third-Party Products on page 1408 and Configuring RNA Settings on page 1756.
- If you import patch information from a third-party application 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 RNA 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 1410.
- Requires: DC + IPS + RNA If you import vulnerabilities from a third-party application and you want to use them for impact correlation, you must map the third-party vulnerability identification string to vulnerabilities in the RNA database. 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 1412 and Configuring RNA Settings on page 1756. To cause services without vendor or version information to map to vulnerabilities, an administrative user must also map vulnerabilities for the service types in the system policy. For more information, see Mapping Vulnerabilities for Services on page 1773.
- If you import service data and you want to use that data for impact correlation, you must map the vendor string for each service to the corresponding RNA service definition. For more information, see Managing Custom Product Mappings on page 1414.

Keysight, EX-1004
Page 1407

Managing Third-Party Product Mappings

Platform: DC

Component: RNA

When you add data from third-party applications to the network map through the user input feature, you must map the vendor, product, and version names used by the third-party application to the RNA product definitions. Mapping the products to RNA definitions assigns vulnerabilities based on those definitions.

Similarly, if you are importing patch information from a third-party application, 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 RNA database.

For more information, see the following sections:

- Mapping Third-Party Products on page 1408
- Mapping Third-Party Product Fixes on page 1410

Mapping Third-Party Products

Platform: DC

Component: RNA

If you import data from a third-party application, you must map the RNA product to the third-party name to assign vulnerabilities and perform impact correlation using that data. Mapping the product associates RNA vulnerability information with the third-party product name, which allows the Sourcefire 3D 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 RNA vulnerabilities during the import.

As an example, if you import data from a third-party application that lists Apache Tomcat as a service 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 a service listing for Apache Tomcat.

Note that for versionless or vendorless services, you must map vulnerabilities for the service types in the system policy. For more information, see Mapping Vulnerabilities for Services on page 1773.

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 1997.

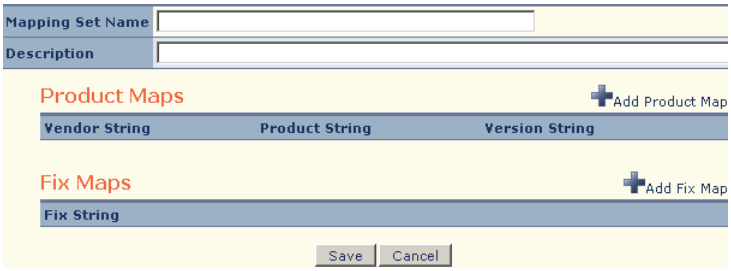
Keysight, EX-1004
Page 1408

To map a third-party product to an RNA product definition:

Access: Admin

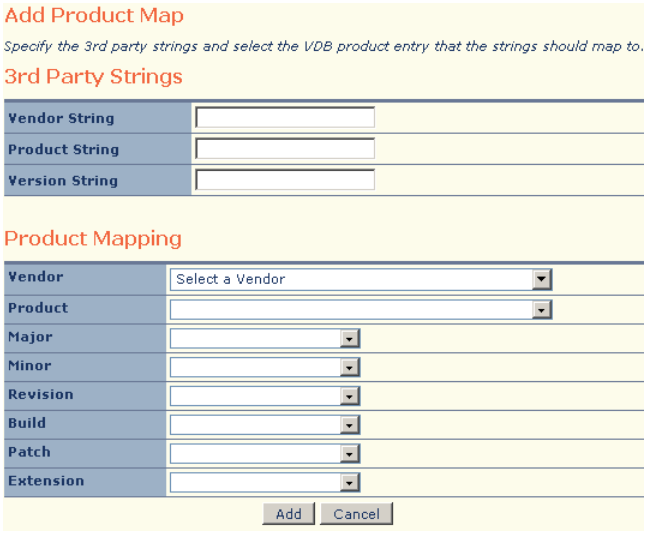
1. Select **Policy & Response** > **RNA** > **Network Map** > **User Third-Party Mappings**.
   The User Third-Party Mappings page appears.
2. You have two choices:
   - To edit an existing map set, click **Edit** next to the map set.
   - To create a new map set, click **Create Product Map Set**.

   The **Edit Third-Party Product Mappings** page appears.



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 pop-up window appears.



Keysight, EX-1004
Page 1409

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 RNA 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 **Add**.

Mapping Third-Party Product Fixes

Platform: DC

Component: RNA

If you map a fix name to a particular set of fixes in the RNA 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, RNA marks all vulnerabilities addressed by the fix as invalid for that host.

To map third-party fixes to RNA fix definitions:

Access: Admin

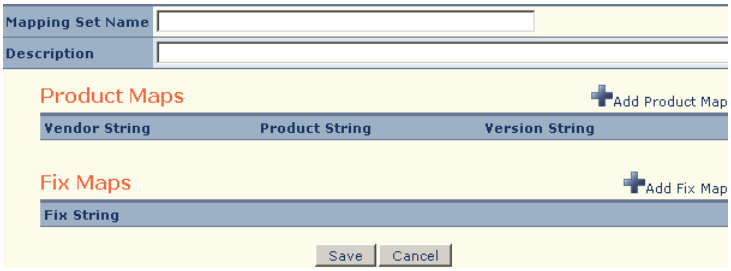
1. Select **Policy & Response** > **RNA** > **Network Map** > **User Third-Party Mappings**.

   The User Third-Party Mappings page appears.

Keysight, EX-1004
Page 1410

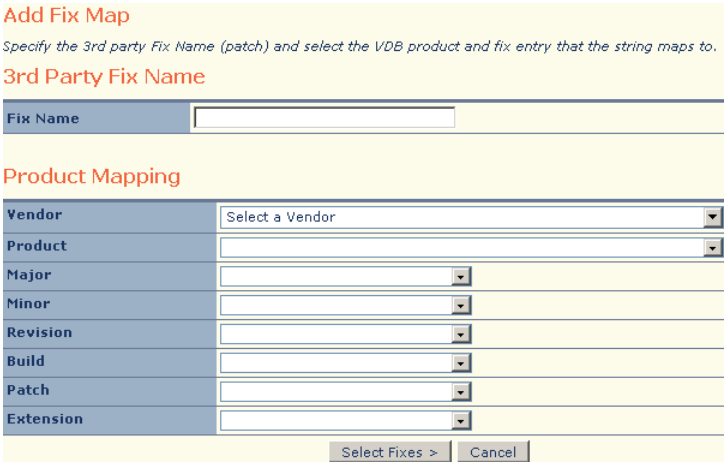
2. You have two choices:
   - To edit an existing map set, click **Edit** next to the map set.
   - To create a new map set, click **Create Product Map Set**.

   The **Edit Third-Party Product Mappings** page appears.



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 pop-up window appears.



6. Type the name of the fix you want to map in the **Fix Name** field.
7. In the Product Mappings section, select the operating system, product, and versions you want to use for RNA fix mapping from the following lists (if applicable):
   - **Vendor**
   - **Product**
   - **Major Version**

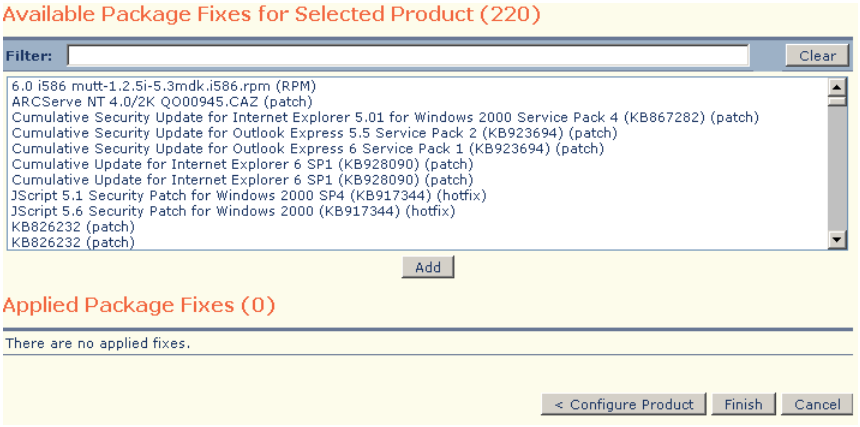
Keysight, EX-1004
Page 1411

- **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 **Select Fixes**.

   The Available Package Fixes for Select Product page appears.



9. Optionally, enter a regular expression in the **Filter** field to filter the list of fixes you can select from. To clear a filter, click **Clear**.
10. Select the fixes you want to automatically apply to the configured product and click **Add**.

> **TIP!** Use Ctrl or Shift while clicking to select multiple fixes. You can click and drag to select multiple adjacent fixes.

   The fixes are added to the Applied Package Fixes list.

11. Click **Finish** to save the fix map.

Mapping Third-Party Vulnerabilities

Platform: DC

Component: RNA

To add vulnerability information from a third-party application to the VDB, you must map the third-party identification string for each imported vulnerability to

Keysight, EX-1004
Page 1412

any existing RNA, 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 in the system policy to allow correlation to occur. For more information, see Configuring RNA Settings on page 1756. For versionless or vendorless services, you must also map vulnerabilities for the service types in the system policy. For more information, see Mapping Vulnerabilities for Services on page 1773.

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 1997.

To map a third-party vulnerability to an existing vulnerability:

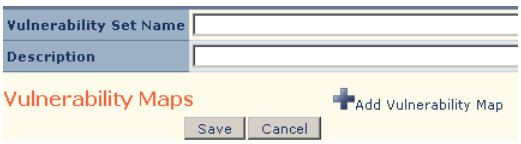
Access: Admin

1. Select **Policy & Response** > **RNA** > **Network Map** > **User Third-Party Mappings**.

   The User Third-Party Mappings page appears.

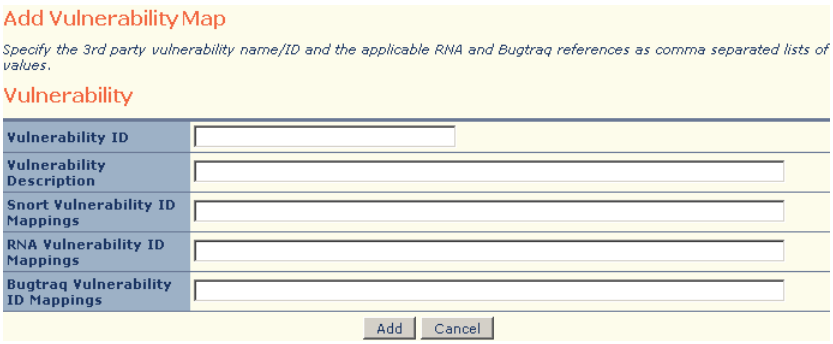
2. You have two choices:
   - To edit an existing vulnerability set, click **Edit** next to the vulnerability set.
   - To create a new vulnerability set, click **Create Vulnerability Map Set**.

   The **Edit Third-Party Vulnerability Mappings** page appears.



3. Click **Add Vulnerability Map**.

   The Add Vulnerability Map pop-up window appears.



Keysight, EX-1004
Page 1413

4. Type the third-party application 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 RNA vulnerability ID in the **RNA 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

Platform: DC

Component: RNA

You can use product mappings to ensure that services or applications input by a third-party application are associated with the appropriate RNA service definitions. Once you define and activate the product mapping, all services or applications 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 services in the network map with a particular vendor string instead of explicitly setting the vendor, product, and version for the service.

For more information, see the following:

- Creating Custom Product Mappings on page 1414
- Editing Custom Product Mapping Lists on page 1416
- Managing Custom Product Mapping Activation State on page 1417

Creating Custom Product Mappings

Platform: DC

Component: RNA

If RNA cannot map a service in the network map to a vendor and product in the VDB, you can manually create the mapping for RNA to use when identifying services. When you activate a custom product mapping, RNA maps

Keysight, EX-1004
Page 1414

vulnerabilities for the selected vendor and product to all services in the network map where that vendor string occurs.

IMPORTANT! Custom product mappings apply to all occurrences of a service, regardless of the source of the service data (such as RNA, Nmap, or the host input feature). 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 1412.

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 an RNA vendor to map to, RNA 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, RNA updates all services 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 service.

If, for example, your company modifies the banner for your Apache Tomcat web services 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 service labelled `Internal Web Server` occurs have the vulnerabilities for Apache Tomcat in the RNA 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:

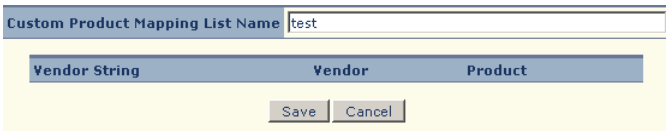
Access: Admin

1. Select **Policy & Response** > **RNA** > **Network Map** > **Custom Product Mappings**.

   The Custom Product Mappings page appears.

2. Click **Create Custom Product Mapping List**.

   The Edit Custom Product Mappings List page appears.

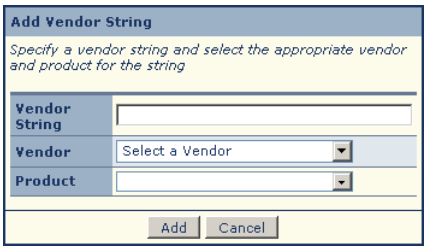


3. Type a name in the **Custom Product Mapping List Name** field.

Keysight, EX-1004
Page 1415

4. Click **Add Vendor String**.

   The Add Vendor String pop-up window appears.



5. In the **Vendor String** field, type the vendor string that identifies the services that should map to the selected RNA vendor and product values.
6. Select the RNA vendor you want to map to from the **Vendor** drop-down list.
7. Select the RNA 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

Platform: DC

Component: RNA

You can modify existing custom product mapping lists by adding or removing vendor strings or changing the list name.

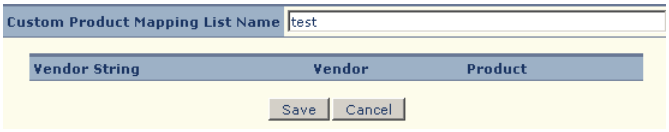
To edit a custom product mapping:

Access: Admin

1. Select **Policy & Response** > **RNA** > **Network Map** > **Custom Product Mappings**.

   The Custom Product Mappings page appears.
2. Click **Edit** next to the product mapping list to edit.

   The Edit Custom Product Mappings List page appears.



3. Make changes to the list as needed. For more information, see Creating Custom Product Mappings on page 1414.

Keysight, EX-1004
Page 1416

4. When you finish, click **Save**.

   The Custom Product Mappings page appears, with the list you updated.

Managing Custom Product Mapping Activation State

Platform: DC

Component: RNA

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 services on hosts in the network map with the specified vendor string, whether detected by RNA or imported through the host input feature.

To activate or deactivate a custom product mapping list:

Access: Admin

1. Select **Policy & Response** > **RNA** > **Network Map** > **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**.

Keysight, EX-1004
Page 1417

Chapter 31
Configuring Active Scanning

RNA 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 service running on an open port but the service has not received or sent traffic during the time that RNA has been active, RNA does not add information about that service to the network map. If you directly scan that host using an active scanner, however, you can detect the presence of the service.

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™ 5.21, an open source active scanner for network exploration and security auditing that can be used to detect operating systems and services running on a host. With an Nmap scan, you can check for detailed information about the operating system and services running on the host and refine RNA's vulnerability reporting based on those results.

IMPORTANT! Some scanning options (such as portscans) can 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 1419
- Managing Nmap Scanning on page 1437
- Managing Scan Targets on page 1442
- Working with Active Scan Results on page 1443

Keysight, EX-1004
Page 1418

Understanding Nmap Scans

Platform: DC

Component: RNA

Nmap allows you to actively scan ports on hosts on your network to determine operating system and service data for the hosts. When you scan a host using Nmap, services on previously undetected open ports are added to the Services list in the host profile for that host. The host profile lists any services 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.

Although you can scan hosts using Nmap and view the scan results in a results file with a Defense Center and a 3D Sensor in your Sourcefire 3D System, to access the full functionality of the Sourcefire 3D System Nmap integration, your Defense Center should have an RNA host license installed. With an RNA license, the results of an Nmap scan enhance your RNA network map and let you fine-tune the accuracy of the vulnerabilities mapped to scanned hosts.

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 RNA recognizes a service identified in an Nmap scan and has a corresponding service definition, RNA maps vulnerabilities for that service to the host. RNA maps the names Nmap uses for services to the corresponding RNA service definitions, and then uses the vulnerabilities mapped to each service in RNA. Similarly, RNA maps Nmap operating system names to RNA operating system definitions. When Nmap detects an operating system for a host, RNA assigns vulnerabilities from the corresponding RNA operating system definition to the host.

Note that the host must exist in the network map before Nmap can append the results to the host profile.

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 1419
- Creating an Nmap Scanning Strategy on page 1422
- Sample Nmap Scanning Profiles on page 1424

Understanding Nmap Remediations

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 compliance policy, run on demand, or scheduled to run at a specific time. In order for the results of an

Keysight, EX-1004
Page 1419

Nmap scan to appear in the network map, the scanned host must already exist in the network map. Note that RNA, NetFlow, and the host input feature can add hosts to the network map.

Note that Nmap-supplied service and operating system data remain static until you run another Nmap scan. If you plan to scan a host for operating system and service data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and service data up-to-date. For more information, see Automating Nmap Scans on page 1722. 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
Scan Which Address(es) From Event?	When you use an Nmap scan as a response to a compliance rule, select an option to control which address in the event is scanned, that of the source host, the destination host, or both.
Scan Types	Select how Nmap scans ports: • 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 remote 3D Sensor 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.
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.

Keysight, EX-1004
Page 1420

Nmap Remediation Options (Continued)

Option	Description
Use Port From Event	If you plan to use the remediation as a response in a compliance policy, enable to cause the remediation to scan only the port specified in the event that triggers the compliance response. **TIP!** You can also control whether Nmap collects information about operating system and service information. Enable the **Use Port From Event** option to scan the port associated with the new service.
Scan from reporting detection engine	Enable to scan a host from the appliance where the detection engine that reported the host resides.
Fast Port Scan	Enable to scan only the TCP ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the sensor that does the scanning, ignoring other port settings.
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.
Probe open ports for vendor and version information	Enable to detect service vendor and version information. If you probe open ports for service vendor and version information, Nmap obtains service data that it uses to identify services. It then replaces the RNA service data for that service.
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.
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 that 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 1360.
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**.

Keysight, EX-1004
Page 1421

Nmap Remediation Options (Continued)

Option	Description
Host Discovery Method	Select to perform host discovery for all hosts in the target range, over the ports listed in the **Host Discovery Port Lis**t, or if no ports are listed, over the default ports for that host discovery method. 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. Select the method to be used when Nmap tests to see if a host is present and available: • 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.
Host Discovery Port List	Specify a customized list of ports, separated by commas, that you want to scan when doing host discovery.
Default NSE Scripts	Enable to run the default set of Nmap scripts for host discovery and service and operating system and vulnerability detection. See http://nmap.org/nsedoc/categories/default.html for the list of default scripts.
Timing Template	Select the timing of the scan process; the higher the number you select, the faster and less comprehensive the scan.

Creating an Nmap Scanning Strategy

Platform: DC

Component: RNA

While active scanning can obtain valuable information, overuse of a tool such as Nmap can 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 1423
- Selecting Appropriate Ports to Scan on page 1424
- Setting Host Discovery Options on page 1424

Keysight, EX-1004
Page 1422

Selecting Appropriate Scan Targets

Platform: DC

Component: RNA

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:

- 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 59.
- 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 RNA is unable to identify, hosts with unidentified services, 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 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 any Nmap-supplied operating system and service data up to date. For more information, see Automating Nmap Scans on page 1722. 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.

Keysight, EX-1004
Page 1423

Selecting Appropriate Ports to Scan

Platform: DC

Component: RNA

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 TCP ports 1-1024. If you plan to use the remediation as a response in a compliance policy, you can cause the remediation to scan only the port specified in the event that triggers the compliance 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

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 service 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

Platform: DC

Component: RNA

The following scenarios provide examples of how Nmap might be used on your network:

- Example: Resolving Unknown Operating Systems on page 1425
- Example: Responding to New Hosts on page 1426

Keysight, EX-1004
Page 1424

Example: Resolving Unknown Operating Systems

Platform: DC

Component: RNA

If RNA 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 service information deactivates RNA monitoring of that data for scanned hosts. If you use Nmap to discover host and service operating system for hosts RNA 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 RNA 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 RNA to continue to monitor the host operating system and update it as needed.

To discover operating systems with Nmap:

Access: P&R Admin/Admin

1. Configure a scan instance for an Nmap module.

   For more information, see Creating an Nmap Scan Instance on page 1428.

2. Create an Nmap remediation using the following settings:
   - Enable **Use Port From Event** to scan the port associated with the new service.
   - Enable **Detect Operating System** to detect operating system information for the host.
   - Enable **Probe open ports for vendor and version information** to detect service 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 1431.

3. Create a compliance rule that triggers when RNA detects a host with an unknown operating system.

   The rule should trigger when **an RNA 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 compliance rules, see Creating Rules for Compliance Policies on page 1214.

4. Create a compliance policy that contains the compliance rule.

   For more information on creating compliance policies, see Creating Compliance Policies on page 1262.

Keysight, EX-1004
Page 1425

5. In the compliance policy, add the Nmap remediation you created in step 2 as a response to the rule you created in step 3.
6. Activate the compliance policy.
7. Purge the hosts on your network map to force RNA to restart and rebuild the network map.
8. After a day or two, search for events generated by the compliance 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 RNA does not recognize.

   For more information on analyzing Nmap results, see Analyzing Nmap Results on page 1451.
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 1364.

Example: Responding to New Hosts

Platform: DC

Component: RNA

When RNA 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 compliance 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 service 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 services 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 service information deactivates RNA monitoring of that data for scanned hosts.

Keysight, EX-1004
Page 1426

To scan in response to the appearance of a new host:

Access: P&R Admin/Admin

1. Configure a scan instance for an Nmap module.

   For more information, see Creating an Nmap Scan Instance on page 1428.

2. Create an Nmap remediation using the following settings:
   - Enable **Use Port From Event** to scan the port associated with the new service.
   - Enable **Detect Operating System** to detect operating system information for the host.
   - Enable **Probe open ports for vendor and version information** to detect service 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 1431.

3. Create a compliance rule that triggers when RNA detects a new host on a specific subnet.

   The rule should trigger when **an RNA event occurs** and **a new host is detected**.

   For information on creating compliance rules, see Creating Rules for Compliance Policies on page 1214.

4. Create a compliance policy that contains the compliance rule.

   For more information on creating compliance policies, see Creating Compliance Policies on page 1262.

5. In the compliance policy, add the Nmap remediation you created in step as a response to the rule you created in step 3.

6. Activate the compliance 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.

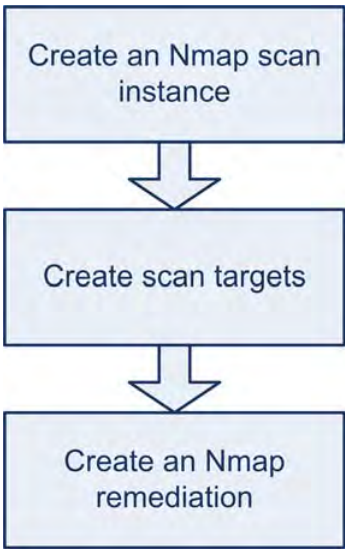
Keysight, EX-1004
Page 1427

Setting up Nmap Scans

Platform: DC

Component: RNA

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 1428
- Creating an Nmap Scan Target on page 1430
- Creating an Nmap Remediation on page 1431

Creating an Nmap Scan Instance

Platform: DC

Component: RNA

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 sensors 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 sensor. 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.

Keysight, EX-1004
Page 1428

Note that you cannot add a scan instance with the same name as any existing scan instance.

To create a scan instance:

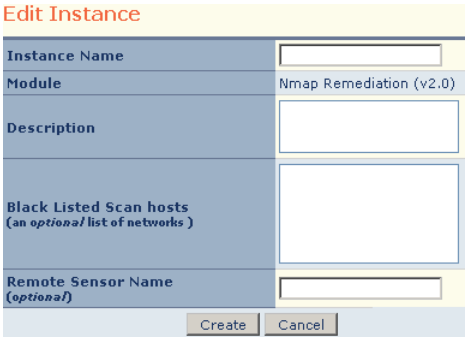
Access: P&R Admin/Admin

1. Select **Operations > Tools > Scanners**.

   The Scanners page appears.

2. Click **Add Nmap Instance**.

   The Instance Detail 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 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:

   - an exact IP address (for example, `192.168.1.101`)
   - 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 59.

   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 sensor instead of the Defense Center, specify the IP address or name of the sensor as it appears in the Information page for the sensor in the Defense Center web interface, in the **Remote Sensor Name** field.

7. Click **Create**.

   The scan instance is created.

Keysight, EX-1004
Page 1429

Creating an Nmap Scan Target

Platform: DC

Component: RNA

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. 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.

Note that 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 any Nmap-supplied operating system and service data up to date. For more information, see Automating Nmap Scans on page 1722. 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:

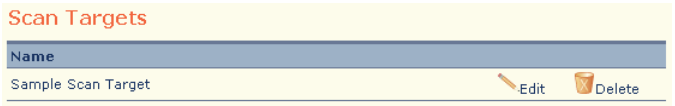
Access: P&R Admin/Admin

1. Select **Operations > Tools > Scanners**.

   The Scanners page appears.

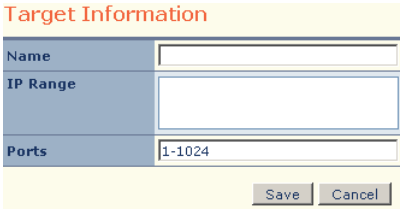
2. On the toolbar, click **Targets**.

   The Scan Target List page appears.



3. Click **Create Scan Target**.

   The Scan Target page appears.



4. In the **Name** field, type the name you want to use for this scan target.

Keysight, EX-1004
Page 1430

5. In the **IP Range** text box, specify the host or hosts you want to scan, using the following syntax:
   - an exact IP address (for example, `192.168.1.101`) or comma-separated list of IP addresses
   - 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 59.
   - 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 6 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)

   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:
   - 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

Platform: DC

Component: RNA

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 compliance 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

Keysight, EX-1004
Page 1431

the network map. Note that RNA, NetFlow, and the host input feature can add hosts to the network map.

For more information on the specific settings in an Nmap remediation, see Understanding Nmap Remediations on page 1419.

Note that Nmap-supplied service and operating system data remains static until you run another Nmap scan. If you plan to scan a host for operating system and service data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and service data up-to-date. For more information, see Automating Nmap Scans on page 1722. 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: P&R Admin/Admin

1. Select **Operations > Tools > Scanners**.

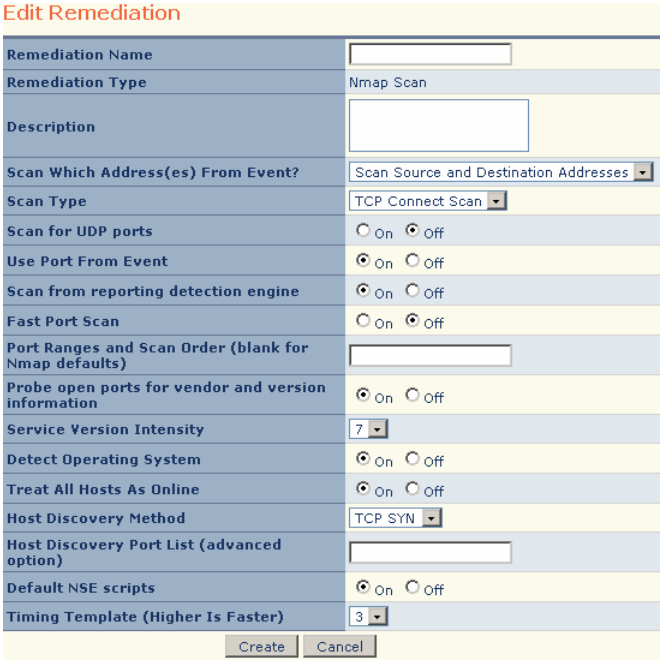
   The Scanners page appears.



Keysight, EX-1004
Page 1432

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.

Keysight, EX-1004
Page 1433

5. If you plan to use this remediation in response to a compliance rule that triggers on an intrusion event, a flow event, or a RUA 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 compliance rule that triggers on an RNA 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 compliance 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.

Keysight, EX-1004
Page 1434

8. If you plan to use this remediation in response to compliance policy violations, configure the **Use Port From Event** option:
   - Select **On** to scan the port in the compliance event, rather than the ports you specify in step 11.

     If you scan the port in the compliance 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 compliance 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 sensor 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.

    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 compliance policy violation, as described in step 8.
12. To probe open ports for service vendor and version information, configure **Probe open ports for vendor and version information:**
    - Select **On** to scan open ports on the host for service information to identify service vendors and versions.
    - Select **Off** to continue using RNA service 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.

Keysight, EX-1004
Page 1435

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 RNA 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 service, 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.

Keysight, EX-1004
Page 1436

20. Click **Save**, then click **Done**.

    The remediation is created.

Managing Nmap Scanning

Platform: DC

Component: RNA

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 1437
- Managing Nmap Remediations on page 1439
- Running an On-Demand Nmap Scan on page 1440

Managing Nmap Scan Instances

Platform: DC

Component: RNA

You can edit or delete Nmap scan instances. For more information, see the following sections:

- Editing an Nmap Scan Instance on page 1437
- Deleting an Nmap Scan Instance on page 1438

Editing an Nmap Scan Instance

Platform: DC

Component: RNA

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: P&R Admin/Admin

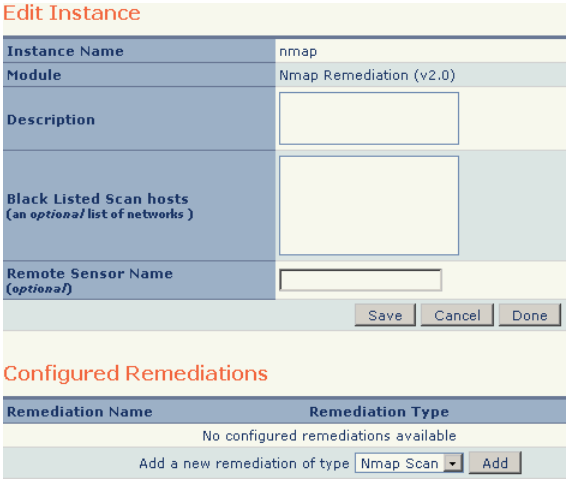
1. Select **Operations > Tools > Scanners**.

   The Scanners page appears.

Keysight, EX-1004
Page 1437

2. Click **View** next to the instance you want to edit.

   The Instance Detail page appears.



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 1439.

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 1440.

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 1439.

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

Platform: DC

Component: RNA

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.

Keysight, EX-1004
Page 1438

To delete a scan instance:

Access: P&R Admin/Admin

1. Click **Operations > Tools > Scanners**.

   The Scanners page appears.

2. Click **Delete** next to the scan instance you want to delete.

   The instance is deleted.

Managing Nmap Remediations

Platform: DC

Component: RNA

You can edit or delete Nmap remediations. For more information, see the following sections:

- Editing an Nmap Remediation on page 1439
- Deleting an Nmap Remediation on page 1440

Editing an Nmap Remediation

Platform: DC

Component: RNA

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: P&R Admin/Admin

1. Select **Operations > Tools > 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 1431.

4. Click **Save**, then click **Done**.

   The remediation is modified.

Keysight, EX-1004
Page 1439

Deleting an Nmap Remediation

Platform: DC

Component: RNA

Delete an Nmap remediation if you no longer need it.

To delete an Nmap remediation:

Access: P&R Admin/Admin

1. Select **Operations > Tools > 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

Platform: DC

Component: RNA

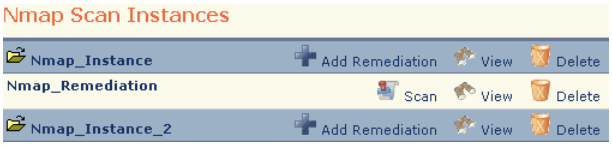
You can launch on-demand Nmap scans whenever needed. You can specify the target for an on-demand scan by entering the IP addresses and ports you want to scan or by selecting an existing scan target.

Note that Nmap-supplied 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 any Nmap-supplied operating system and service data up to date. For more information, see Automating Nmap Scans on page 1722. 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: P&R Admin/Admin

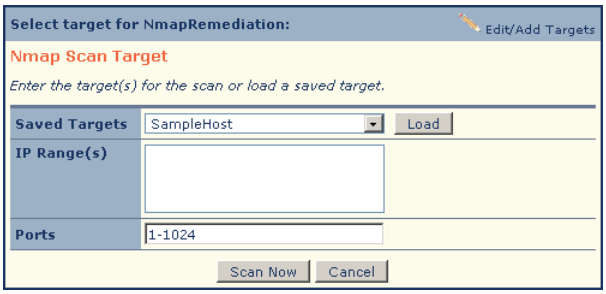
1. Select **Operations > Tools > Scanners**.
   The Scanners page appears.



Keysight, EX-1004
Page 1440

2. Next to the Nmap remediation you want to use to perform the scan, click **Scan**.

   The Nmap Scan Target dialog box appears.



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 1430.

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.

   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 59.

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 1558.

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.

Keysight, EX-1004
Page 1441

Managing Scan Targets

Platform: DC

Component: RNA

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.

Once 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 1430
- Editing a Scan Target on page 1442
- Deleting a Scan Target on page 1443

Editing a Scan Target

Platform: DC

Component: RNA

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 compliance policy violation that launched the remediation.

To edit an existing scan target:

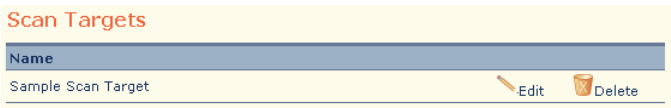
Access: P&R Admin/Admin

1. Select **Operations > Tools > Scanners**.

   The Scanners page appears.

2. On the toolbar, click **Targets**.

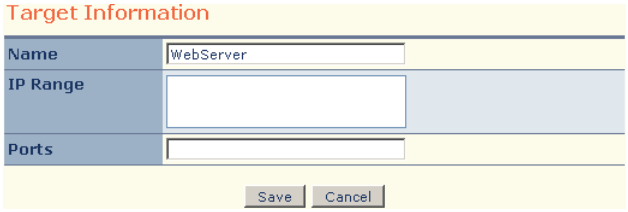
   The Scan Target List page appears.



Keysight, EX-1004
Page 1442

3. Click **Edit** next to the scan target you want to edit.

   The Scan Target page appears.



4. Make modifications as necessary and click **Save**.

   The scan target is updated.

Deleting a Scan Target

Platform: DC

Component: RNA

Delete a scan target if you no longer want to scan the hosts listed in it.

To delete a scan target:

Access: P&R Admin/Admin

1. Select **Operations > Tools > Scanners**.

   The Scanners page appears.

2. On the toolbar, click **Targets**.

   The Scan Target List page appears.



3. Next to the scan target you want to delete, click **Delete**.

   The scan target is deleted.

Working with Active Scan Results

Platform: DC

Component: RNA

For information on how to monitor Nmap scans in progress, import results from scans previously performed through the Sourcefire 3D System or results

Keysight, EX-1004
Page 1443

preformed outside the Sourcefire 3D System, and view and analyze scan results, see the following sections:

- Monitoring Scans on page 1444
- Importing Results on page 1445
- Viewing Scan Results on page 1445
- Understanding the Scan Results Table on page 1448
- Searching for Scan Results on page 1449
- Analyzing Nmap Results on page 1451

Monitoring Scans

Platform: DC

Component: RNA

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, once 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: P&R Admin/Admin

1. Select **Operations > Tools > Scan Results**.

   The first page of the default scan results workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of scan results, from the **Workflows** menu on the toolbar, select **Scan Results**.

Keysight, EX-1004
Page 1444

2. 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 Results

Platform: DC

Component: RNA

You can import XML results files created by and 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 1444.

To import results:

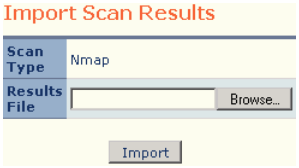
Access: P&R Admin/Admin

1. Select **Operations > Tools > Scanners**.

   The Scan Instances page appears.

2. On the toolbar, click **Import Results**.

   The Import Results page appears.



3. Click **Browse** to navigate to the results file.
4. After you return to the Import Results page, click **Import** to import the results.

   The results file is imported.

Viewing Scan Results

Platform: DC

Component: RNA

You can view a table of scan results, and then manipulate the event view depending on the information you are looking for.

Keysight, EX-1004
Page 1445

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 1618.

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 1448.
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 1597.
sort scan results	click the column title. Click the column title again to reverse the sort order.
constrain the columns that appear	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, use the Expand arrow () to expand the search constraints, then click the column name under Disabled Columns.

Keysight, EX-1004
Page 1446

Scan Results Table Functions (Continued)

To...	You can...
drill down to the next page in the workflow, constraining on a specific value	use one of the following methods: • 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
configure scan instances and remediations	Click **Scanners** in the toolbar. For more information, see Setting up Nmap Scans on page 1428.
navigate within and between workflow pages	find more information in Using Workflow Pages on page 1593.
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 1614.
search for scan results	click **Search**. For more information, see Searching for Scan Results on page 1449.

To view scan results:

Access: P&R Admin/Admin

▶ Select **Operations > Tools > Scan Results**.

The first page of the default scan results workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

Keysight, EX-1004
Page 1447

Understanding the Scan Results Table

Platform: DC

Component: RNA

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 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: • `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.

Keysight, EX-1004
Page 1448

Searching for Scan Results

Platform: DC

Component: RNA

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 1556 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 1556 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 1557 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: • 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 1550.

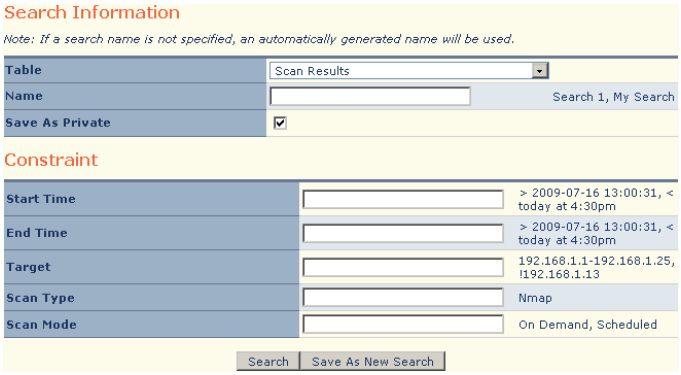
Keysight, EX-1004
Page 1449

To search for scan results:

Access: P&R Admin/Admin

1. Select **Analysis & Reporting > Searches > Scan Results**.

   The Scan Results search page appears.



   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.

   TIP! 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

4. 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.

Keysight, EX-1004
Page 1450

5. 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.

Analyzing Nmap Results

Platform: DC

Component: RNA

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 service information detected by Nmap in host profiles and in the network map. If a scan of a host produces service information for services on filtered or closed ports, or if a scan collects information that cannot be included in the operating system information or the services section, the host profile includes those results in an Nmap Scan Results section. For more information, see Viewing Host Profiles on page 948.

Keysight, EX-1004
Page 1451

Chapter 32
Using Sourcefire RUA

Sourcefire Real-time User Awareness, also called RUA, allows your organization to correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, RUA can help you to identify the source of policy breaches, attacks, or network vulnerabilities, as well as 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.

All RUA deployments require a Defense Center that has an RUA feature license installed. If your organization uses LDAP, you can use the user information on your LDAP servers to augment the Defense Center's database of user identity information with available metadata.

In addition, you have three choices in how you collect user login information:

- You can configure 3D Sensors (except 3D9800 and Crossbeam-based software sensors) to observe your organization's network traffic and communicate various types of user logins to the sensors' managing Defense Center.
- If your organization uses Active Directory LDAP servers, you can install Sourcefire RUA Agents on them. The agents communicate LDAP user logins to the Defense Center.
- You can use both 3D Sensors and RUA Agents.

Keysight, EX-1004
Page 1452

The following sections explain more about RUA, how to choose an implementation and set up RUA, and how to view, interpret, and use the data that RUA collects.

- Understanding RUA on page 1453
- Configuring RUA on page 1474
- Working with RUA Users on page 1489
- Working with RUA Events on page 1500

Understanding RUA

Platform: DC

Component: RUA

RUA, when paired with IPS or RNA (or both), allows you to correlate network activity with user identity information. By linking network behavior, traffic, and events directly to individual users, RUA can help you to identify the source of policy breaches, attacks, or network vulnerabilities. In other words, RUA 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 red (vulnerable) 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.

Depending on your RUA implementation, you can detect various types of user activity.

- Sourcefire RUA Agents monitor LDAP users as they log into the network or when they authenticate against Active Directory credentials for any other reason.
- 3D Sensors with RUA can detect not only LDAP user logins, but also POP3, IMAP, SMTP, and AIM logins.
- 3D Sensors with RNA can detect Oracle and Session Initiation Protocol (SIP) logins; SIP is used for communications sessions such as VoIP.

Note that regardless of the mechanism used to detect user logins, your Defense Center **must** have an RUA feature license to add user identity information in the database.

Keysight, EX-1004
Page 1453

For more information, see the following sections:

- How Does RUA Work? on page 1454
- How Do I Choose an RUA Implementation? on page 1463
- What Information Does RUA Provide? on page 1466
- How Can You Use RUA? on page 1468
- What are the Limitations of Sourcefire RUA? on page 1473

How Does RUA Work?

Platform: DC

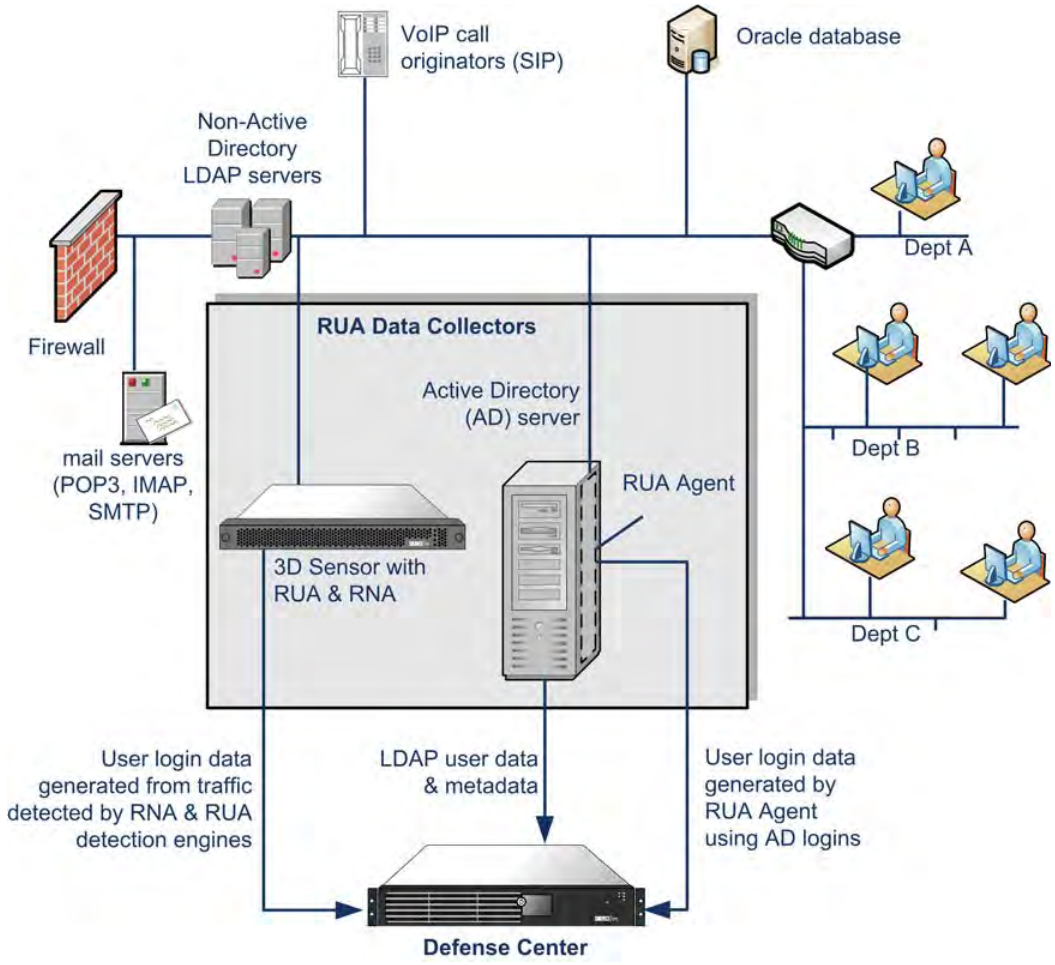
Component: RUA

An RUA deployment includes:

- a Defense Center that you can use to configure RUA and view and analyze the data that RUA collects
- 3D Sensors or Sourcefire RUA Agents (or both) that detect user logins
- optionally, LDAP servers that you can use to obtain information about the users detected by RUA

Keysight, EX-1004
Page 1454

The following diagram illustrates a typical RUA deployment.



The following sections describe the role that each component plays in your RUA deployment:

- Understanding LDAP Servers and RUA on page 1456
- Understanding Sourcefire RUA Agents on page 1458
- Understanding 3D Sensors and RUA on page 1459
- Understanding the Defense Center and RUA on page 1460

Keysight, EX-1004
Page 1455

Understanding LDAP Servers and RUA

Platform: DC

Component: RUA

If your organization uses LDAP, you can configure the Defense Center to connect to LDAP servers to obtain metadata for RUA-detected LDAP users. RUA supports connections to LDAP servers running Microsoft Active Directory on Windows Server 2003, Sun Directory Services 5.2 P4 on Windows 2003 and XP, or OpenLDAP on Linux. You must have TCP/IP access from the Defense Center to the LDAP servers.

To create a Defense Center-LDAP server connection, you must create and configure an *authentication object*, which contains your settings for connecting to and retrieving user data from an LDAP server. Note that you do not need to enable the object in a system policy to use the object with RUA.

When either a 3D Sensor or an RUA Agent on an Active Directory server detects a user login for a user who is not already in the database, an RUA user is added to the Defense Center user database.

Then, every five minutes, the Defense Center queries your LDAP servers 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 can take five to ten minutes for the Defense Center database to update with user metadata after RUA 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

Note that if you remove an LDAP user from your servers, the Defense Center does not remove that user from its local database. You must manually purge the user from the Defense Center database.

You must make sure that your LDAP servers use the default LDAP field names shown in the Field Names for Mapping LDAP Fields to RUA Fields table on page 1457. For example, if you are using an Active Directory server, RUA 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 one of the fields, the

Keysight, EX-1004
Page 1456

Defense Center cannot populate its user database with the information in that field.

Field Names for Mapping LDAP Fields to RUA Fields

LDAP Server Field Name	Defense Center Field Name	Servers Where Field Exists by Default
cn	Username	• Microsoft Active Directory on Windows Server 2003 • Sun Directory Services 5.2 P4 on Windows 2003 and XP • OpenLDAP on Linux
samaccountname	Username	• Microsoft Active Directory on Windows Server 2003
uid	Username	• Sun Directory Services 5.2 P4 on Windows 2003 and XP • OpenLDAP on Linux
givenname	First Name	• Microsoft Active Directory on Windows Server 2003 • Sun Directory Services 5.2 P4 on Windows 2003 and XP • OpenLDAP on Linux
sn	Last Name	• Microsoft Active Directory on Windows Server 2003 • Sun Directory Services 5.2 P4 on Windows 2003 and XP • OpenLDAP on Linux
mail	E-mail	• Sun Directory Services 5.2 P4 on Windows 2003 and XP • OpenLDAP on Linux
userprincipalname	email	• Microsoft Active Directory on Windows Server 2003
department	Department	• Sun Directory Services 5.2 P4 on Windows 2003 and XP
distinguishedname	Department	• Microsoft Active Directory on Windows Server 2003
telephonenumber	Phone	• OpenLDAP on Linux

If there is no field on a particular LDAP server type, and you want that metadata, you must create the field on the server. For example, only OpenLDAP has a phone number field by default. If you are using an Active Directory server and want your user database on the Defense Center to include phone numbers, you must create a field called `telephonenumber` on the Active Directory server. This allows the Defense Center to pull the metadata from that field.

For information on configuring a Defense Center-LDAP server connection, see Connecting to LDAP Servers on page 1478.

Keysight, EX-1004
Page 1457

Understanding Sourcefire RUA Agents

Platform: DC

Component: RUA

If your organization uses Microsoft Active Directory to implement LDAP, you can install Sourcefire RUA Agents on your Active Directory servers. RUA Agents monitor users as they log into the network or when they authenticate against Active Directory credentials for any other reason (for example, your organization may use services or applications that rely on Active Directory for centralized authentication).

IMPORTANT! If your organization does not use Active Directory, you **must** configure 3D Sensors with RUA to obtain user login data. For more information, see Understanding 3D Sensors and RUA on page 1459.

The RUA Agent is less than 1MB in file size and runs as a background Microsoft .NET Framework service, typically using less than 5% of CPU resources, depending on the processing power of the Active Directory server and the frequency and quantity of authentications. The agent uses approximately 15MB of RAM. If the Microsoft .NET Framework version 2.0 is not installed on your Active Directory servers, the RUA Agent setup will guide you through the .NET Framework installation process. You should install RUA Agents on all your Active Directory servers to get complete coverage of your network.

After you install RUA Agents on your Active Directory servers, you must configure your Defense Center to communicate with the agents. Then, when a login or other authentication occurs, an RUA Agent sends the following information to the Defense Center:

- the user's LDAP username
- the time of the login or other authentication
- the IP address of the user's host

When the Defense Center receives the information, it logs it to the RUA events database as a User Login event. If the user is already in the Defense Center database, the Defense Center correlates the login information with the user identity data in its database, and logs the information to the RUA history database so that you can view a history of the user's logins over time. If the user is not in the database, the Defense Center adds it.

For information on installing and configuring RUA Agents, see Configuring an RUA Agent on an Active Directory Server on page 1483.

Keysight, EX-1004
Page 1458

Understanding 3D Sensors and RUA

Platform: DC

Component: RUA

You can use the RUA and RNA components on most 3D Sensors (3D9800 and Crossbeam-based software sensors do not support RUA) instead of—or in addition to— the Sourcefire RUA Agents on your Active Directory servers:

- 3D Sensors with RUA can detect LDAP, POP3, IMAP, SMTP, and AIM logins.
- 3D Sensors with RNA can detect Oracle and SIP logins.

Setting up 3D Sensors to detect user logins is a two-step process. First, you must decide which sensors you are going to use. In general, you will want to monitor your internal network. Depending on your organization and network infrastructure, you may also want to monitor your DMZ.

If you are adding RUA to an existing deployment, consider the placement of your current 3D Sensors. Do they provide you with complete coverage of the networks you want to monitor? Do you have enough free resources on these sensors to add RUA and RNA detection engines? If not, you may need to deploy additional sensors.

Note that the number of detection engines you can create is limited by the 3D Sensor model and by how many other detection engines of different types you have already created. For example, on the 3D500, you can create one RUA detection engine or one RNA detection engine, but not both. For more information, see Understanding Detection Resources and 3D Sensor Models on page 202.

If you are setting up the Sourcefire 3D System for the first time, you must address the same issues. That is, make sure you deploy 3D Sensors so that you have complete coverage of the networks you want to monitor, and deploy sensors that can run the components of the Sourcefire 3D System you want to use. For detailed information on where and how to deploy 3D Sensors, refer to the *Sourcefire 3D Sensor Installation Guide*.

After you decide which sensors you are going to use in your RUA deployment, you must create RUA or RNA detection engines on those sensors. RUA detection engines passively observe your organization's network traffic and detect user logins. When a login occurs, the 3D Sensor sends the following information to the Defense Center:

- the username identified in the login event
- 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)

Keysight, EX-1004
Page 1459

- the user's email address (for POP3, IMAP, and SMTP logins)
- the name of the detection engine that detected the login, as well as the name of the sensor it resides on

The Defense Center logs the information in the RUA events database as a User Login event, as well as in the RUA history database, which stores records of individual user logins.

IMPORTANT! 3D Sensors with RUA interpret only Kerberos logins for LDAP connections as LDAP authentications. 3D Sensors with RUA cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.

The Defense Center also attempts to correlate the information with its user database. For example, if a 3D Sensor with RUA detects a POP3 login for a user with the same email address as a user that already exists in the database, the POP3 login will be associated with that user. If there is no data in the login event that the Defense Center can correlate with its user database, the Defense Center creates a new user record. AIM, SIP, and Oracle logins, for example, always create new user records.

IMPORTANT! SMTP logins are not logged unless there is already a user with a matching email address in the database.

For information on configuring 3D Sensors with RUA, see Setting up 3D Sensors and RUA on page 1487.

Understanding the Defense Center and RUA

Platform: DC

Component: RUA

You must use a Defense Center to configure RUA; you cannot configure it on a standalone 3D Sensor. Although 3D Sensors with RUA and RNA can detect user logins and collect user identity data, the sensors send that data to their managing Defense Center where you can view and analyze it. You must also use the Defense Center to view user identity data and RUA events. Note, however, that you can use the Master Defense Center to view user identity data correlated with intrusion, compliance, and white list events.

This section explains some important things to keep in mind when you are using the Defense Center to configure RUA.

Install an RUA feature license on the Defense Center.

For you to view user identity data and RUA events, your Defense Center must have a RUA feature license. In a high-availability environment, both Defense

Keysight, EX-1004
Page 1460

Centers must have RUA licenses. If one Defense Center does not have a RUA license, you will not be able to use it to view user identity data or RUA events.

IMPORTANT! An RUA Agent can only connect to one Defense Center at a time. In an high-availability environment, if the primary Defense Center fails, you must make sure that your RUA Agents can communicate with the secondary Defense Center. For more information, see Configuring an RUA Agent on an Active Directory Server on page 1483.

RUA feature licenses specify the number of users you can monitor with RUA. After you reach your licensed limit, RUA stops adding new users to the Defense Center database.

If your licensed limit drops to a number that is below the number of users in the Defense Center database (for example, if you install a new, smaller RUA license), the Defense Center automatically purges the database until it reaches the limit; oldest records are purged first.

After you reach your licensed limit, to detect additional users, you have several options:

- You can purchase and install a larger RUA license.
- You can manually delete old or inactive users from the Defense Center database; see Viewing RUA Users on page 1490.
- You can purge all users from the Defense Center database and start over; see Purging the RNA and RUA Databases on page 2014.

For information on setting up RUA licensing, see Licensing RUA on page 1475.

If necessary, restrict RUA logging to preserve RUA licenses.

To make sure you do not surpass your RUA feature license limit, you can either purchase an RUA license that allows you to monitor more than the actual number of users in your organization, or you can use the system policy to restrict RUA from logging usernames obtained through specific protocols.

Restricting RUA helps minimize username clutter and preserve RUA licenses. For example, obtaining usernames through protocols such as AIM, POP3, and IMAP can introduce usernames not relevant to your organization due to network access from contractors, visitors, and other guests. In addition, AIM, SIP, and Oracle logins always create duplicate user records. This is because these logins are not associated with any of the user metadata that RUA obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your 3D Sensors detect.

Keysight, EX-1004
Page 1461

For more information, see Restricting RUA Logging on page 1476.

IMPORTANT! Sourcefire RUA Agents installed on Microsoft Active Directory LDAP servers collect **only** LDAP user login information. Therefore, unless your RUA implementation includes 3D Sensors, filtering non-LDAP logins has no effect. For more information on RUA Agents and 3D Sensors and RUA, see How Do I Choose an RUA Implementation? on page 1463.

Set realistic database limits.

RUA stores data on the Defense Center in three places: in the users database, the RUA events database, and the RUA history database. You should make sure that you configure the Defense Center to store an adequate amount of data for your analysis.

User identity data is stored in the Defense Center users database, with one record for each unique user. As explained above, your RUA feature license determines the number of users you can store in the database. This database is populated when RUA detects user logins, either via an RUA Agent on an Active Directory server or via a 3D Sensor with RUA. The Defense Center uses information in the login event to create a record in the user database (if there is not already a record of that user). Your optional Defense Center-LDAP server connections (see Understanding LDAP Servers and RUA on page 1456) augments user records with available metadata. For more information on RUA users, see Working with RUA Users on page 1489.

RUA events, which track not only user logins, but also additions and deletions of users from the Defense Center database, are stored in the RUA event database, which by default stores 10 million events. For more information on RUA events, see Working with RUA Events on page 1500.

Individual user logins are stored in the RUA history database. This data is used to generate host histories, which track the users that have logged into an individual host over a period of time, and user histories, which track the hosts that an individual user has logged into over a similar period of time. By default, the RUA history database stores 10 million user logins. For more information on host histories and user histories, see Working with User History in the Host Profile on page 973 and Understanding User Details and Host History on page 1495.

For information on changing the number of events you can store in the RUA events and RUA history database, see Configuring Database Event Limits on page 1747.

Configure a connection with LDAP servers.

If your organization uses LDAP, you can set up connections between the Defense Center and LDAP servers. The Defense Center uses the information on the LDAP servers to augment existing user records with available metadata. For more information, see Understanding LDAP Servers and RUA on page 1456 and Connecting to LDAP Servers on page 1478.

Keysight, EX-1004
Page 1462

Configure connections with Sourcefire RUA Agents.

If your organization uses Active Directory, you can set up connections between the Defense Center and the Sourcefire RUA Agents you can install on your Active Directory servers. RUA Agents send LDAP login information to the Defense Center. For more information, see Understanding Sourcefire RUA Agents on page 1458 and Configuring an RUA Agent on an Active Directory Server on page 1483.

Configure connections with 3D Sensors with RUA.

You can set up 3D Sensors with RUA to collect LDAP, POP3, IMAP, SMTP, and AIM user login information and send it to the Defense Center. For more information, see Understanding 3D Sensors and RUA on page 1459 and Setting up 3D Sensors and RUA on page 1487.

Configure connections with 3D Sensors with RNA.

If you add an RUA feature license to your Defense Center and have not restricted RUA logging to exclude Oracle and SIP, 3D Sensors with RNA automatically collect Oracle and SIP user login information from the monitored networks specified in your RNA detection policy. For information on how to set up an RNA detection policy, see What is an RNA Detection Policy? on page 889,

How Do I Choose an RUA Implementation?

Platform: DC

Component: RUA

As explained in detail in Understanding RUA on page 1453, an RUA deployment includes:

- a Defense Center that you can use to configure RUA and view and analyze the data that RUA collects
- Either 3D Sensors or Sourcefire RUA Agents (or both) that detect user logins
- Optionally, LDAP servers that you can use to augment user records in the Defense Center's user database with available metadata

The main decision you must make when choosing an RUA implementation is whether you want to use RUA Agents or 3D Sensors (or both) to collect user login data. You must also decide whether you want to use LDAP servers to augment user records in the Defense Center's user database with available metadata. For information on the advantages and disadvantages of these choices, see the following sections:

- Using LDAP Servers on page 1464
- Using the Sourcefire RUA Agent on page 1464
- Using 3D Sensors and RUA on page 1465
- Using Both the Sourcefire RUA Agent and 3D Sensors on page 1466

Keysight, EX-1004
Page 1463

Using LDAP Servers

Platform: DC

Component: RUA

If your organization uses LDAP, Sourcefire recommends that you create a connection between the Defense Center and one or more LDAP servers, so you can augment the records in the Defense Center's user database with information about your LDAP users.

RUA supports connections to LDAP servers running Microsoft Active Directory on Windows Server 2003, Sun Directory Services 5.2 P4 on Windows 2003 and XP, or OpenLDAP on Linux. You must have TCP/IP access from the Defense Center to the LDAP server. Also, you should work closely with your LDAP administrators to obtain the configuration information you need to set up Defense Center-LDAP server connections.

For more information on using LDAP servers in your RUA deployment, see Understanding LDAP Servers and RUA on page 1456.

Using the Sourcefire RUA Agent

Platform: DC

Component: RUA

If you use Microsoft Active Directory LDAP servers, Sourcefire **strongly** recommends that you install Sourcefire RUA Agents on all of your Active Directory servers to collect user login data and send that data to the Defense Center. There are two advantages to using RUA Agents:

- If installed on all your Active Directory servers, RUA Agents provide comprehensive coverage of LDAP logins across your entire organization, so you do not have to rely on or deploy additional 3D Sensors to cover your network segments.
- You can dedicate the hardware resources of your 3D Sensors to running IPS and RNA.

In general, you should not be concerned about running the RUA Agent on your Active Directory servers. The RUA Agent is less than 1MB in file size and runs as a background Microsoft .NET Framework service, typically using less than 5% of CPU resources, depending on the processing power of the Active Directory server and the frequency and quantity of authentications. The agent uses

Keysight, EX-1004
Page 1464

approximately 15MB of RAM. However, there are a few potential reasons for not using RUA Agents:

- Your organization may have a policy against installing third-party software on your mission-critical Active Directory servers.
- Your organization may require certification testing of the RUA Agent, which can be a lengthy process. If you want to benefit from RUA right away, you can implement it using 3D Sensors and then later replace (or supplement) them with RUA Agents.
- RUA Agents only detect LDAP logins. If you want to detect other types of logins, you must use 3D Sensors.

Note that you will have to work closely with your Active Directory administrators to install the RUA Agents and obtain the configuration information you need to set up Defense Center-RUA Agent connections.

TIP! There are advantages to using both RUA Agents and 3D Sensors in the same deployment. For more information, see Using Both the Sourcefire RUA Agent and 3D Sensors on page 1466.

For more information on RUA Agents, see Understanding Sourcefire RUA Agents on page 1458.

Using 3D Sensors and RUA

Platform: DC

Component: RUA

If your organization does not use Microsoft Active Directory LDAP servers and thus you cannot use Sourcefire RUA Agents, you **must** use 3D Sensors to collect user login information. In addition, because RUA Agents only detect LDAP logins, if you want to detect other types of logins, you must use 3D Sensors.

If you use 3D Sensors and RUA, there are several questions you must answer:

- Can you achieve complete network coverage with your existing 3D Sensors?
- Do your 3D Sensors have enough available resources to devote to RUA and RNA?
- Can you create RUA detection engines on your 3D Sensors to detect POP3, IMAP, SMTP, and AIM logins? The 3D9800 and Crossbeam-based software sensors do not support RUA detection engines.
- If your organization uses LDAP, do you use the Kerberos authentication protocol for your LDAP connections? 3D Sensors with RUA cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.

If you cannot answer "yes" to these questions, you may need to purchase additional 3D Sensors or revisit your network encryption.

Keysight, EX-1004
Page 1465

Note that AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that RUA obtains from LDAP servers, nor are they associated with any of the information contained in the other types of logins that your 3D Sensors detect. Therefore, if you are using 3D Sensors in your RUA deployment, you may want to purchase a larger RUA feature license, or restrict RUA logging.

TIP! There are advantages to using both RUA Agents and 3D Sensors in the same RUA deployment. For more information, see the next section, Using Both the Sourcefire RUA Agent and 3D Sensors.

For more information on 3D Sensors and RUA, see Understanding 3D Sensors and RUA on page 1459.

Using Both the Sourcefire RUA Agent and 3D Sensors

Platform: DC

Component: RUA

If your organization uses Microsoft Active Directory LDAP servers and you use Sourcefire RUA Agents to collect LDAP user login information, but you also want to detect other types of logins, you must use 3D Sensors in addition to your RUA Agents.

There are no compatibility issues with the use of both methods. The issues that you must consider when deciding to implement both methods are the same as if you were deciding to implement either method by itself.

Note that even though both RUA Agents and 3D Sensors with RUA detect LDAP logins, you will not see duplicate login events on the Defense Center—the Defense Center records the first time that a user logs into a specific host and disregards subsequent logins. If another user logs into that host, however, the Defense Center will record the new login. Then, if the original user logs in again, his or her new login is recorded.

For specific information about using either component, see Using the Sourcefire RUA Agent on page 1464 and Using 3D Sensors and RUA on page 1465.

What Information Does RUA Provide?

Platform: DC

Component: RUA

RUA analyzes the data it collects and provides you with the information described in the following sections.

Keysight, EX-1004
Page 1466

RUA Users

The Defense Center maintains a database of *RUA users* on your network. The database is populated when RUA detects user logins, either via an RUA Agent on an Active Directory server or via a 3D Sensor. The Defense Center uses information in the login event to create a record in the user database (if there is not already a record of that user). Your optional Defense Center-LDAP server connections augment user records with available metadata. The number of users the Defense Center can store in its database depends on your RUA feature license.

For more information, see Working with RUA Users on page 1489.

RUA Events

RUA events are triggered by user logins, additions and deletions of RUA users from the Defense Center database, and failures of RUA to add a user to the database (because you have reached your licensed limit). The events generated by RUA can alert you to the user activity on your network and provide you with the information you need to respond appropriately.

Whenever possible the Sourcefire 3D System correlates RUA events 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.

RUA provides a predefined workflow that you can use to analyze the events that are generated for your network. You can use this predefined workflow, or you can create custom workflows that display only the information that matches your specific needs.

For more information, see Working with RUA Events on page 1500.

Host Histories and User Histories

RUA uses its record of individual user logins to generate *host histories*, which track the hosts that a user has logged into. The host history provides a graphic representation of the last twenty-four hours of the user's activity. For more information, see Understanding User Details and Host History on page 1495.

RUA also uses its record of individual user logins to generate *user histories*, which track the users that have logged into an individual host. The user history provides a graphic representation of the last twenty-four hours of the logins on the host. For more information, see Working with User History in the Host Profile on page 973.

Keysight, EX-1004
Page 1467

How Can You Use RUA?

Platform: DC

Component: RUA

You can use RUA not only to view a complete record of your users and their activity your network, but you can also take advantage of the Sourcefire 3D System's correlation of RUA events 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 RUA events in compliance rules. Based on the type of RUA event generated as well as other criteria that you specify, you can build compliance rules that, when used in a compliance policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria.

The following sections provide examples of how you can use RUA to correlate threat, endpoint, and network intelligence with user identity information:

- Example: Identifying the Owner of an Attacked Host on page 1468
- Example: Unauthorized Access of High-Criticality Host on page 1469
- Example: Identifying Excessive Bandwidth Users on page 1470
- Example: Identifying Instant Messenger Users on page 1471

Example: Identifying the Owner of an Attacked Host

Platform: DC

Component: IPS + RNA + RUA

Impact flags are indicators of the correlation between intrusion data, RNA network discovery events, and vulnerability information. A red impact flag means that a host is vulnerable to an attack represented by an intrusion event.

You can use RUA to determine who was logged into the attacked host at the time of the attack. With this knowledge, you may be able to lessen the time it takes to mitigate the effects of the attack.

To identify the owners of attacked hosts:

Access: Any IPS/Admin

1. Search for intrusion events with red impact flags.

   For more information, see Searching for Intrusion Events on page 309.

Keysight, EX-1004
Page 1468

2. View the search results in a workflow that contains a table view of intrusion events.

   All of the predefined intrusion event workflows contain a table view of events; after you select a workflow click **Table View of Events** in the top left corner of the event view.

3. Evaluate the events.

   The **Destination User** column tells you who was logged into each host at the time of each attack.

Example: Unauthorized Access of High-Criticality Host

Platform: DC

Component: IPS + RNA + RUA

Your organization's critical servers need extra protection and monitoring. You can use the Sourcefire 3D System to track and notify you of unauthorized access attempts to your critical servers from internal users. When you receive a notification, you can use RUA to determine who attempted the access.

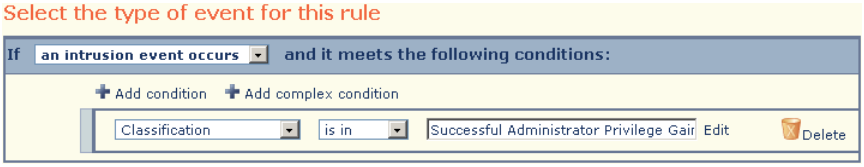
To identify unauthorized users accessing critical servers:

Access: RNA/Admin

1. Use RNA to assign a high host criticality to your critical servers.

   For more information, see Setting Host Attributes for Specific Hosts on page 1036.

2. Create a compliance rule that triggers on intrusion events with classifications that suggest unauthorized access or other suspicious activity.



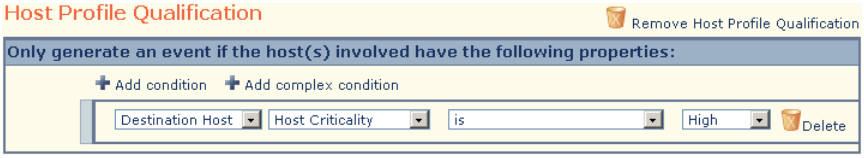
   These classifications can include:

   - Attempted User Privilege Gain
   - Unsuccessful User Privilege Gain
   - Successful User Privilege Gain
   - Attempted Administrator Privilege Gain
   - Successful Administrator Privilege Gain
   - An Attempted Login Using a Suspicious Username was Detected
   - Attempt to Login By a Default Username and Password

   For information on creating compliance rules, see Creating Rules for Compliance Policies on page 1214. For a list of classifications, see Defining the Event Classification on page 772.

Keysight, EX-1004
Page 1469

3. Add a host profile qualification that constrains the compliance rule so it only triggers if the target host has a high host criticality.

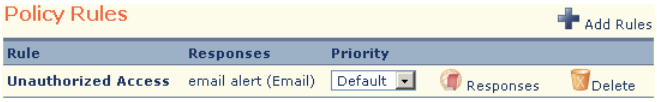


   For information on adding a host profile qualification to a compliance rule, see Adding a Host Profile Qualification on page 1233.

4. Save the compliance rule.
5. Create and activate an alert.

   For more information, see Creating Alerts on page 1285

6. Create a compliance policy. Add the rule you just created to the policy, and assign the alert you created as a response to that rule.



   For more information on creating compliance policies, see Creating Compliance Policies on page 1262.

7. Save and activate the compliance policy.

   When an internal user attempts unauthorized access or other suspicious activity targeted at one of your critical hosts, you are notified via the alert you created.

8. Use the information in the alert to search compliance events for the event that concerns you.

   For more information, see Searching for Compliance Events on page 1276.

9. View the search results in a workflow that contains a table view of compliance events.

   The predefined compliance workflow contains a table view of events.

10. Evaluate the event.

    The **Source User** column tells you who initiated the attack.

Example: Identifying Excessive Bandwidth Users

Platform: DC

Component: RNA + RUA

Your organization may prohibit excessive bandwidth use. You can use the Sourcefire 3D System to view which hosts use the most bandwidth, and then examine the host profiles of those hosts to determine who is logged into those hosts.

Keysight, EX-1004
Page 1470

To identify excessive bandwidth users:

Access: Any RNA/Admin

1. View a flow graph of the Top Ten Initiators by Traffic, which is graph of your 10 most active hosts on your monitored network, based on the total number of kilobytes transmitted by the hosts.

   For more information, see Viewing Flow Data on page 1093.

2. On the graph that appears, click any of the bars associated with a host that is using excessive bandwidth, then select **View Host Profile**.

   The host profile for that host appears in a new window. The User History section of the host profile tells you which users have logged into that host in the last 24 hours.

Example: Identifying Instant Messenger Users

Platform: DC

Component: RNA + RUA

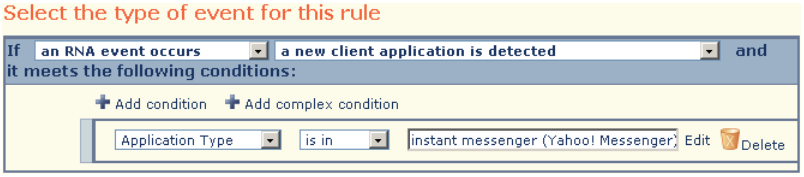
Use of instant messaging software may be a violation of your company policy. You can use the Sourcefire 3D System to track and notify you of unauthorized use of this software. When you receive a notification, you can use RUA to determine who is logged into the host running the instant messaging client application.

Note that this example outlines how to use RNA to detect instant messaging client use, then associate that use with RUA users. RNA can detect the use of multiple instant messaging clients, including AIM, Microsoft MSN Messenger, Yahoo! Messenger, and Skype. RUA, on the other hand, specifically detects AIM login attempts and records the usernames and IP addresses of AIM users. To see all the AIM users on your network, search for RUA users with a **User Type** of `aim`. For more information, see Searching for RUA Users on page 1497.

To identify users of instant messaging software:

Access: Any Analyst/Admin

1. Create a compliance rule that triggers when RNA detects instant messaging activity.

   For information on creating compliance rules, see Creating Rules for Compliance Policies on page 1214.

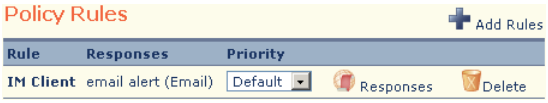
2. Save the compliance rule.

Keysight, EX-1004
Page 1471

3. Create and activate an alert.

   For more information, see Creating Alerts on page 1285

4. Create a compliance policy. Add the rule you just created to the policy, and assign the alert you created as a response to that rule.



   For more information on creating compliance policies, see Creating Compliance Policies on page 1262.

5. Save and activate the compliance policy.

   When RNA detects the use of an instant messaging client application on one of the hosts on your network, you are notified via the alert you created.

6. Use the information in the alert to search compliance events for the event that concerns you.

   For more information, see Searching for Compliance Events on page 1276.

7. View the search results in a workflow that contains a table view of compliance events.

   The predefined compliance event workflow contains a table view of events.

8. Evaluate the event.

   The **Source User** column tells you who was logged into the host at the time that RNA detected the instant messaging application.

Keysight, EX-1004
Page 1472

What are the Limitations of Sourcefire RUA?

Platform: DC

Component: RUA

The limitations of Sourcefire RUA are described in the Sourcefire RUA Limitations table.

Sourcefire RUA Limitations

Limitation	Description
non-Kerberos logins for LDAP connections	3D Sensors with RUA interpret only Kerberos logins for LDAP connections as LDAP authentications. 3D Sensors with RUA cannot detect encrypted LDAP authentications if they use other protocols, such as SSL or TLS. On the other hand, if your organization uses Active Directory LDAP servers, you can install Sourcefire RUA Agents on them. RUA Agents use the security logs on Active Directory servers to collect user login data and thus have no such limitations.
non-Active Directory LDAP servers	The Sourcefire RUA Agent is only supported on Active Directory. It is **not** supported on other types of LDAP servers, such as Sun Directory Services or OpenLDAP.
logoff detection	RUA detects user logins only. It does not detect logoffs.
multiple logins to the same host by different users	RUA assumes that only one user is logged into any given host at a time, and that the current user of a host is the user who logged in last.
multiple logins to the same host by the same user	The Defense Center 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 Defense Center records is the original login. If another user logs into that host, however, the Defense Center will record the new login. Then, if the original user logs in again, his or her new login is recorded.
double-byte characters	RUA does not detect usernames with double-byte characters (Chinese and Japanese).

Keysight, EX-1004
Page 1473

Sourcefire RUA Limitations (Continued)

Limitation	Description
disabled/deleted LDAP user accounts	If you remove or disable an LDAP user on your LDAP servers, the Defense Center does not remove that user from its local database, and that user continues to count against your licensed limit. You must manually purge the user from the Defense Center database (see RUA Users Table Functions on page 1491).
AOL Instant Messenger (AIM) login detection	3D Sensors with RUA can detect AIM logins using the OSCAR protocol only. While most AIM clients use OSCAR, some use TOC2.

Configuring RUA

Platform: DC

Component: RUA

Before you can correlate your network traffic with user identity data, you must license and configure RUA. The steps you take to configure RUA depends on how you plan to implement RUA data collection. For more information on deciding how to implement RUA, see How Do I Choose an RUA Implementation? on page 1463.

To configure RUA:

Access: Admin

1. Make sure your Defense Center has an RUA feature license installed. You cannot view RUA user identity data or RUA events without a license.

   For more information, see Licensing RUA on page 1475.

2. Do you want to restrict RUA logins based on protocol, to help minimize username clutter and preserve RUA licenses?
   - If **yes**, see Restricting RUA Logging on page 1476.
   - If **no**, continue with the next step.

Keysight, EX-1004
Page 1474

3. Does your organization use LDAP?
   - If **yes**, configure a connection between the Defense Center and your LDAP servers. The Defense Center will use the information on the LDAP servers to augment the user records in its database with available metadata.

     For more information, see Connecting to LDAP Servers on page 1478. If you have already configured an LDAP connection, see Managing RUA LDAP Connections on page 1482 for more information on enabling that connection.
   - If **no**, skip to step 5.
4. Does your organization use Active Directory LDAP servers?
   - If **yes**, configure Sourcefire RUA Agents on your Active Directory servers to send LDAP login information to the Defense Center. For more information, see Configuring an RUA Agent on an Active Directory Server on page 1483.
   - If **no** (or if you choose not to install RUA Agents on your Active Directory servers), continue with the next step.
5. Optionally, set up 3D Sensors with RUA to collect LDAP, POP3, IMAP, SMTP, and AIM login information. Note that you **must** configure 3D Sensors with RUA to collect user login information if:
   - Your organization does not use Active Directory LDAP servers.
   - Your organization uses Active Directory LDAP servers, but you did not configure an RUA Agent.
   - You want to obtain POP3, IMAP, SMTP, and AIM login information.

   For more information, see Setting up 3D Sensors and RUA on page 1487.

Licensing RUA

Platform: DC

Component: RUA

For you to view user identity data and RUA events, your Defense Center must have an RUA feature license. For more information, see Understanding the Defense Center and RUA on page 1460.

If your organization has purchased an RUA feature license but you do not have the license file, you can request it from the Defense Center web interface. Before beginning, you should have the 12-digit feature license serial number provided by Sourcefire when you purchased RUA. If you do not have the serial number, you can find it by logging into the Sourcefire Support Site (https://support.sourcefire.com/), clicking **Account**, then clicking **Products & Contracts**. The serial number appears in the Sourcefire Software & Licenses section.

Keysight, EX-1004
Page 1475

To add an RUA feature license:

Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **License**.

   The License page appears.

3. Click **Add New License**.

   The Add Feature License page appears.

4. Do you already have the license file?

   - If **yes**, skip to step 6.
   - If **no**, click **Get License**.

   The Licensing Center web site appears.

   IMPORTANT! If your web browser cannot access the Internet, you must switch to a host that can access it. Copy the license key at the bottom of the page and browse to https://keyserver.sourcefire.com/.

5. Follow the on-screen instructions for a feature license to obtain your license file, which will be sent to you in an email.

6. Copy the license file from the email you received, paste it into the **License** field, and click **Submit License**.

   If the license is valid, the RUA license is installed on the Defense Center.

7. Do you want to restrict RUA logins based on protocol, to help minimize username clutter and preserve RUA licenses?

   - If **yes**, continue with the next section, Restricting RUA Logging.
   - If **no**, and your organization uses LDAP, skip to Connecting to LDAP Servers on page 1478.
   - If **no**, and your organization does not use LDAP, skip to Setting up 3D Sensors and RUA on page 1487.

Restricting RUA Logging

Platform: DC

Component: RUA

You can use the RUA settings in the system policy to filter which types of network activity cause RUA to add users to the database.

The RUA feature license on the Defense Center (see Licensing RUA on page 1475) specifies the number of users you can monitor with RUA. After you

Keysight, EX-1004
Page 1476

reach your licensed limit, RUA stops adding new users to the Defense Center database.

Restricting RUA helps minimize username clutter and preserve RUA licenses. For example, obtaining usernames through protocols such as AIM, POP3, and IMAP can introduce usernames not relevant to your organization due to network access from contractors, visitors, and other guests. In addition, AIM, SIP, and Oracle logins always create duplicate user records. This is because these logins are not associated with any of the user metadata that RUA obtains from LDAP servers, nor are they associated with any of the information contained in the other types of login that your 3D Sensors detect.

IMPORTANT! Sourcefire RUA Agents installed on Microsoft Active Directory LDAP servers collect **only** LDAP user login information. Therefore, unless your RUA implementation includes 3D Sensors, filtering non-LDAP logins has no effect. For more information on RUA Agents and 3D Sensors and RUA, see How Do I Choose an RUA Implementation? on page 1463.

To filter RUA users based on network activity type:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the RUA settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the RUA 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 1733, and click **Save**.

   In either case, the Access List page appears.

3. Click **RUA Settings**.

   The RUA Detection Settings page appears.

4. Select the check boxes that correspond to the types of logins that will create RUA users.

   By default, all login types cause RUA to add users to the database.

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 1736 for more information.

Keysight, EX-1004
Page 1477

6. Does your organization use LDAP?
   - If **yes**, continue with the next section, Connecting to LDAP Servers.
   - If **no**, skip to Setting up 3D Sensors and RUA on page 1487.

Connecting to LDAP Servers

Platform: DC

Component: RUA

Optionally, you can configure connections between your Defense Center and your LDAP servers. When RUA detects an LDAP, POP3, IMAP, or AIM user login for a user who is not already in the database, an RUA user is added to the Defense Center user database. Then, every five minutes, the Defense Center queries the LDAP servers you configured to obtain 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. For more information, see Understanding LDAP Servers and RUA on page 1456.

IMPORTANT! Make sure your LDAP servers use the default LDAP field names.

To create an Defense Center-LDAP server connection, you must create and configure an *authentication object*, which contains your settings for connecting to and retrieving user data from the LDAP server.

You use the same method to configure an authentication object as you would if you were setting up external authentication for user accounts on your Sourcefire appliances. Note, however, that RUA does not require all of the information that you must provide when configuring external authentication, nor does it support all of the functionality of an external authentication object. You do not need to supply a user name template for RUA objects, and RUA does not support encryption. In addition, you do not need to enable the object in a system policy to use the object with RUA.

When the local appliance 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 will have 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 once you identify a primary server, you can automatically retrieve a list of available base DNs from the server and select the appropriate base DN.

To allow the local appliance to access the user objects, you must supply user credentials for a user with appropriate rights to the authentication objects you

Keysight, EX-1004
Page 1478

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.

For the authentication method specific parameters, you can use the LDAP naming standards and filter and attribute syntax defined in 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 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.

You control the encryption level for the connection to the LDAP server. You can choose Transport Layer Security (TLS) or Secure Sockets Layer (SSL) encryption, or leave the connection unencrypted. Note that if you change the encryption method after specifying the port, the port may change to the default port for an encrypted connection. Remember to check that access to the server is not blocked by a firewall and that the port you have configured in the object is open.

Also 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. If you enter `10.10.10.250` in the authentication profile and `computer1.example.com` in the certificate, the connection fails, even if computer1.example.com has an IP address of 10.10.10.250. Changing the name of the server in the authentication profile to `computer1.example.com` causes the connection to succeed.

After you create the authentication object, you must enable the connection in your RUA configuration. You must work closely with your LDAP administrators to obtain the configuration information you need to set up Defense Center-LDAP server connections.

TIP! For more information on authentication objects, see Creating Advanced LDAP Authentication Objects on page 1642. For sample configurations showing how you might use different configuration options for connections to different LDAP server types, see LDAP Authentication Object Examples on page 1657.

Keysight, EX-1004
Page 1479

To connect to an LDAP server for RUA:

Access: Admin

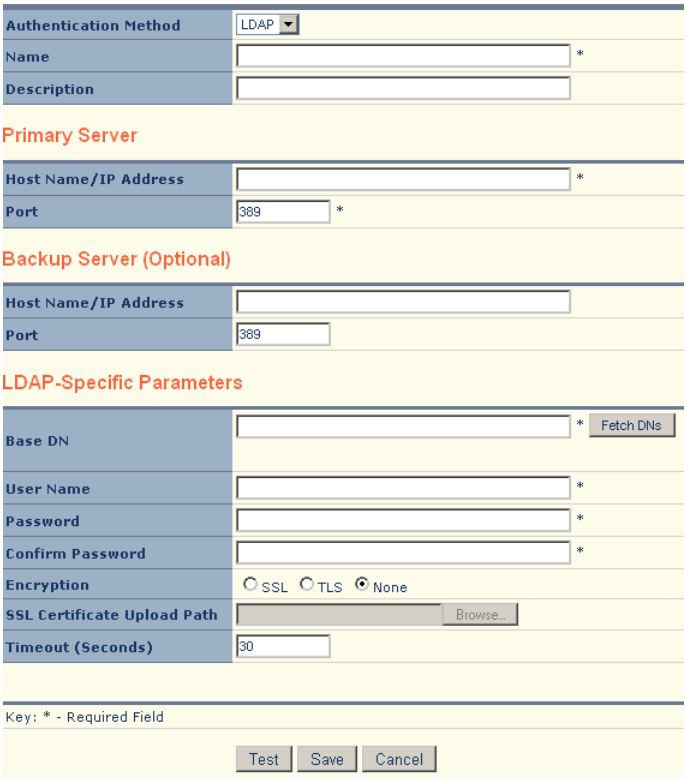
1. Select **Operations > Configuration > RUA.**

   The RUA configuration page appears.



2. Click **Add LDAP Connection**.

   The Create Authentication Object page appears.



3. Select **LDAP** from the **Authentication Method** drop-down list.
4. Type a name and description for the authentication server in the **Name** and **Description** fields.

Keysight, EX-1004
Page 1480

5. 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.
6. Optionally, modify the port used by the primary authentication server in the **Primary Server Port** field.
7. 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.
8. Optionally, modify the port used by the primary authentication server in the **Backup Server Port** field.

IMPORTANT! You cannot connect to an LDAP server over an encrypted connection (using SSL or other encryption) when connecting from RUA.

9. 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`.
10. 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 would type `uid=NetworkAdmin,ou=security,dc=example,dc=com.`
11. Re-type the password in the **Confirm Password** field.
12. 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.

Keysight, EX-1004
Page 1481

13. 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.

14. Type the number of seconds that should elapse before rolling over to the backup connection in the **Timeout** field.

15. To test the authentication object, select **Show Details** and then click **Test**.

16. To save the authentication object, click **Save**.

    The authentication object is listed under LDAP Connections on the RUA configuration page.

17. To use the LDAP object, click **Enable**.

18. Repeat steps 2 through 17 to add additional Defense Center-LDAP server connections.

19. To get RUA events, you must configure at least one RUA Agent or a 3D Sensor with RUA. You have two options.

    - If you are using Active Directory and want to configure RUA Agents on your Active Directory servers, continue with Configuring an RUA Agent on an Active Directory Server.
    - If you are not using Active Directory, or do not want to configure RUA Agents, skip to Setting up 3D Sensors and RUA on page 1487.

    You can use both RUA Agents and 3D Sensors with RUA. For more information, see How Do I Choose an RUA Implementation? on page 1463.

    For information on managing RUA LDAP connections, see Managing RUA LDAP Connections on page 1482.

Managing RUA LDAP Connections

Platform: DC

Component: RUA

You can enable, disable or delete existing LDAP connections from the RUA configuration page. Note that you can delete an enabled connection without disabling it first.

To manage RUA LDAP connections:

Access: Admin

1. Select **Operations > Configuration > RUA.**

   The RUA configuration page appears.

Keysight, EX-1004
Page 1482

2. You have several choices:
   - To enable a disabled LDAP connection, click **Enable**.
   - To disable an enabled LDAP connection, click **Disable**.
   - To delete an LDAP connection, click **Delete** and confirm the deletion.

Configuring an RUA Agent on an Active Directory Server

Platform: DC

Component: RUA

If your organization uses Microsoft Active Directory, you can install Sourcefire RUA Agents on your Active Directory servers. RUA Agents monitor users as they log into the network or when they authenticate against Active Directory credentials for any other reason (for example, your organization may use services or applications that rely on Active Directory for centralized authentication).

RUA Agents run as a background Microsoft .NET Framework service on your Active Directory server. If the Microsoft .NET Framework version 2.0 is not installed on your Active Directory servers, the RUA Agent setup will guide you through the .NET Framework installation process. You should install RUA Agents on all your Active Directory servers to get complete coverage of your network. For more information, see Understanding Sourcefire RUA Agents on page 1458.

Note that installation of RUA agents is supported only on Microsoft Windows Server 2003 SP2 and Microsoft Windows Server 2008 R2.

To collect Active Directory user login information using RUA Agents, you must configure the Defense Center to connect to the Active Directory servers where the RUA Agents will reside. Then, you can install and configure the agents.

For more information, see

- Configuring the Defense Center to Connect to RUA Agents on page 1483
- Installing an RUA Agent on an Active Directory Server on page 1484

Configuring the Defense Center to Connect to RUA Agents

Platform: DC

Component: RUA

The first step in collecting LDAP user login information using Sourcefire RUA Agents is to configure the Defense Center to connect to the RUA Agents you plan to install on your Active Directory servers.

To configure the Defense Center to connect to an RUA Agent:

Access: Admin

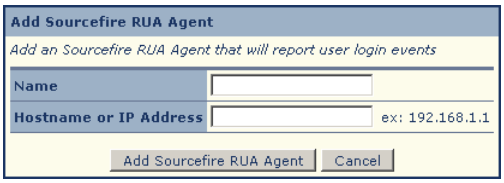
1. Select **Operations > Configuration > RUA**.

   The RUA configuration page appears.

Keysight, EX-1004
Page 1483

2. Click **Add Sourcefire RUA Agent**.

   The Add Sourcefire RUA Agent pop-up window appears.



3. Type a descriptive name for the Agent in the **Name** field.
4. Type the IP address or host name of the Active Directory server where the RUA Agent will reside in the **Hostname or IP Address** field.
5. Click **Add Sourcefire RUA Agent**.

   The Defense Center is configured to connect to the RUA Agent on the Active Directory server.

   TIP! To delete the Defense Center-RUA Agent connection, click **Delete** next to the connection you want to delete.

6. Repeat steps 2 through 5 to add connections to RUA Agents on additional Active Directory servers.
7. Continue with the next section, Installing an RUA Agent on an Active Directory Server.

Installing an RUA Agent on an Active Directory Server

Platform: DC

Component: RUA

After you configure the Defense Center to connect to your Active Directory servers, you must install and configure the Sourcefire RUA Agents. You should install RUA Agents on all your Active Directory servers to get complete coverage of your network. The RUA Agent is supported on Microsoft Windows Server 2003 SP2 and Microsoft Windows Server 2008 R2.

The RUA Agent requires that Microsoft .NET Framework version 2.0, which is available from Microsoft as the .NET Framework Version 2.0 Redistributable Package (`dotnetfx.exe`), be installed on the Active Directory server. If it is not already installed, the RUA Agent setup will prompt you to download and install it.

If Sourcefire releases a new version of the RUA Agent, you must uninstall the old agent from your Active Directory servers using the Microsoft Windows Add or

Keysight, EX-1004
Page 1484

Remove Programs tool. Then, you can install the new version of the RUA Agent using the procedure in this section.

TIP! Before you uninstall the old version of the RUA Agent, select **Start > All Programs > Sourcefire RUA Agent > Agent Settings** on the Active Directory server to determine the Defense Center and server timeout settings. When you install the new RUA Agent, use the same values.

To install an RUA Agent on an Active Directory server:

Access: Any

1. Download the RUA Agent setup file (`SourcefireRUAAgentSetup.zip`) 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 Active Directory server where you want to install the RUA Agent.
3. Unzip the setup file and open the Windows Installer file (`SourcefireRUAAgentSetup.msi`) that appears.

   If Microsoft .NET Framework version 2.0 is not installed on your Active Directory server, you are prompted to download and install it. After you download and install the framework, open the RUA Agent installer file again.

   The RUA Agent setup wizard appears.
4. Follow the prompts in the setup wizard to install the RUA Agent and to read the release notes.
5. On the Active Directory server, enable security logging.

   The RUA Agent does not record user logins unless security logging is enabled. Use the Security log in the Event Viewer to verify that logging is enabled (**Start > All Programs > Administrative Tools > Event Viewer**).
6. On the Active Directory server, select **Start > All Programs > Accessories > Administrative Tools> Services**.

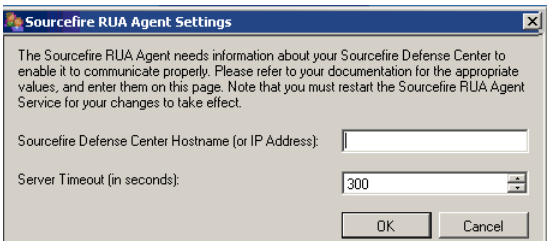
   The service manager appears.
7. In the services list, right-click **Sourcefire RUA Agent** and select **Stop**.

   The RUA Agent stops.

Keysight, EX-1004
Page 1485

8. On the Active Directory server, select **Start > All Programs > Sourcefire RUA Agent > Agent Settings**.

   The Sourcefire RUA Agent Settings window appears.



9. Type the IP address or host name of your Defense Center in the **Sourcefire Defense Center Hostname (or IP Address)** field.
10. Optionally, modify the number of seconds in the **Server Timeout (in seconds)** field.
11. Click **OK** to save your settings.
12. In the service manager, in the services list, right-click **Sourcefire RUA Agent** and select **Properties**.

    The Sourcefire RUA Agent Properties window appears.
13. Ensure that the RUA Agent starts automatically when the Active Directory server boots by selecting **Automatic** from the **Startup type** drop-down list on the **General** tab.
14. Click OK.

    The Sourcefire RUA Agent Properties window closes.
15. In the service manager, in the services list, right-click **Sourcefire RUA Agent** and select **Start**.

    The RUA Agent starts as a service on the Active Directory server and begins reporting logins to the server to your Defense Center.
16. Repeat steps 1 through 15 to add install RUA Agents on additional Active Directory servers.
17. Do you want to configure 3D Sensors with RUA to supplement the LDAP logins detected by the RUA Agent with data gathered from POP3, IMAP, SMTP, and AIM logins?
    - If **yes**, continue with the next section, Setting up 3D Sensors and RUA.
    - If **no**, you are finished configuring RUA. If you have installed an RUA feature license on your Defense Center, you can begin viewing and analyzing RUA data, including Oracle and SIP user logins detected by your 3D Sensors with RNA. You do **not** need to manage RUA with a policy.

Keysight, EX-1004
Page 1486

Setting up 3D Sensors and RUA

Platform: DC

Component: RUA

If you choose to use the RUA component on 3D Sensors instead of (or in addition to) the Sourcefire RUA Agent on your Active Directory server, you must first decide which sensors you are going to use. Then, you must create RUA detection engines on those sensors. For more information, see Understanding 3D Sensors and RUA on page 1459.

IMPORTANT! RUA detection engines on 3D Sensors detect POP3, AIM, SMTP, and LDAP logins. RNA detection engines on 3D Sensors with RNA automatically collect Oracle and SIP user login information from the monitored networks specified in your RNA detection policy; see What is an RNA Detection Policy? on page 889.

You can create an RUA detection engine on any 3D Sensor (except 3D9800 and Crossbeam-based software sensors, which do not support RUA) if you have an available interface set and detection resource; an RUA detection engine uses one detection resource.

You can create as many RUA detection engines on one 3D Sensor as you need to monitor the interface sets that you have configured. In general, however, you should only create multiple RUA detection engines if you need to monitor multiple interface sets. Keep in mind that you can monitor the same interface set with multiple detection engines. For example, if you have a passive interface set that includes all of the interfaces on the 3D Sensor, you can monitor that set with IPS, RNA, and RUA. For general information on detection engines and interface sets, see Using Detection Engines and Interface Sets on page 198.

Note that the number of RUA detection engines you can create on a 3D Sensor is limited by the sensor model and by how many other detection engines of different types you have already created. For example, on the 3D500, you can only create one RUA detection engine. For more information, see Understanding Detection Resources and 3D Sensor Models on page 202.

To create an RUA detection engine:

Access: Admin

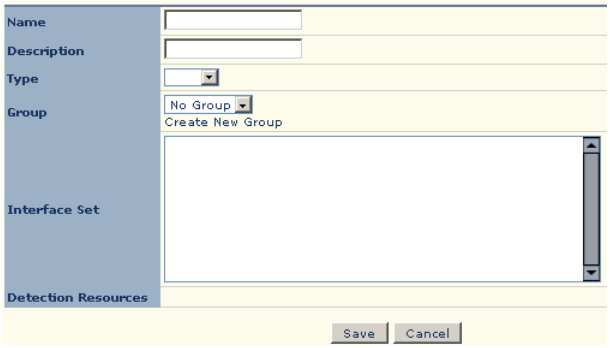
1. On the Defense Center, select **Operations > Configuration > Detection Engines > Detection Engines**.

   The Detection Engines page appears.

Keysight, EX-1004
Page 1487

2. Click **Create Detection Engine**.

   The Create Detection Engine page appears.



3. In the **Name** and **Description** fields, enter a name and description for the new detection engine.

   You can use alphanumeric characters, punctuation, and spaces.

4. From the **Type** drop-down list, specify that you want to create a **RUA** detection engine.

5. Optionally, add the detection engine to an existing detection engine group.

   See Using Detection Engine Groups on page 211 for information on creating and modifying detection engine groups.

6. Select the interface set that you want to assign to this detection engine.

   You can monitor any available interface set. See Using Interface Sets on page 222 for more information about creating and modifying interface sets.

IMPORTANT! On a 3D3800 or 3D5800 sensor, if you plan to use RNA to monitor either an inline or inline with fail open interface set, you **must** either configure an IPS detection engine that uses that interface set, as well as apply an intrusion policy to that detection engine, or configure the interface set in tap mode. Otherwise, the RNA detection engine monitoring that interface set will not see any traffic. If you are monitoring the same inline interface set with both IPS and RNA or RUA, and the IPS detection engine fails for any reason, the RNA or RUA detection engine monitoring that interface set will not see any traffic until the IPS detection engine restarts. Neither RNA nor RUA are supported on the 3D9800 sensor.

---

7. Specify that you want to use one detection resource for this detection engine.

   RUA detection engines only use one detection resource.

8. Click **Save**.

   The detection engine is created.

Keysight, EX-1004
Page 1488

9. Repeat steps 2 through 8 to create additional RUA detection engines on additional 3D Sensors.

   You are finished configuring RUA. If you have installed an RUA feature license on your Defense Center, you can begin viewing and analyzing RUA data, including Oracle and SIP user logins detected by your 3D Sensors with RNA.

Working with RUA Users

Platform: DC

Component: RUA

If you have installed an RUA feature license, when either an RUA Agent or a 3D Sensor detects a user login for a user who is not already in the database, an RUA user is added to the Defense Center user database, unless you have specifically restricted that login type (see Restricting RUA Logging on page 1476).

IMPORTANT! Although RUA detects SMTP logins, the Defense Center does not record them unless there is already a user with a matching email address in the database; RUA users are **not** added to the database based on SMTP logins.

The type of login that RUA detected determines what information is stored about the new user, as described in the following table.

RUA Login Types and User Data Stored

Login Type	User Data Stored
LDAP AIM Oracle SIP	• username • current IP address • login type (`aim`, `ldap`, `oracle`, or `sip`)
POP3 IMAP	• username • current IP address • email address • login type (`pop3` or `imap`)

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 can take five to ten minutes for the Defense Center database to update with user metadata after RUA detects a

Keysight, EX-1004
Page 1489

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 RUA license; see Licensing RUA on page 1475. Note that AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that RUA obtains from LDAP servers.

You can search, view, and delete users from the database; you can also purge all RUA users from the database. For more information, see the following sections:

- Viewing RUA Users on page 1490
- Understanding the RUA Users Table on page 1493
- Understanding User Details and Host History on page 1495
- Searching for RUA Users on page 1497

Viewing RUA Users

Platform: DC

Component: RUA

You can view a table of RUA users, and then manipulate the event view depending on the information you are looking for.

The page you see when you access RUA 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 1618.

Keysight, EX-1004
Page 1490

The RUA Users Table Functions table below describes some of the specific actions you can perform on an RUA users workflow page.

RUA Users Table Functions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the RUA Users Table on page 1493.
view a host's host profile	click the host profile icon ( ) that appears next to the IP address. For more information, see Using Host Profiles on page 944.
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 1495.
sort user data	click the column title. Click the column title again to reverse the sort order.
constrain the columns that appear	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, use the Expand arrow ( ) to expand the search constraints, then click the column name under Disabled Columns.
drill down to the next page in the workflow, constraining on a specific value	use one of the following methods: • 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.

Keysight, EX-1004
Page 1491

RUA Users Table Functions (Continued)

To...	You can...
delete RUA users from the system	use one of the following methods: • To delete some RUA users, select the check boxes next to the users you want to delete, then click **Delete**. • To delete all RUA users in the current constrained view, click **Delete All**, then confirm you want to delete all the users. RUA users remain deleted until they are detected again. See Purging the RNA and RUA Databases on page 2014 for information on deleting all RUA users from the database.
navigate within and between workflow pages	find more information in Using Workflow Pages on page 1593.
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 1614.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for RUA users	click **Search**. For more information, see Searching for RUA Users on page 1497.

Keysight, EX-1004
Page 1492

To view RUA users:

Access: Any Analyst/Admin

▶ Select **Analysis & Reporting > RUA > Users**.

The first page of the default RUA users workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

TIP! If you are using a custom workflow that does not include the table view of RUA users, from the **Workflows** menu on the toolbar, select **Users**.

Understanding the RUA Users Table

Platform: DC

Component: RUA

When RUA discovers a user, it collects data about that user and stores it in the Defense Center database. The fields in the RUA users table are described in the RUA Users Fields table.

RUA Users Fields

Field	Description
User	One of: • the first name, last name, and username of the RUA 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 RUA 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).
Username	The username of the RUA user, collected from the login event.

Keysight, EX-1004
Page 1493

RUA Users Fields (Continued)

Field	Description
Current IP	The IP address of the host that the RUA user is logged into. This field is blank if another user logs into the host. (RUA associates the last user that logged in with the host.) Note that because unsuccessful AIM login attempts are recorded, the Defense Center may indicate that an AIM user is logged in, even if that information is incorrect.
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

Keysight, EX-1004
Page 1494

RUA Users Fields (Continued)

Field	Description
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 RUA user. For example, for users added to the database when RUA 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

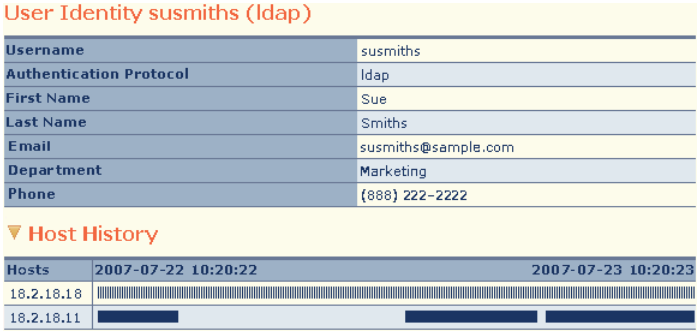
Platform: DC

Component: RUA

From any event view that associates user identity data with other kinds of events, as well as from a table view of RUA users, you can display the User Identify

Keysight, EX-1004
Page 1495

pop-up window to learn more about a specific user. User identity information also appears in the terminating page for RUA users workflows.



The user identity data you see is the same as you would see in the table view of users; for more information, see the RUA Users Fields table on page 1493.

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 approximates login and logout times with bar graphs. A typical user might log on to 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.

The data used to generate the host history is stored in the RUA 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 1747.

To view RUA user details and host history:

Access: Any Analyst/Admin

▶ You have two options:

- In any event view that lists RUA users, click the user icon ( )that appears next to a user identity.
- In any RUA users workflow, click the Users terminating page.

User details appear.

Keysight, EX-1004
Page 1496

Searching for RUA Users

Platform: DC

Component: RUA

You can search for specific RUA users. You may want to create searches customized for your network environment, then save them to re-use later. The search criteria you can use are described in the RUA Users Search Criteria table.

RUA Users Search Criteria

Field	Search Criteria Rules
User	Enter all or part of the user identity you want to search for. The user is one of: • the first name, last name, and username of the RUA user as collected via the optional Defense Center-LDAP server connections • the username only, if you have not configured a Defense Center-LDAP server connection, or for users that the Defense Center cannot correlate with an LDAP record You can specify multiple user identities in a comma-separated list. This field is case-insensitive.
Username	Enter the full username of the user you want to search for. This search field is case-insensitive. You can use an asterisk (*) as a wildcard character in this field.
First Name	Enter the full first name of the user you want to search for. This search field is case-insensitive. You can use an asterisk (*) as a wildcard character in this field.
Last Name	Enter the full last name of the user you want to search for. This search field is case-insensitive. You can use an asterisk (*) as a wildcard character in this field.
E-Mail	Enter the full email address of the user you want to search for. This search field is case-insensitive. You can use an asterisk (*) as a wildcard character in this field.
Department	Enter the full department name of the user you want to search for. This search field is case-insensitive. You can use an asterisk (*) as a wildcard character in this field.

Keysight, EX-1004
Page 1497

RUA Users Search Criteria (Continued)

Field	Search Criteria Rules
Phone	Enter the telephone number of the user you want to search for. Note that you must enter the telephone number exactly as it is stored on your LDAP servers and thus in the Defense Center database. For example, if you search for 800-555-1212 and the number is stored as (800) 555-1212, you will not get any results. You can, however, use an asterisk (*) as a wildcard character in this field. For example, you could search for *800*555*1212.
User Type	Enter the protocol used to detect the user you want to search for. Valid search criteria are `ldap`, `pop3`, `imap`, and `aim`; because RUA users are not added to the database based on SMTP logins, entering `smtp` will not return any results.

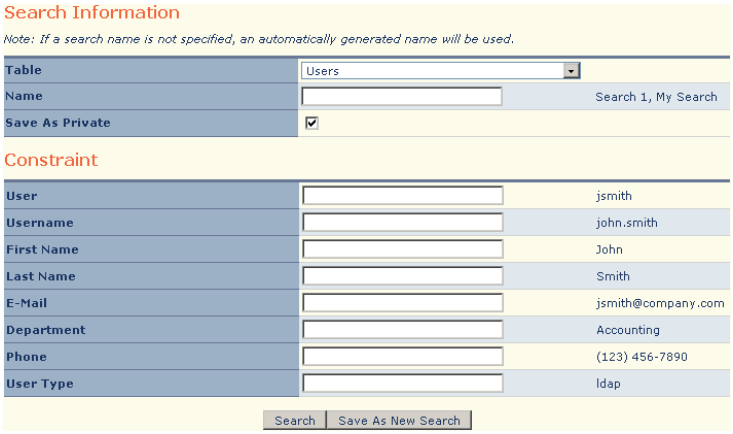
For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

To search for RUA users:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Searches > 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.

Keysight, EX-1004
Page 1498

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 created automatically when you save the search.

3. Enter your search criteria in the appropriate fields, as described in the RUA Users Search Criteria table on page 1497. If you enter multiple criteria, the search returns only the records that match all the criteria.

   TIP! 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

4. You have the following options:
   - Click **Search** to start the search.

     Your search results appear in the default RUA users workflow. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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.

Keysight, EX-1004
Page 1499

Working with RUA Events

Platform: DC

Component: RUA

If you have installed an RUA feature license on your Defense Center, RUA generates events that communicate the details of user activity on your network. There are four types of RUA events, as described in the RUA Event Types table.

RUA Event Types

Event Name	Description
New User Identity	This event is generated when RUA 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 RUA Agent that you installed on an Active Directory server detects an LDAP login • an RUA detection engine on a 3D Sensor detects an LDAP, POP3, IMAP, SMTP, or AIM login • an RNA detection engine on a 3D Sensor detects an 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. • All AIM login attempts are recorded, even if the login is unsuccessful. • A user login is not recorded if you have specifically restricted its login type; see Restricting RUA Logging on page 1476.
Delete User Identity	This event is generated when you manually delete a user from the Defense Center database.
User Identity Dropped: User Limit Reached	This event is generated when RUA 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 RUA feature license.

When an RUA event is generated, it is logged to the database. You can view, search, and delete RUA events; you can also purge all RUA events from the database.

Whenever possible the Sourcefire 3D System correlates RUA events with other types of events. For example, intrusion events can tell you the users who were

Keysight, EX-1004
Page 1500

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 RUA events in compliance rules. Based on the type of RUA event generated as well as other criteria that you specify, you can build compliance rules that, when used in a compliance policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria. For more information on how you can use RUA events, see How Can You Use RUA? on page 1468.

For more information, see the following sections:

- Viewing RUA Events on page 1501
- Understanding the RUA Events Table on page 1505
- Searching for RUA Events on page 1506

Viewing RUA Events

Platform: DC

Component: RUA

You can view a table of RUA events, and then manipulate the event view depending on the information you are looking for.

The page you see when you access RUA events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of RUA events 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 1618.

Keysight, EX-1004
Page 1501

The RUA Events Table Functions table below describes some of the specific actions you can perform on an RUA users workflow page.

RUA Events Table Functions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the RUA Events Table on page 1505.
modify the time and date range for displayed events	click the time range link. For more information, see Setting Event Time Constraints on page 1597. 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.
view a host's host profile	click the host profile icon ( ) that appears next to the IP address. For more information, see Using Host Profiles on page 944.
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 1495.
sort user data	click the column title. Click the column title again to reverse the sort order.
constrain the columns that appears	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, use the Expand arrow ( ) to expand the search constraints, then click the column name under Disabled Columns.

Keysight, EX-1004
Page 1502

RUA Events Table Functions (Continued)

To...	You can...
drill down to the next page in the workflow, constraining on a specific value	use one of the following methods: • 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 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
delete RUA events from the system	use one of the following methods: • To delete some RUA events, select the check boxes next to the users you want to delete, then click **Delete**. • To delete all RUA events in the current constrained view, click **Delete All**, then confirm you want to delete all the users. See Purging the RNA and RUA Databases on page 2014 for information on deleting all RUA events from the database.
navigate within and between workflow pages	find more information in Using Workflow Pages on page 1593.
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 1614.

Keysight, EX-1004
Page 1503

RUA Events Table Functions (Continued)

To...	You can...
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.
search for RUA events	click **Search**. For more information, see Searching for RUA Events on page 1506.

To view RUA events:

Access: Any Analyst/Admin

- Select **Analysis & Reporting > RUA > RUA Events**.

  The first page of the default RUA events workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range; see Setting Event Time Constraints on page 1597.

  TIP! If you are using a custom workflow that does not include the table view of RUA events, from the **Workflows** menu on the toolbar, select **RUA Events**.

Keysight, EX-1004
Page 1504

Understanding the RUA Events Table

Platform: DC

Component: RUA

When RUA generates an event, it is logged to the database. The fields in the RUA users table are described in the RUA Events Fields table.

RUA Events Fields

Field	Description
Time	The time that RUA generated the event.
Event	The event type. For more information, see the RUA Event Types table on page 1500.
User	The user associated with the RUA event. At a minimum, this field contains a username and the protocol used to detect the RUA user. If there is LDAP metadata on the user, this field can also contain the first name and last name of the user.
User Type	The protocol used to detect the RUA user. For example, for users added to the database when RUA detects a POP3 login, the user type is `pop3`.
IP Address	For User Login events, 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). For other types of RUA events, this field is blank.
Description	For Delete User Identity and User Identity Dropped events, the username of the user who was deleted from the database or failed to be added to the database. For other types of RUA events, this field is blank.
Detection Engine	For events generated by an RUA detection engine on a 3D Sensor, the name of the detection engine and the sensor it resides on. For other types of `RUA` events, `Local Detection Engine`.
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.

Keysight, EX-1004
Page 1505

Searching for RUA Events

Platform: DC

Component: RUA

You can search for specific RUA users. You may want to create searches customized for your network environment, then save them to re-use later. The search criteria you can use are described in the RUA Events Search Criteria table.

RUA Events Search Criteria

Field	Search Criteria Rules
Event	Enter all or part of the event type you want to search for. For more information, see the RUA Event Types table on page 1500.
IP Address	Specify the IP address of the host for which you want to see User Login events. See Specifying IP Addresses in Searches on page 1557 for detailed information about how to enter IP addresses.
User	Enter all or part of the user identity you want to search for. The user is one of: • the first name, last name, and username of the RUA user as collected via the optional Defense Center-LDAP server connections • the username only, if you have not configured a Defense Center-LDAP server connection, or for users that the Defense Center cannot correlate with an LDAP record You can specify multiple user identities in a comma-separated list. This field is case-insensitive.
Description	Enter the username of a user who was deleted from the database or failed to be added to the database.

Keysight, EX-1004
Page 1506

RUA Events Search Criteria (Continued)

Field	Search Criteria Rules
Detection Engine	Enter one of: • the name of the RUA detection engine that generated the event or the name of the sensor where that detection engine resides, to search for events generated by that detection engine or by any one of the RUA detection engines on a 3D Sensor • `Local Detection Engine`, to search for all other types of RUA events See Specifying Detection Engine/Appliance Names on page 1559 for detailed information about how to enter detection engine names.
User Type	Enter the protocol used to detect the user you want to search for. Valid search criteria are `ldap`, `pop3`, `imap`, and `aim`; because RUA users are not added to the database based on SMTP logins, entering `smtp` will not return any results.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

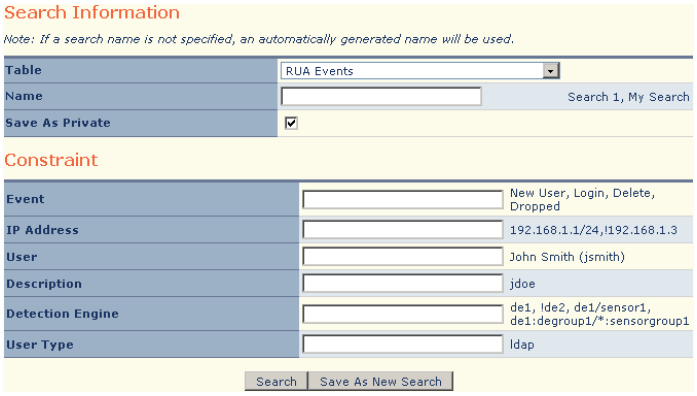
Keysight, EX-1004
Page 1507

To search for RUA events:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Searches > Events**.

   The RUA Events search page appears.



   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 created automatically when you save the search.

3. Enter your search criteria in the appropriate fields, as described in the RUA Events Search Criteria table on page 1506. If you enter multiple criteria, the search returns only the records that match all the criteria.

   TIP! 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 save a search as a restriction for restricted event analyst users, you **must** save it as a private search.

4. You have the following options:

   - Click **Search** to start the search.

     Your search results appear in the default RUA events workflow, constrained by the current time range. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

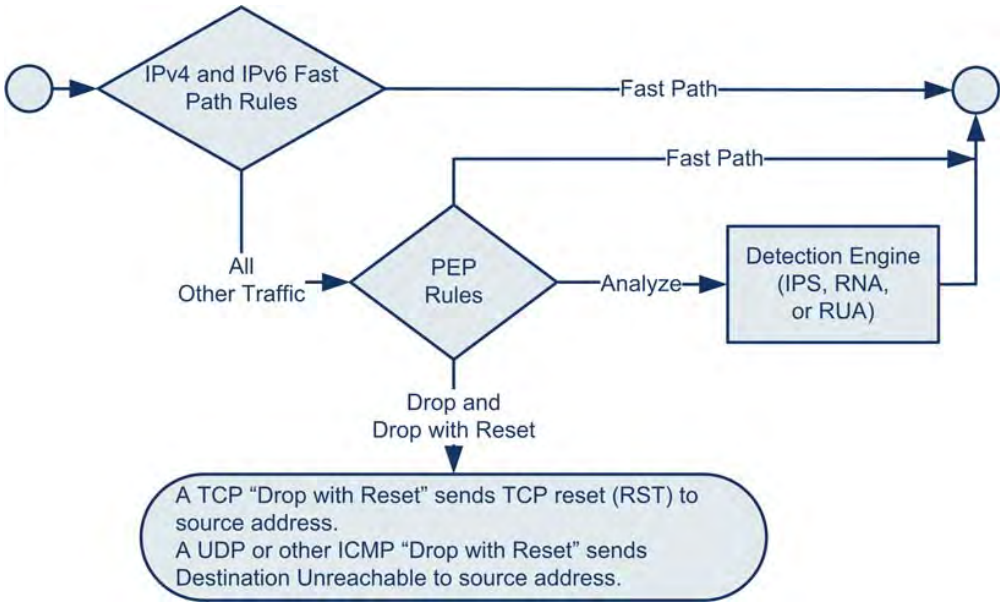
Keysight, EX-1004
Page 1508

- 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.

Keysight, EX-1004
Page 1509

Chapter 33
Using PEP to Manage Traffic

PEP is a technology based on the hardware capabilities of the Sourcefire 3D9900 and Series 3 sensors. PEP allows you to create rules to block, analyze, or send traffic directly through these sensors with no further inspection, as shown below.



PEP traffic management enhances your Sourcefire 3D Sensor's efficiency by allowing you to pre-select traffic to bypass or to drop instead of analyzing.

Keysight, EX-1004
Page 1510

The following sections describe PEP traffic management and how you can use it in your Sourcefire 3D System deployment:

- Understanding PEP Traffic Management on page 1511 explains the fundamentals of PEP policies and provides examples of their use.
- Creating PEP Policies on page 1513 explains how to use PEP traffic filters and PEP rules to create PEP policies.
- Working with PEP Policies on page 1525 explains how to edit and apply PEP policies.
- Viewing Active PEP Policy Details on page 1529 explains how to view the details of PEP policies applied to interface sets.

Understanding PEP Traffic Management

PEP is a rule-driven detection technology that uses two stages to manage traffic:

- The first stage uses fast path rules; see Fast Path Rule Fundamentals on page 1511.
- The second stage uses PEP rules; see PEP Rule Fundamentals on page 1511.

If you understand the characteristics of the PEP traffic management stages, you can better determine how to use them. For PEP traffic management examples, see Using PEP Traffic Management on page 1512.

Fast Path Rule Fundamentals

Fast path rules divert traffic that does not need to be analyzed to bypass the sensor. Fast path rules either send traffic to the fast path (out of the interface) or allow it to continue into the 3D Sensor for further analysis. They use 3D9900 and Series 3 hardware capabilities. 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. These rules, therefore, are restricted. For example, you can specify individual ports in a fast path rule, but you cannot use a range of ports.

PEP Rule Fundamentals

PEP rule functions are more complex than fast path rules. For example, you can configure additional rule actions in PEP rules. The disposition or action taken because of a PEP rule offers more options than the one-of-two paths determination a fast path rule makes. You can set a PEP rule to:

- fast path all traffic
- drop all traffic
- drop all traffic with reset

Keysight, EX-1004
Page 1511

- analyze all traffic
- analyze or fast path traffic based on the detection type of the traffic

Using PEP Traffic Management

The examples that follow describe situations in which you could use PEP to more efficiently manage the traffic passing through an appliance.

The Example1 company performs a nightly backup to high-speed Network Attached Storage (NAS) from several critical systems. The backup software uses a specific set of TCP and UDP ports. To ensure that backups are fast, the interconnecting network segments between the critical systems and NAS are all 1GB fiber. In this situation, you can use fast path rules to send traffic directly to and from the critical systems and NAS. In these fast path rules you specify the critical system and NAS network segments as well as the specific backup ports.

The Example2 company has a huge campus and provides Voice over IP (VoIP) services to its employees. They recently replaced several devices in their network that introduced latency and rendered their VoIP system useless. This is a candidate for fast path rules to send traffic directly to and from VoIP system network segments. Because each of the protocols used in VoIP have one or more well-known port numbers associated with them, you could also use specific ports in your fast path rules.

To contrast with the preceding examples, consider an example where PEP rules are more appropriate. The Example3 company policy inspects all standard text communications, transferred files, and web browsing on specific network segments. On those segments they do not permit encrypted text communications (SSH on port 22), encrypted file transfers (FTPS on ports 889 and 890), or encrypted web browsing (HTTPS on port 443). Example3 only permits unencrypted FTP, telnet, and HTTP protocols. Due to the probable payloads, those types of traffic, as well as other traffic, are analyzed.

Employees at Example3 routinely use web browsers and FTP clients to access internal servers as well as conduct business with outside contractors. This is a candidate for PEP rule sets that drop traffic using ports 22, 889, 890, and 443, fast path all other RNA or RUA traffic, and analyze all other IPS traffic. For traffic from external networks, you could use a simple drop action. For traffic from internal networks, a drop with reset is a better choice so that Example3 employees can more quickly use their web browsers and FTP clients again after they accidently use secured protocols.

Keysight, EX-1004
Page 1512

Creating PEP Policies

Platform: DC/MDC

Component: None

You can create PEP policies from the PEP Policy Management page.

TIP! Instead of creating a new policy, you can export a PEP 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 1997.

To create a PEP policy:

Access: P&U Admin/Admin

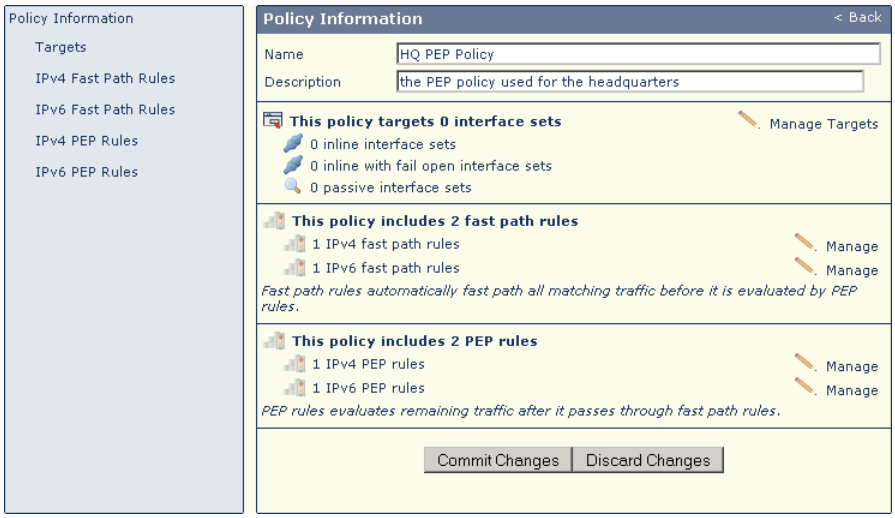
1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

2. Click **Create Policy**.

   When another policy has unsaved changes, you are prompted whether to continue. Click **Cancel** to abandon the new policy and return to the Policy Management page, or click **OK** to discard changes to the other policy and continue.

   The Policy Information page appears.



3. Enter a name for your new policy in the **Name** field to uniquely identify the policy.

4. Optionally, enter a description in the **Description** field.

Keysight, EX-1004
Page 1513

5. You can then:
   - select your target interface set. For more information, see Selecting and Changing Target PEP Interface Sets on page 1514.

   TIP! If you do not select and associate a valid target interface set with your PEP policy, each time you create or edit that policy and commit the changes, a confirmation dialog box reminds you that you need to add interface sets.

   - build your IPv4 fast path rules. For more information, see Building and Editing IPv4 Fast Path Rules on page 1516.
   - build your IPv6 fast path rules. For more information, see Building and Editing IPv6 Fast Path Rules on page 1518.
   - build your IPv4 PEP rules. For more information, see Building and Editing IPv4 PEP Rules on page 1520.
   - build your IPv6 PEP rules. For more information, see Building and Editing IPv6 PEP Rules on page 1522.
6. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.
7. After you create your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

   Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

   When you save your changes, the new PEP policy is available but does not go into effect until you apply it.

See the following for more information:

- Using PEP to Manage Traffic on page 1510.
- Editing PEP Policies on page 1526.
- Applying PEP Policies on page 1527.
- Exporting a PEP Policy on page 2002.
- Deleting PEP Policies on page 1529.
- Viewing Active PEP Policy Details on page 1529.

Selecting and Changing Target PEP Interface Sets

Platform: DC/MDC

Component: None

You can manually apply a PEP policy to interface sets, groups of interface sets, or sensors regardless of whether you specify targets within the policy. However, you can also target specific interface sets, groups of interface sets, or sensors so

Keysight, EX-1004
Page 1514

you can routinely apply the policy to your specified targets without having to identify the targets each time.

IMPORTANT! You can only apply a PEP policy to interface sets and sensors capable of enforcing a PEP policy. When you manually apply a PEP policy, the list only contains those interface sets and sensors.

You can specify any of the following targets in your policy:

- one or more interface sets
- one or more interface set groups
- all of the interface sets on one or more sensors

To select or change a target PEP interface set:

Access: P&R Admin/Admin

1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

2. Click **Create Policy** to create a new policy or click **Edit** in the PEP Policy Management page to edit an existing policy.

   The Policy Information page appears.

3. Click **Interface Sets** in the navigation panel.

   The Interface Sets view appears. Any existing targeted sensors, interface set groups, or interface sets are listed.

4. To add interface sets click the appropriate **Add** for sensors, interface set groups, or interface sets.

   A pop-up window with a selection list appears.

   IMPORTANT! If you select a sensor, all interface sets on that sensor become the policy target. If you select an interface set group, all interfaces in that group that are PEP-capable become the policy target.

5. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.

6. After you modify your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

   Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

   When you save your changes, the changes to the PEP policy are available but do not go into effect until you apply it. For more information, see Applying PEP Policies on page 1527.

Keysight, EX-1004
Page 1515

Building and Editing IPv4 Fast Path Rules

Platform: DC/MDC

Component: None

Fast path rules send traffic to the fast path (out of the interface) or into the 3D Sensor 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 (select the specific protocol number associated with the protocol in `/etc/protocols`, type a protocol number, or select **All** to match traffic from any protocol on the list)
- initiator or responder port, for TCP or UDP protocols
- VLAN ID
- bidirectional option

Note that the outermost ID is used for fast path rules on the 3D9900 and Series 3 appliances.

To build or edit IPv4 fast path rules:

Access: P&R Admin/Admin

1. Select **Policy & Response > PEP > Policy Management**.

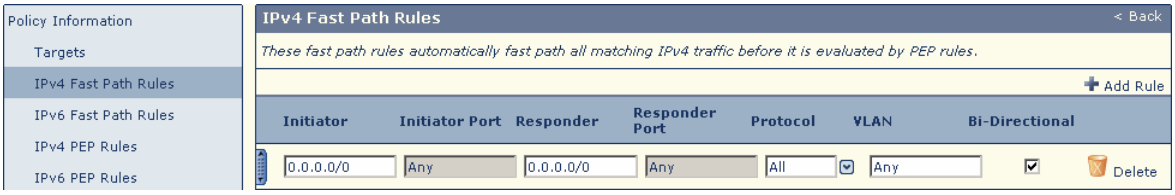
   The PEP Policy Management page appears.

2. Click **Create Policy** to create a new policy or click **Edit** in the PEP Policy Management page to edit an existing policy.

   The Policy Information page appears.

3. Click **IPv4 Fast Path Rules** in the navigation panel.

   The IPv4 Fast Path Rules view appears, listing any existing IPv4 fast path rules for the PEP policy.



4. Click **Add Rule** to add a fast path rule.

   The default rule appears.

WARNING! The default rule permits all traffic to bypass further analysis.

Keysight, EX-1004
Page 1516

5. 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 59.

6. Optionally, click the **Bi-Directional** field.

   The **Bi-Directional** option is selected as the default. When the **Bi-Directional** option is selected, you filter all traffic traveling between the specified initiator and responder IP addresses. When unselected you filter just traffic from the specified initiator IP address to the specified responder IP address.

7. Optionally, set the protocol on which you want the rule to act. You have several options:
   - Select the protocol from the Protocol drop-down list.
   - Type the protocol name into the Protocol field and select a match from the drop-down list.
   - Type the protocol number for the protocol into the Protocol field and press Enter or Tab.

     Note that when you type a protocol number for a protocol that is not in `/etc/protocols`, `Protocol (`*`protocol_number`*`)` populates the Protocol field.

8. Optionally, if you chose the TCP or UDP protocol in step 7, you can enter initiator and responder ports in the **Initiator Port** and the **Responder Port** fields to designate ports.

   TIP! You can enter a list of comma-separated port numbers. You cannot use port ranges in IPv4 fast path rules.

   If the bi-directional option is selected, your filter criteria is narrowed to packets from those initiator ports or packets to those responder ports.

9. Optionally, you can enter a VLAN ID in the **VLAN** field.

   Your rule matches only traffic for that VLAN.

10. To add another fast path rule, click **Add Rule** and return to step 5.

    TIP! The rule order makes no difference in the IPv4 fast path rule list.

11. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.

Keysight, EX-1004
Page 1517

12. After you modify your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

    Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

    When you save your changes, the changes to the PEP policy are available but do not go into effect until you apply it. For more information, see Applying PEP Policies on page 1527.

Building and Editing IPv6 Fast Path Rules

Platform: DC/MDC

Component: None

Fast path rules send traffic to the fast path (out of the interface) or into the 3D Sensor 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 (select the specific protocol number associated with the protocol in `/etc/protocols`, type a protocol number, or select **All** to match traffic from any protocol on the list)
- 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 on the 3D9900 and Series 3 appliances.

To build or edit IPv6 fast path rules:

Access: P&R Admin/Admin

1. Select **Policy & Response > PEP > Policy Management**.

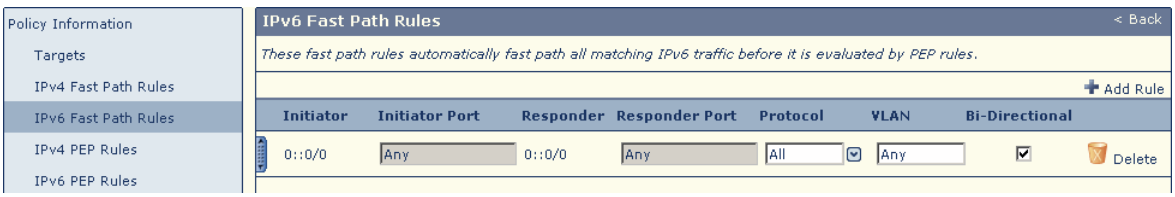
   The PEP Policy Management page appears.

2. Click **Create Policy** to create a new policy or click **Edit** in the PEP Policy Management page to edit an existing policy.

   The Policy Information page appears.

3. Click **IPv6 Fast Path Rules** in the navigation panel.

   The IPv6 Fast Path Rules view appears, listing any existing IPv6 fast path rules.



Keysight, EX-1004
Page 1518

4. Click **Add Rule** to add a fast path rule.

   The default rule appears. Note that the initiator and responder fields are fixed and indicate that the filter applies to IPv6 packets from any initiator or responder.

   WARNING! The default rule permits all traffic to bypass further analysis.

5. Optionally, set the protocol on which you want the rule to act. You have several options:
   - Select the protocol from the Protocol drop-down list.
   - Type the protocol name into the Protocol field and select a match from the drop-down list.
   - Type the protocol number for the protocol into the Protocol field and press Enter or Tab.

     Note that when you type a protocol number for a protocol that is not in `/etc/protocols`, `Protocol (`*`protocol_number`*`)` populates the Protocol field.

   Your fast path rule matches only the selected protocol's packets.

6. Optionally, click the **Bi-Directional** field. The **Bi-Directional** option is selected as the default.

   When the **Bi-Directional** option is not selected, 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 bi-directional option only functions with initiator and responder ports.

7. Optionally, if you chose the TCP or UDP protocol in step 5, you can enter initiator and responder ports in the **Initiator Port** and the **Responder Port** fields to designate ports.

   TIP! You can enter a list of comma-separated port numbers. You cannot use port ranges in IPv6 fast path rules.

8. Optionally, you can enter a VLAN ID in the **VLAN** field.

   Your rule matches only traffic for that VLAN.

9. To add another fast path rule, click **Add Rule**. Return to step 5.

   TIP! Although you add new rules to the top of the list, the rule order makes no difference in the IPv6 fast path rule list.

Keysight, EX-1004
Page 1519

10. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.
11. After you modify your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

    Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

    When you save your changes, the changes to the PEP policy are available but do not go into effect until you apply it. For more information, see Applying PEP Policies on page 1527.

Building and Editing IPv4 PEP Rules

Platform: DC/MDC

Component: None

PEP rules evaluate traffic after it passes through packet filters. PEP rules can send traffic to the fast path (out of the interface) or into the 3D Sensor for further analysis. PEP rules can also drop traffic or drop traffic and return a reset to the initiator address. You can build IPv4 PEP rules based on:

- initiator or responder IP address or CIDR block
- protocol (select the specific protocol number associated with the protocol in `/etc/protocols`, type a protocol number, or select **All** to match traffic from any protocol on the list)
- initiator or responder port, for TCP or UDP protocols
- VLAN ID
- action

You can set a DE-specific action to process traffic detected by different detection engine types differently. For example, you can create a rule that analyzes IPS traffic but sends RNA and RUA traffic to the fast path.

Note that the outermost VLAN ID is used for PEP rules on the 3D9900 and Series 3 appliances.

PEP rules run in a hierarchal order. You list PEP rules in priority order. If you have specific traffic that you do not want to analyze or you want to drop first, put rules for that traffic higher in your list. The default is to send all remaining traffic on for analysis.

To build or edit IPv4 PEP rules:

Access: P&R Admin/Admin

1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

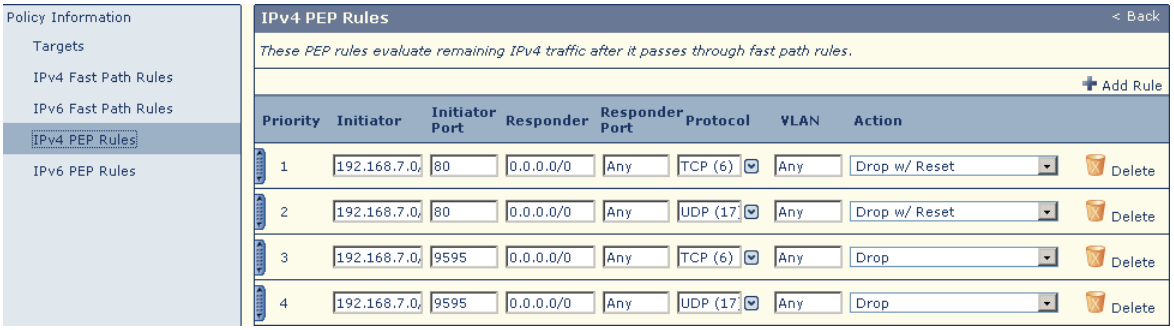
Keysight, EX-1004
Page 1520

2. Click **Create Policy** to create a new policy or click **Edit** in the PEP Policy Management page to edit an existing policy.

   The Policy Information page appears.

3. Click **IPv4 PEP Rules** in the navigation panel.

   The IPv4 PEP Rules view appears and any existing IPv4 PEP rules for the PEP policy are listed.



4. Click **Add Rule** to add a rule to your PEP rules.

   The default rule appears.

   TIP! The default rule sends all traffic on to further analysis.

5. Use CIDR notation in the **Initiator** and the **Responder** fields to designate the IP address ranges of initiators or responders that you want to use in the rule.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

6. Optionally, set the protocol on which you want the rule to act. You have several options:

   - Select the protocol from the Protocol drop-down list.
   - Type the protocol name into the Protocol field and select a match from the drop-down list.
   - Type the protocol number for the protocol into the Protocol field and press Enter or Tab.

     Note that when you type a protocol number for a protocol that is not in `/etc/protocols`, `Protocol (`*`protocol_number`*`)`populates the Protocol field.

Keysight, EX-1004
Page 1521

7. Optionally, if you chose the TCP or UDP protocol, you can enter initiator and responder ports in the **Initiator Port** and the **Responder Port** fields to designate ports on which you want the rule to act.

   TIP! You can enter individual port numbers or a list of comma-separated port numbers or port ranges. Sourcefire **strongly** recommends that you limit the number of ports and use port ranges sparingly. Port ranges increase rule complexity and if your rule is overly complex, you cannot apply it.

8. Optionally, you can enter a VLAN ID in the **VLAN** field.

   Your rule is narrowed to that VLAN.

9. You have several options for the action to carry out against flows matching the rule:
   - Select **Analyze** to analyze the traffic.
   - Select **Drop** to drop the traffic.
   - Select **Drop w/ Reset** to drop the traffic and send a
   - Select **Fast Path** to divert the traffic to the fast path.
   - Select **DE-Specific** and then select **Fast Path** for each type of traffic (**IPS**, **RNA**, or **RUA**) you want to divert to the fast path.
   - Select **DE-Specific** and then select **Analyze** for each type of traffic (**IPS**, **RNA**, or **RUA**) you want to analyze.

10. To add another rule to your IPv4 PEP rules, click **Add Rule** and return to step 5.

    Arrange the PEP rules in priority order by clicking and dragging them.

11. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.

12. After you modify your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

    Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

    When you save your changes, the changes to the PEP policy are available but do not go into effect until you apply it. For more information, see Applying PEP Policies on page 1527.

Building and Editing IPv6 PEP Rules

Platform: DC/MDC

Component: None

PEP rules evaluate traffic after it passes through fast path rules. PEP rules can send traffic to the fast path (out of the interface) or into the 3D Sensor for further

Keysight, EX-1004
Page 1522

analysis. PEP rules can also drop traffic or drop traffic and return a reset to the initiator address. You can build IPv6 PEP rules based on:

- initiator or responder IP address or CIDR block
- protocol (select the specific protocol number associated with the protocol in `/etc/protocols`, type a protocol number, or select **All** to match traffic from any protocol on the list)
- initiator or responder port, for TCP or UDP protocols
- VLAN ID
- action

You can set a DE-specific action to process traffic detected by different detection engine types differently. For example, you can create a rule that analyzes IPS traffic but sends RNA and RUA traffic to the fast path.

Note that the outermost VLAN ID used for PEP rules on the 3D9900 and Series 3 appliances. In addition, note that you should not apply a policy containing IPv6 rules to a non-Series 3 sensor.

PEP rules run in a hierarchal order. You list PEP rules in priority order. If you have specific traffic that you do not want to analyze or you want to drop first, put rules for that traffic higher in your list. The default is to sends all remaining traffic on for analysis.

Remember to create rules for all addressing schemes used on your network; for example, to match both external and internal traffic coming to a host you need to create a separate rule with a responder IP address range including the unicast address for the host and a separate rule with a responder IP address range including the link-local address for the host.

IMPORTANT! You cannot apply a policy created on an appliance running version 4.10 to an appliance running 4.9.1.x if it contains IPv6 PEP rules.

To build or edit IPv6 PEP rules:

Access: P&R Admin/Admin

1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

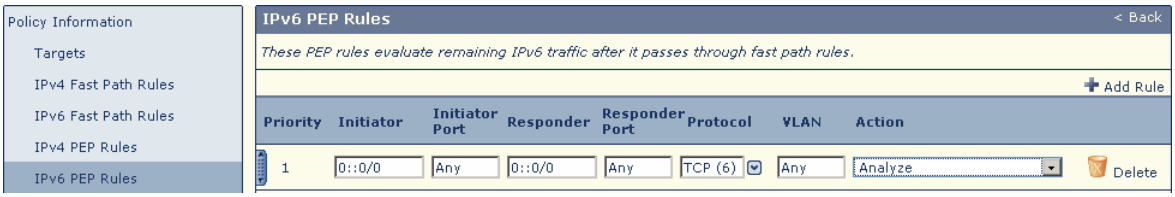
2. Click **Create Policy** to create a new policy or click **Edit** in the PEP Policy Management page to edit an existing policy.

   The Policy Information page appears.

3. Click **IPv6 PEP Rules** in the navigation panel.

Keysight, EX-1004
Page 1523

4. The IPv6 PEP Rules view appears, and any existing IPv6 PEP rules for the PEP policy are listed.



5. Click **Add Rule** to add a rule to your IPv6 PEP rules.

   The default rule appears.

6. If you are using a Series 3 sensor, use CIDR notation in the **Initiator** and the **Responder** fields to designate the IP address ranges of initiators or responders that you want to use in the rule.

   For information on using CIDR notation in the Sourcefire 3D System, see IP Address Conventions on page 59.

7. Optionally, set the protocol on which you want the rule to act. You have several options:

   - Select the protocol from the Protocol drop-down list.
   - Type the protocol name into the Protocol field and select a match from the drop-down list.
   - Type the protocol number for the protocol into the Protocol field and press Enter or Tab.

     Note that when you type a protocol number for a protocol that is not in `/etc/protocols`, `Protocol (`*`protocol_number`*`)` populates the Protocol field.

8. Optionally if you chose the TCP or UDP protocol in step 7, you can enter initiator and responder ports in the **Initiator Port** and the **Responder Port** fields to designate ports.

   TIP! You can enter individual port numbers or a list of comma-separated port numbers or port ranges. Sourcefire **strongly** recommends that you limit the number of ports and use port ranges sparingly. Port ranges increase rule complexity and if your rule is overly complex, you cannot apply it.

9. Optionally, you can enter a VLAN ID in the **VLAN** field.

   Your rule is narrowed to that VLAN.

10. To add another rule to your IPv6 PEP rules, click **Add Rule** and return to step 7.

    Arrange the PEP rules in priority order by clicking and dragging them.

11. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.

Keysight, EX-1004
Page 1524

12. After you modify your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

    Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

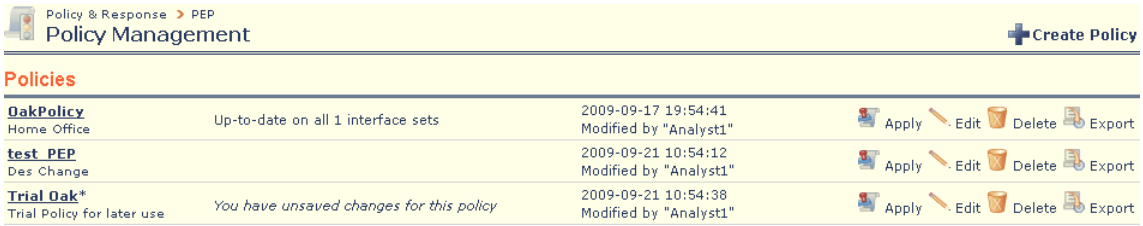
    When you save your changes, the changes to the PEP policy are available but do not go into effect until you apply it. For more information, see Applying PEP Policies on page 1527.

Working with PEP Policies

Platform: DC/MDC

Component: None

On the Policies page (**Policy & Response > PEP > Policy Management**) you can view all your PEP policies, listed by name with the time and date the policy was last modified and the user who modified it.



Options on this page also allow you to create a new policy, apply a policy, edit a policy, export a policy, and delete a policy. You can associate policies with individual interface sets, all interfaces in an interface set group, or with all interfaces on a 3D Sensor. After you create or edit a PEP policy, your settings take effect only after you apply the policy. The text style varies for descriptions of each of the following conditions:

- Plain black text indicates on how many targeted interface sets the most recently saved policy changes have been applied.
- Plain red text indicates on how many targeted interfaces sets the most recently saved policy changes have not been applied.
- *Italicized* black text indicates that the policy has unsaved changes.

See the following for more information about PEP policies:

- Understanding PEP Traffic Management on page 1511
- Creating PEP Policies on page 1513
- Editing PEP Policies on page 1526
- Applying PEP Policies on page 1527
- Exporting a PEP Policy on page 2002

Keysight, EX-1004
Page 1525

- Deleting PEP Policies on page 1529
- Viewing Active PEP Policy Details on page 1529

Editing PEP Policies

Platform: DC/MDC

Component: None

You can edit your PEP policies to meet changing network requirements.

To edit a PEP policy:

Access: P&R Admin/Admin

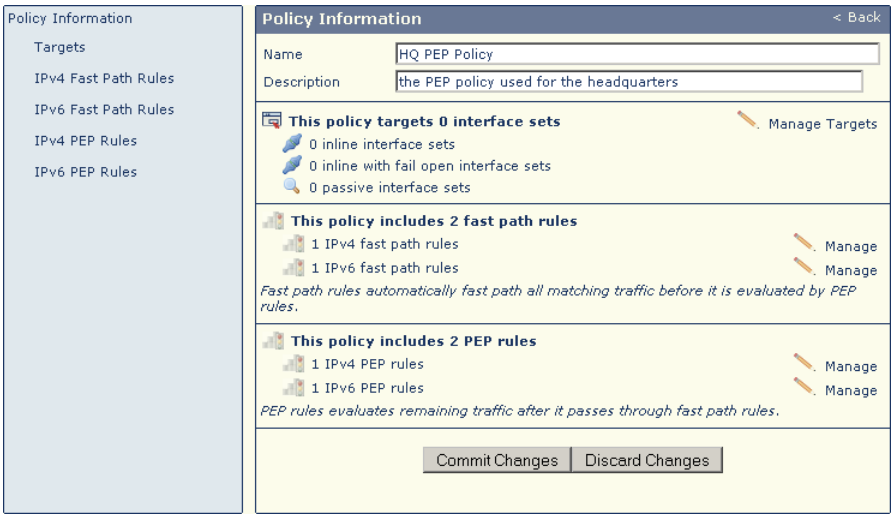
1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

2. Click **Edit** in the PEP Policy Management page to edit an existing policy.

   When another policy has unsaved changes, you are prompted whether to continue. Click **Cancel** to abandon the new policy and return to the PEP Policy Management page, or click **OK** to discard changes to the other policy and continue.

   The Policy Information page appears.



3. You can modify the name and description of the policy using the **Name** and **Description** fields.

Keysight, EX-1004
Page 1526

4. You can also:
   - change your target interface set. For more information, see Selecting and Changing Target PEP Interface Sets on page 1514.
   - edit your IPv4 fast path rules. For more information, see Building and Editing IPv4 Fast Path Rules on page 1516.
   - edit your IPv6 fast path rules. For more information, see Building and Editing IPv6 Fast Path Rules on page 1518.
   - edit your IPv4 PEP rules. For more information, see Building and Editing IPv4 PEP Rules on page 1520.
   - edit your IPv6 PEP rules. For more information, see Building and Editing IPv6 PEP Rules on page 1522.
5. Click the policy change icon ( ) in the navigation panel to return to the Policy Information page.
6. After you modify your policy, you can click **Commit Changes** to save your changes, click **Discard Changes** to discard any unsaved changes, or exit the policy without saving changes.

   Note that the system holds your unsaved changes in memory when you exit the policy or log out of the system.

   When you save your changes, the changes to the PEP policy are available but do not go into effect until you apply it.

See the following for more information about PEP policies:

- Understanding PEP Traffic Management on page 1511
- Creating PEP Policies on page 1513
- Applying PEP Policies on page 1527
- Exporting a PEP Policy on page 2002
- Deleting PEP Policies on page 1529
- Viewing Active PEP Policy Details on page 1529.

Applying PEP Policies

Platform: DC/MDC

Component: None

You can apply the policy to applicable and to targeted interface sets.

IMPORTANT! If you are applying a PEP policy to a clustered sensor pair and the software versions are not the same on the two sensors in that pair (for example, if an upgrade on one of the sensors failed), then the policy apply will not succeed.

Keysight, EX-1004
Page 1527

To apply PEP policies:

Access: P&R Admin/Admin

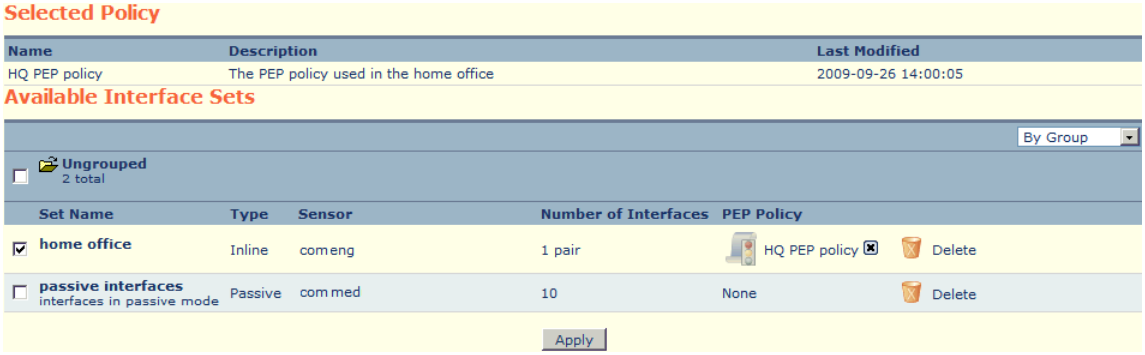
1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

2. Click **Apply**.
   - If the policy is associated with an interface set, continue to step 3.
   - If the policy is not associated with an interface set, skip to step 4.

3. If the policy is associated with an interface set, an Apply to Interface Sets dialog box appears.
   - Click **Quick Apply** to apply the policy to the interface set or sets the policy targets. The PEP Policy is applied.
   - Click **Custom Apply** to apply the policy to interface sets other than the policy targets. The PEP Policy Apply page appears; skip to step 4.
   - Click **Cancel** to return to the PEP Policy Management page.

4. If you attempt to apply a PEP policy that is not associated with an interface set or you use **Custom Apply** to apply it to a different interface set, the PEP Policy Apply page appears.



   You can sort the available interface sets by:
   - Interface Set Group
   - Interface Set Type (inline fail-open, inline, or passive)
   - Sensor
   - PEP policy

5. Select the interface set name, then click **Apply**.

   Your PEP policy is applied to the interface set.

See the following for more information about PEP policies:

- Understanding PEP Traffic Management on page 1511
- Creating PEP Policies on page 1513

Keysight, EX-1004
Page 1528

- Editing PEP Policies on page 1526
- Exporting a PEP Policy on page 2002
- Deleting PEP Policies on page 1529
- Viewing Active PEP Policy Details on page 1529

Deleting PEP Policies

Platform: DC/MDC

Component: None

You can delete a PEP policy that is no longer in use.

To delete a PEP policy:

Access: P&R Admin/Admin

1. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

2. Click **Delete** then **OK** on the confirmation dialog box to delete the PEP policy.

   The system deletes the policy and returns to the PEP Policy Management page.

See the following for more information about PEP policies:

- Understanding PEP Traffic Management on page 1511
- Creating PEP Policies on page 1513
- Editing PEP Policies on page 1526
- Applying PEP Policies on page 1527
- Exporting a PEP Policy on page 2002
- Viewing Active PEP Policy Details on page 1529

Viewing Active PEP Policy Details

Platform: 3D9900

Component: None

You can view a PEP policy that is applied to an interface set. Note that you cannot make changes to an active policy from this view.

Note that you cannot view an active PEP policy from the Defense Center; you must access the policy from the web interface of the sensor where the policy is applied.

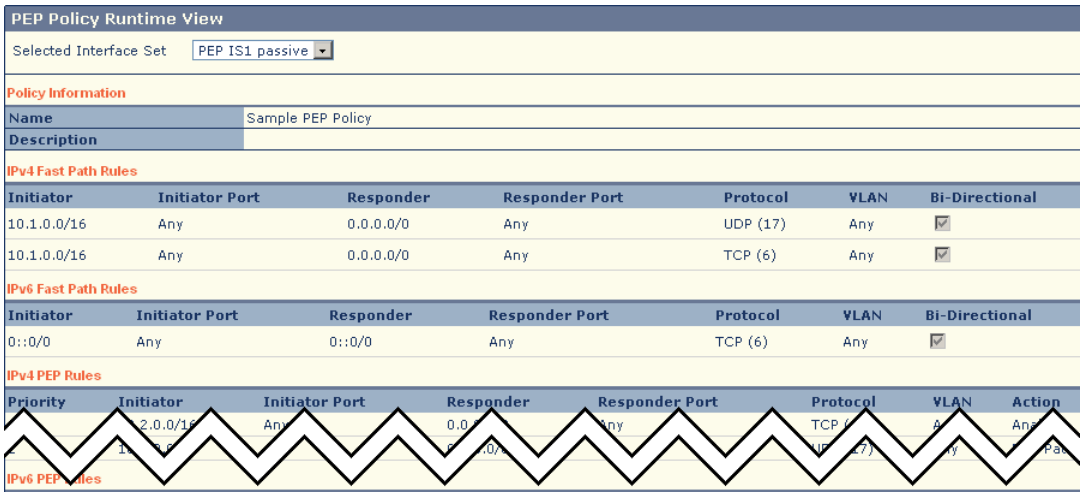
Keysight, EX-1004
Page 1529

To view an active PEP Policy:

Access: Admin

1. Select **Operations > Configuration > Detection Engines > Interface Sets**. Click on the PEP policy name.

   The Policy Runtime View page appears and you can see the PEP policy details.



2. To view the PEP policy on other interface sets, use the Selected Interface Set drop-down.

   Choose the interface set with the policy that you want to view.

See the following for more information about PEP policies:

- Understanding PEP Traffic Management on page 1511
- Creating PEP Policies on page 1513
- Editing PEP Policies on page 1526
- Applying PEP Policies on page 1527
- Exporting a PEP Policy on page 2002
- Deleting PEP Policies on page 1529

Keysight, EX-1004
Page 1530

Chapter 34
Working with Event Reports

The Sourcefire 3D System provides a flexible reporting system that you can use to generate a variety of reports that include the data you see on the event view pages. You can create and generate a report at any time, or you can also use predefined templates, called report profiles, to generate reports. You can also use profiles as templates to create your own new profiles.

See the following sections for more information:

- Generating Reports from an Event View on page 1531
- Using Report Profiles on page 1538

Generating Reports from an Event View

Platform: DC/MDC/Standalone Sensor

Component: None

You can generate reports based on any subset of events in an event view. You can then view, download, or delete reports, as well as move reports to a remote storage location. If you are using a Defense Center, you can run reports remotely from the Defense Center using the data on the sensors for the reports. You can also use your Defense Center to store reports locally or remotely. For more information, see the following sections:

- Generating a Basic Report on page 1532
- Viewing Generated Reports on page 1534
- Downloading Generated Reports on page 1535
- Deleting Generated Reports on page 1535

Keysight, EX-1004
Page 1531

- Running Remote Reports on page 1536
- Using Remote Storage for Reports on page 1537

Generating a Basic Report

Platform: DC/MDC/Standalone Sensor

Component: None

You can generate reports on any subset of events in an event view.

To generate a report for a specific set of events:

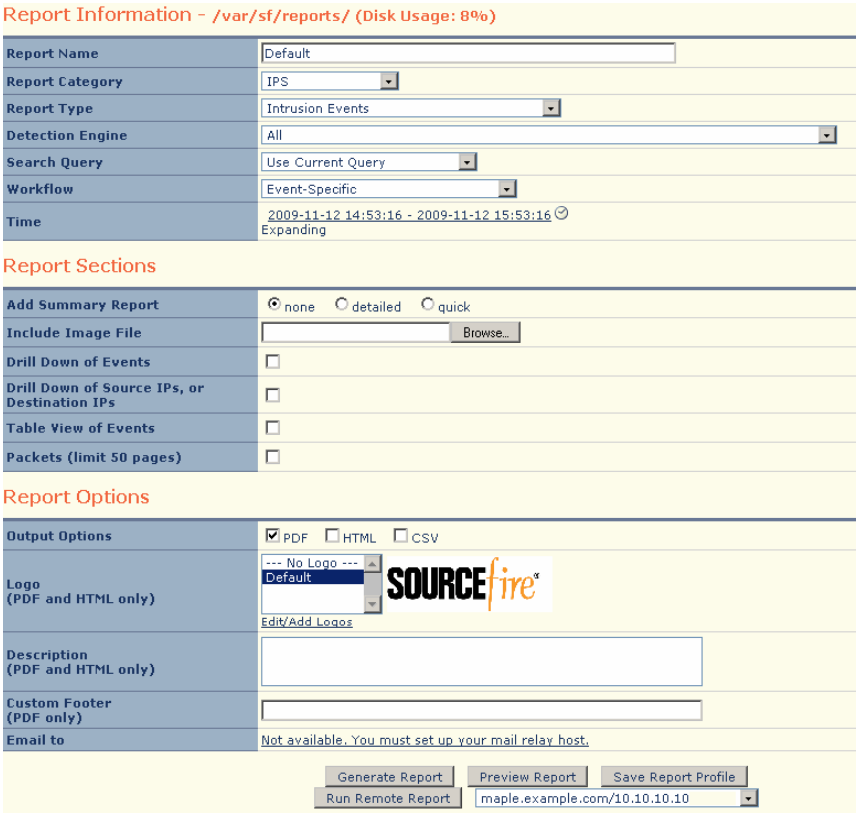
Access: Any Analyst/Admin

1. Populate an event view with the events you want to include in the report. You can do this several ways:
   - Use an event search to define the type of events you want to view. For details on using the event search, see Searching for Events on page 1550.
   - Drill down through a workflow until you have the proper events in your event view. For details on using workflows and constraining events within a workflow, see Understanding and Using Workflows on page 1573.

Keysight, EX-1004
Page 1532

2. Click **Report Designer** in the toolbar.

   The Report Designer page appears. The settings on the page reflect the event parameters that you selected. The following graphic shows the Defense Center version of the page.



TIP! If you need to go back to the drill-down page where you opened the Report Designer, click **Return to Calling Page** at the bottom of the Report Designer page.

3. Change any of the parameters as necessary to meet your needs.

   For details on the parameters for a report, see Creating a Report Profile on page 1542.

4. Select the check boxes next to the output options you want in the report: **PDF**, **HTML**, or **CSV**. Note that you may select more than one format.

5. Click **Generate Report**.

Keysight, EX-1004
Page 1533

6. Click **OK** to confirm that you want to save the current parameters as a report profile.

   The report profile is saved and the report generates in the output formats you selected.

7. To view the report, click **Reports** in the toolbar, then click the report name on the Reporting page that appears.

   The report appears.

Viewing Generated Reports

Platform: DC/MDC/Standalone Sensor

Component: None

The Reports page lists all your previously generated event reports. Each report is listed with the report name, the date and time the report was generated, who generated it, and whether it is stored locally or remotely. The default location for report storage is listed at the top of the page; for local, NFS, and SMB storage, the appliance provides the disk usage of the storage device.

Each report has one of the following file extensions appended to the report name:

- `.csv` for comma-separated value reports
- `.pdf` for PDF reports
- `.zip` for HTML reports (HTML reports are zipped along with the necessary graphics)

You can also view the status of each report, which indicates whether it has yet to be generated (for example, for scheduled tasks), it has already been generated, or whether the generation failed (for example, due to lack of disk space).

Note that only Series 2 Defense Centers support remote storage of reports. You can enable or disable remote storage using the **Enable Remote Storage for Reports** check box. If you disable remote storage, the Defense Center hides any previously generated remotely stored reports. In addition, if you change the remote storage location, the Defense Center hides reports not stored in the new location. To configure remote storage, click **Remote Storage** on the toolbar. For more information, see Managing Remote Storage on page 1817.

Use the following procedure to view generated reports. Note that users with the Administrator role access can view all the reports; other users can view only the reports they generated.

To view a generated report:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Report Profiles**.

   The Report Profiles page appears.

Keysight, EX-1004
Page 1534

2. On the toolbar, click **Reports**.

   The Reports page appears.

3. You have two options:

   - Enable the check box next to the report you want to view, then click **View**.
   - Click the name of the report.

   In either case, the report opens.

Downloading Generated Reports

Platform: DC/MDC/Standalone Sensor

Component: None

Use the following procedure to download generated reports.

To download generated reports:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Report Profiles**.

   The Report Profiles page appears.

2. On the toolbar, click **Reports**.

   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.

   The reports are downloaded in a single `.zip` file.

Deleting Generated Reports

Platform: DC/MDC/Standalone Sensor

Component: None

Use the following procedure to delete generated reports.

Keysight, EX-1004
Page 1535

To delete generated reports:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Report Profiles**.

   The Report Profiles page appears.

2. On the toolbar, click **Reports**.

   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 that you want to delete the reports.

   The reports are deleted.

Running Remote Reports

Platform: DC

Component: None

If you use a Defense Center to manage your sensors, you have the option of running reports remotely from the Defense Center using the data on the sensors. For example, if you use your Defense Center to manage a 3D Sensor with IPS, and you store IPS data on the sensor in addition to sending it automatically to the Defense Center, you can run the report on the data that is resident on the sensor.

There are several limitations that you need to keep in mind:

- If you do not store data on the sensor, then the remote report will be empty.
- If your report uses a logo or image file, the logo or image file must exist on both the Defense Center and the managed sensor where you run the report.
- You cannot run incident reports remotely on managed 3D Sensors with IPS.
- You cannot run remote reports on 3Dx800 or Crossbeam-based software sensors.

To run a remote report:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Report Profiles**.

   The Report Profiles page appears.

Keysight, EX-1004
Page 1536

2. Click **Create Report Profile**.

   The Report Designer page appears.

3. Create the report that you want to run on the managed sensor.

   See Generating a Basic Report on page 1532 for details.

4. From the drop-down list at the bottom of the page, select the sensor where you want to run the report and click **Run Remote Report**.

   A prompt appears asking you to confirm that you want to run the report remotely.

5. Click **OK**.

   The report is run on the sensor that you selected.

6. In the toolbar, click **Reports**.

   The Reporting page appears, listing the report you just generated on the managed sensor. Note that `remote-` is prepended to the name of the report.

7. You can view or download the remote report as you would with any other locally generated report.

TIP! You can also use report profiles as the basis for remote reports by creating a profile as described in Creating a Report Profile on page 1542. When you run the report, make sure you select the name of the sensor and click **Run Report Remotely**.

Using Remote Storage for Reports

Platform: DC/MDC

Component: None

On Series 2 Defense Centers, you can move locally stored reports to a remote storage location. Note that after you move a report to a remote location, you cannot move it back. For information on configuring a remote storage location and enabling remote storage of reports, see Managing Remote Storage on page 1817.

To move generated reports:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Report Profiles**.

   The Report Profiles page appears.

2. On the toolbar, click **Reports**.

   The Reports page appears.

Keysight, EX-1004
Page 1537

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.

Using Report Profiles

Platform: DC/MDC/Standalone Sensor

Component: None

Report profiles are templates that you can use to help you generate event reports. You can use one of the Sourcefire-provided report profiles or you can create one of your own. You can use a report profile to generate a report any time you need it. You can also use a profile as the basis for a regularly scheduled report.

See the following sections for more information:

- Using a Predefined Report Profile on page 1538
- Creating a Report Profile on page 1542
- Setting Up the Report Options on page 1547
- Editing Report Profiles on page 1549
- Deleting Report Profiles on page 1549

Using a Predefined Report Profile

Platform: DC/MDC/Standalone Sensor

Component: None

A predefined report profile provides you with predefined setting for event reports. As with custom report profiles that you create (see Creating a Report Profile on page 1542), you can use a predefined report profile as a template for an event report. You can modify field settings as appropriate, save the report with the new values, and run the report manually or automatically.

Sourcefire has included three predefined report profiles:

- Blocked Events
- High Priority Events
- Host Audit

Keysight, EX-1004
Page 1538

You can use these profiles as is, or you use them as the basis for your own profiles. The tables that follow list the default settings for each of the predefined profiles.

The Blocked Events report profile provides information on intrusion events that indicate attacks on your network were dropped by the 3D Sensor. This report profile is available only on a Defense Center that manages 3D Sensors with RNA and IPS.

Blocked Events Report Profile

Field	Setting
Report Name	Blocked Events
Report Category	IPS
Report Type	Intrusion Events
Detection Engine	All
Search Query	Blocked Events
Workflow	Impact and Priority (on the Defense Center) Destination Port (on the 3D Sensor)
Time	Last day, sliding time window
Add Summary Report	Quick
Impact Based Event Summary (on the Defense Center)	Enabled
Drill Down of Source and Destination IPS (on the Defense Center)	Enabled
Drill Down of Destination Port (on the 3D Sensor)	Enabled
Drill Down of Events (on the 3D Sensor)	Enabled
Table View of Events	Disabled
Packets (limit 50 pages)	Disabled

Keysight, EX-1004
Page 1539

The High Priority Events report profile provides information on intrusion events correlated with the host criticality of the hosts involved in the events. This report profile is available only on a Defense Center that manages 3D Sensors with RNA and IPS.

High Priority Events Report Profile

Field	Setting
Report Category	IPS
Report Type	Intrusion Events with Destination Criticality
Detection Engine	All
Search Query	High Priority Events
Workflow	Events by Impact, Priority, and Host Criticality
Time	Last day, sliding time window
Add Summary Report	Quick
Impact to Criticality Summary	Enabled
Source Destination Drill Down	Enabled
Intrusion Events with Destination Criticality	Enabled
Packets (limit 50 pages)	Disabled

The Host Audit report profile provides operating system details for the hosts less than two network hops away from 3D Sensors with RNA. This report profile is available only on a Defense Center that manages 3D Sensors with RNA.

Default Settings for the Host Audit Report Profile

Field	Setting
Report Category	RNA
Report Type	RNA Hosts
Detection Engine	All

Keysight, EX-1004
Page 1540

Default Settings for the Host Audit Report Profile (Continued)

Field	Setting
Search Query	Local Systems
Workflow	Operating System Summary
Time	Last week, sliding time window
Add Summary Report	summary
Summary of OS Names	Enabled
Summary of OS Versions	Enabled
OS Details with IP, NetBIOS, Criticality	Enabled
Table View of Events	Disabled
Packets (limit 50 pages)	Disabled

To generate a report using a predefined report profile:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Reporting**.

   The Reporting page appears. All predefined report profiles are listed, including any that you have previously created.

2. Click the name of the report profile you want to use.

   The Report Designer page appears with the parameters previously defined for the profile.

3. Click **Generate Report**.

   Your report is generated.

4. Click **Reports** in the toolbar.

   The Reports page appears, listing the report that you generated as well as any other previously generated reports. For information on managing generated reports, see Viewing Generated Reports on page 1534.

Keysight, EX-1004
Page 1541

Creating a Report Profile

Platform: DC/MDC/Standalone Sensor

Component: None

You can create a report profile by defining the report category and type, and then specifying which detection engines to search, the criteria for the search, and which workflows to examine. Not all options are available for all reports. For example, in the IPS report category, selecting the Intrusion Events report type gives you the option to select which detection engines to search.

The Report Category defines the type of events you want to include in the report. The following table describes the report categories.

Report Categories

Select...	If you want to create a profile for...
IPS	intrusion events or the SEU import log.
RNA	host attributes, RNA client applications, vulnerabilities, flow data, hosts with services, RNA hosts, RNA events, RNA services, or scan results.
RUA	RUA events and users.
Compliance	white list violations, remediation status, compliance events, or white list events.
Third-Party Vulnerabilities	third-party vulnerabilities.
Health Monitoring	health events.
Audit Log	audit log events.

The Report Types table describes the reports you can create and the components required for producing them.

Report Types

Report Type	Report Category	Requires
Intrusion Events	IPS	DC + IPS
Intrusion Events with Destination Criticality	IPS or RNA	DC + RNA + IPS

Keysight, EX-1004
Page 1542

Report Types (Continued)

Report Type	Report Category	Requires
Intrusion Events with Source Criticality	IPS or RNA	DC + RNA + IPS
SEU Import Log	IPS	DC + IPS
Host Attributes	RNA	DC + RNA
RNA Hosts	RNA	DC + RNA
Scan Results	RNA	DC + RNA
RNA Client Applications	RNA	DC + RNA
RNA Events	RNA	DC + RNA
RNA Services	RNA	DC + RNA
Vulnerabilities	RNA	DC + RNA
Hosts with Services	RNA	DC + RNA
Flow Data	RNA	DC + RNA
RUA Events	RUA	DC + RUA
Users	RUA	DC + RUA
White List Violations	Compliance	DC + RNA
Compliance Events	Compliance	DC + RNA
White List Events	Compliance	DC + RNA
Remediation Status	Compliance	DC + RNA
Third-Party Vulnerabilities	Third-Party Vulmerabilities	DC + RNA
Health Events	Health Monitoring	DC
Audit Log Events	Audit Log	Any

Keysight, EX-1004
Page 1543

You can include a summary of the event data in your profile. The Comparison of Quick and Detail Summary Reports table describes the kind of information included in the intrusion event summaries.

Comparison of Quick and Detail Summary Reports

Report Information	Quick Summary	Detail Summary
Pie chart showing the percentage of events in each event type (which maps to the rule category for the rule that generated the event)	X	X
List of the 10 most active and 10 least active events	X	X
Graph showing the number of events over time	X	X
Pie charts showing the percentage of events by protocol (for example, TCP, UDP, or ICMP) and event classification (which maps to the value for the `classtype` keyword in the rule that generated the event)	X	X
Tables listing the 50 most active and least active events	X	X
Tables listing the 50 most active source and destination ports	X	X
Tables listing the 25 most active source and destination hosts and host combinations.	X	X
Tables listing the 25 most active source and destination hosts as well as the 25 most active source and host combinations		X
Tables listing the most active events for each of the 25 most active destination hosts		X
Tables listing the most active events for the 25 most active source and destination host combinations		X

The RNA Summary includes:

- RNA event statistics including total number of events, events in the last day and hour, total services, total hosts, total routers, total bridges, and host limit usage
- a list of events divided by event type with counts for the last hour and total number within the report range
- pie charts showing the percentage of events by protocol (for example, TCP, UDP, or ICMP), service, and operating system

Keysight, EX-1004
Page 1544

To create a report profile:

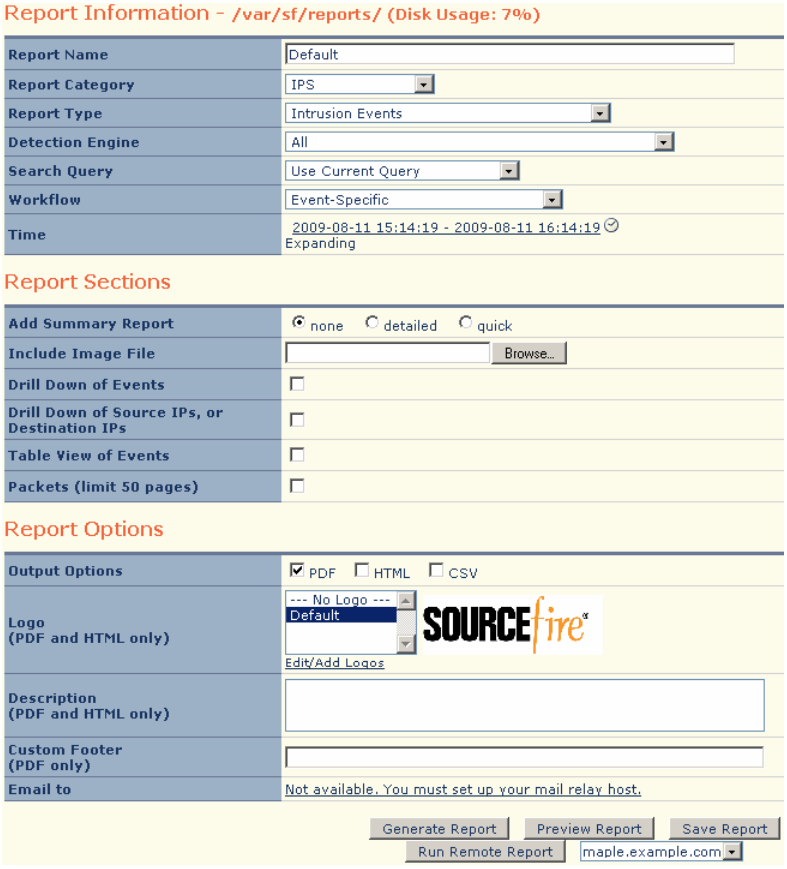
Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Reporting**.

   The Reporting page appears.

2. Click **Create Report Profile**.

   The Report Designer page appears. The following graphic shows the Defense Center version of the page.



3. In **Report Name**, enter a name for you report profile. You can use up to 80 alphanumeric characters, periods, dashes, parentheses, and spaces in the name.

4. From the **Report Category** drop-down list, select the report category for which you want to create a report.

   See the Report Categories table on page 1542 for more information.

Keysight, EX-1004
Page 1545

5. From the **Report Type** drop-down list, select the type of report you want to create.

   The options for report type are based on your selection for report category. see the the Report Types table on page 1542 for more information.

6. Optionally, if the report type you selected includes the **Detection Engine** option, select a specific detection engine on which to report.

7. Optionally, if you are reporting on health events on the Defense Center, select a specific sensor or sensor group from the **Sensor** drop-down list.

8. From the **Search Query** drop-down list, either use the **Use Current Query** option (which retains any query parameters you specified on the search page or event page) or select one of the existing search queries.

   Note that if you did not previously specify a search query, the **Use Current Query** option places no constraints on the events.

9. From the **Workflows** list, select the workflow you want to use to build the report.

   For information on workflows, see Understanding and Using Workflows on page 1573.

10. Specify the time range for the report.

    Depending on your default time window, the time range matches either the time window for the event view you are using to building the report profile, or the global time window. You can change time range by clicking it and using the Date/Time pop-up window to select a new time range. For more information, see Setting Event Time Constraints on page 1597.

IMPORTANT! For report profiles that you plan to use multiple times, such as in scheduled tasks, Sourcefire **strongly** recommends that you use a sliding time range. If you create a report profile with a static time range, the appliance will generate a report using the same time range (and therefore the same events) every time you use the report profile.

Keysight, EX-1004
Page 1546

11. If a summary is available for the report type you selected, specify whether you want to include it as part of your report.
    - To include a summary with intrusion event-based reports, select **quick** or **detailed**. For a description of the information provided in Quick and Detailed summaries, see the Comparison of Quick and Detail Summary Reports table on page 1544.
    - On a Defense Center with an RNA host license, to include a summary with an RNA-based report, select **summary**.
    - To exclude the summary, select **none**, which is the default.
12. If you want to include an image, which will be displayed between the summary report and the drill-down or table views, type the path to the image in the **Include Image File** text box, or navigate to a JPEG, PNG, or TIFF file.

    This can be useful for providing information best displayed in a non-textual format or simply as a break between sections.
13. Select the check boxes next to the sections of the workflow you want to include in the profile.

    The options in this section depend on the workflow you selected in step 9.

    TIP! Note that if you select a table view of events, the report is limited to 10,000 records, regardless of the number of events specified by the search query.
14. Continue with Setting Up the Report Options on page 1547.

Setting Up the Report Options

Platform: DC/MDC/Standalone Sensor

Component: None

Use the options in the Report Options section of the page to specify how your report will be produced and whether you want to automatically email the report.

You can generate a report in PDF, HTML or comma-separated value (CSV) formats. Note that graphics are not generated in the CSV format. You can include a logo on your report. In PDF formats, the logo is included on every page. In HTML formats, the logo is included at the top of the report. You can add a description which will be included on the front page summary of the report.

Keysight, EX-1004
Page 1547

To define the report options:

Access: Any Analyst/Admin

1. Select the check boxes next to one or more output options for your report: **PDF**, **HTML**, or **CSV**.
2. Optionally, for PDF and HTML reports, select a logo from the list of image files that were previously added to the system.
3. Optionally, for PDF and HTML reports, type a description in the **Description** field. You can use alphanumeric characters and spaces. The description appears in the report header.
4. Optionally, for PDF reports, type the text you want to include as the footer in the **Custom Footer** field. You can use 1 - 80 alphanumeric characters and spaces.
5. Optionally, you can specify that reports are automatically emailed after they are generated. To email a report, type one or more email addresses in a comma-separated list in the **Email to** field.

   IMPORTANT! You must make sure that the mail host is identified: If you have not previously set up a mail host, click **Not available. You must set up your mail relay host.**

   The report is emailed from *host_name*@*domain_name*, where *host_name* is the host name of the appliance and *domain_name* is the name of the domain where you deployed the appliance.
6. You have the following options:
   - To save the report profile, click **Save Report Profile**.

     When prompted, follow the instructions for your browser to save the report profile.

     The report profile is saved with the name you specified in the **Report Name** field.
   - To save the report profile and, at the same time, generate a report based on the profile, click **Generate Report**.

     At the prompt, follow the instructions to generate the report and save the report profile.
   - To see a PDF preview of your report, click **Preview Report**.

     At the prompt, follow the instructions to display the report in your browser.
   - On a Defense Center, to generate the report remotely, select the sensor where you want to run the report and click **Run Remote Report**.

     At the prompt, follow the instructions to generate the report and save the report profile.

Keysight, EX-1004
Page 1548

Editing Report Profiles

Platform: DC/MDC/Standalone Sensor

Component: None

You can create a new report profile by using an existing report profile as a basis, modifying the settings as appropriate, and saving the profile with a different name. You can also edit a report profile to make changes to the resulting report.

To edit a report profile:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Reporting**.

   The Reporting page appears.

2. Click **Edit** next to the profile that you want to edit.

   The Report Designer page appears and contains the current settings for the profile.

3. Make changes to the profile as needed.

   See the following sections for information:

   - Creating a Report Profile on page 1542
   - Setting Up the Report Options on page 1547

   IMPORTANT! If you are creating a new report profile from a predefined or existing report profile, remember to change the name of the report profile in the **Report Name** field.

4. Click **Save Report Profile**. When prompted, follow the instructions for your browser to save the report profile. The report profile is saved with the name you specified in the **Report Name** field.

Deleting Report Profiles

Platform: DC/MDC/Standalone Sensor

Component: None

Use the following procedure to delete a report profile.

To delete a report profile:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Reporting**.

   The Reporting page appears.

2. Click **Delete** next to the profile that you want to delete.

   The report profile is deleted.

Keysight, EX-1004
Page 1549

Chapter 35
Searching for Events

Sourcefire 3D Sensors and Defense Centers 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 re-use later. You can also create your own search criteria.

TIP! On the Defense Center, you can search both predefined and custom tables for specific events or other data.

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 1551
- Using Wildcards and Symbols in Searches on page 1555
- Specifying Time Constraints in Searches on page 1556
- Specifying IP Addresses in Searches on page 1557
- Specifying Ports in Searches on page 1558
- Specifying Detection Engine/Appliance Names on page 1559

Keysight, EX-1004
Page 1550

Performing and Saving Searches

Platform: DC/MDC/3D Sensor

Component: None

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 on the appliance. For more information, see the following sections:

- Performing a Search on page 1551 explains how to perform a search and, optionally, save it.
- Loading a Saved Search on page 1554 explains how to load a previously saved search.
- Deleting a Saved Search on page 1554 explains how to delete a previously saved search.

IMPORTANT! To search a custom table on the Defense Center, you follow a slightly different procedure. For more information, see Searching Custom Tables on page 1569.

Performing a Search

Platform: DC/MDC/3D Sensor

Component: None

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 re-use later. You can also create your own search criteria. For more information, see Loading a Saved Search on page 1554.

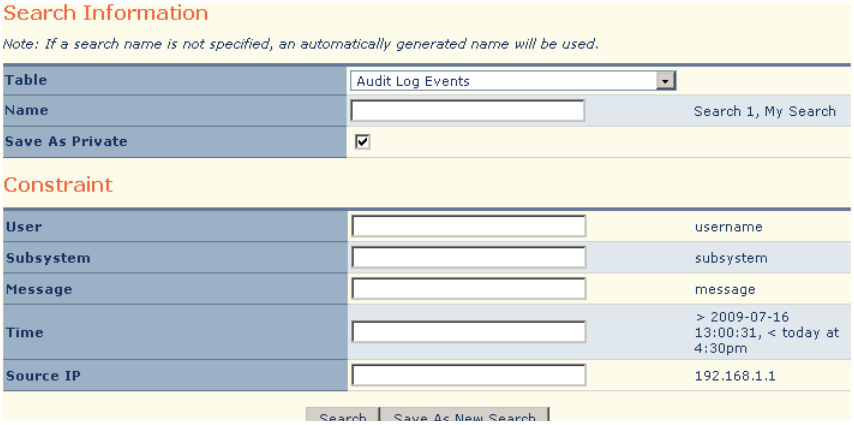
Keysight, EX-1004
Page 1551

To perform a search:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Searches**, then select the type of event or data you want to search for.

   The Search page appears. The following graphic shows the search page for the audit log.



   TIP! To search the database for a different kind of event or data, select it from the **Table** drop-down list.

2. 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.

3. Enter your search criteria in the appropriate fields.

   If you enter multiple criteria, the search returns only the records that match all the criteria. See the sections listed in the following table for detailed information on the search criteria you can use.

For search criteria for...	See...
Intrusion Events	Searching for Intrusion Events on page 309
RNA Events	Searching for RNA Network Discovery Events on page 1016
Hosts	Searching for Hosts on page 1027
Host Attributes	Searching for Host Attributes on page 1038

Keysight, EX-1004
Page 1552

For search criteria for...	See...
Services	Searching for Services on page 1045
Client Applications	Searching for Client Applications on page 1054
Flow Data	Searching for Flow Data on page 1094
Sourcefire Vulnerabilities	Searching for Sourcefire Vulnerabilities on page 1062
Third-Party Vulnerabilities	Searching for Third-Party Vulnerabilities on page 1068
Compliance Events	Searching for Compliance Events on page 1276
White List Events	Searching for Compliance White List Events on page 1201
White List Violations	Searching for White List Violations on page 1208
Remediation Status Events	Searching for Remediation Status Events on page 1349
SEU Import Log	Searching the SEU Import Log on page 1865
Audit Log Events	Searching Audit Records on page 1978
Health Events	Searching for Health Events on page 1965
Scan Results	Searching for Scan Results on page 1449
RUA Users	Searching for RUA Users on page 1497
RUA Events	Searching for RUA Events on page 1506

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.

Keysight, EX-1004
Page 1553

5. You have the following options:
   - Click **Search** to start the search.

     Your search results appear in the default workflow for the data you are searching, constrained by time (if applicable). To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. 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

Platform: DC/MDC/3D Sensor

Component: None

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: Any Analyst/Admin

1. You have two options:
   - From any page on a workflow, click **Search**.
   - Select **Analysis & Reporting > Searches**, 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

Platform: DC/MDC/3D Sensor

Component: None

If you have saved searches, you can delete them from the Search page.

Keysight, EX-1004
Page 1554

To delete a saved search:

Access: Any Analyst/Admin

1. You have two options:
   - From any page on a workflow, click **Search**.
   - Select **Analysis & Reporting > Searches**, then select the event type 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

Platform: DC/MDC/3D Sensor

Component: None

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*`.

Keysight, EX-1004
Page 1555

Specifying Time Constraints in Searches

Platform: DC/MDC/3D Sensor

Component: None

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 Time Specification in Search Fields 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 2:23pm, March 22, 2006.
>	`> today at 2:45pm`	Returns events with a timestamp later than today at 2:45pm.

Keysight, EX-1004
Page 1556

Specifying IP Addresses in Searches

Platform: DC/MDC/3D Sensor

Component: None

When specifying IP addresses in searches, you can enter an individual IP address, a list of IP addresses, or a range of IP addresses. You can also use negation. The following table contains examples of valid ways to enter IP addresses.

Acceptable IP Address Syntax

To specify...	Type...	For example...
a single IP address	the IP address.	`192.168.1.1`
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`
a range of IP addresses that can be specified with a CIDR block	the IP address range in CIDR 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 59.
a range of IP addresses that cannot be specified with a CIDR block	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`
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 or IP address range (including CIDR blocks).	`!192.168.1.10`
on a Defense Center with RNA, all IP addresses in the networks you are monitoring with RNA	`local`	`local`
on a Defense Center with RNA, all IP addresses that are **not** in the networks you are monitoring with RNA	`remote`	`remote`

Keysight, EX-1004
Page 1557

Specifying Ports in Searches

Platform: DC/MDC/3D Sensor

Component: None

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.

Keysight, EX-1004
Page 1558

Specifying Detection Engine/Appliance Names

Platform: DC/MDC/3D Sensor

Component: None

Your Sourcefire 3D System uses a specific syntax for detection engines, group names, and appliance names in an event search. If you enter a string in the **Detection Engine** field, the search is conducted in the following order for:

- an event detected by a detection engine with that name
- an event detected by a detection engine that belongs to a detection engine group with that name
- on a Defense Center, an event detected by a 3D Sensor with that name
- on a Defense Center, an event detected by a 3D Sensor that belongs to a sensor group with that name

That is, if there is a detection engine whose name matches the string you entered, the search does not use detection engine groups, sensors, or sensor groups with that name as a search criterion.

Note that the context for detection engine group searches is local. For example, assume that you create two detection engine groups on a 3D Sensor and name them `apple` and `orange`. Next you create two detection engine groups on the sensor's managing Defense Center and name them `horse` and `lion`. If you conduct a search based on detection engine groups on the Defense Center, it will not return any results for a search that includes `apple` or `orange`. A search on the Defense Center will only return results for a search that uses `horse` or `lion` (assuming there are valid results that can be returned).

WARNING! On a 3D Sensor with IPS, if you search using a term that does not match any detection engines, you retrieve all events on the sensor. This occurs because a standalone sensor does not have a sensor name, and the search engine defaults to returning all events in case the supplied search term was intended as a sensor name.

If you need to refine your search further, use the following syntax:

`DE_name``:`*`DE_group`*

For example:

```
Default*:DE_GROUP_1
```

On a Defense Center, you can use the detection engine name, the sensor name and sensor group:

`DE_name``:`*`DE_group`*`/`*`sensor_name`*`:`*`sensor_group`*

For example:

```
Default*:DE_GROUP_1/Sensor13:NOC4_Sensors
```


Keysight, EX-1004
Page 1559

This syntax allows you to search for events detected by a detection engine whose name starts with `Default`, that is in the detection engine group `DE_GROUP_1`, which is on the sensor `Sensor13` which is in the sensor group `NOC4_Sensors`.

Keep in mind that you do not have to use all the parameters. For example:

```
*:DE GROUP 1/*
```

searches for events detected by any detection engine in the detection engine group `DE_GROUP_1`

As another example:

```
Primary:*/*:NOC4_Sensors
```

searches for events detected by detection engines named called `Primary` that exist on sensors in the `NOC4_Sensors` group.

On a Master Defense Center, you can use the detection engine name, the sensor name, the Defense Center (or appliance) name and the Defense Center (or appliance) group name:

DE_name:*DE_group/sensor_name*:*sensor_group/app_name:app_group*

For example:

```
Default*:DE_GROUP_1/Sensor13:NOC4_Sensors/DC_Alpha:NOC_DCs
```

This syntax allows you to search for events detected by a detection engine whose name starts with `Default`, that is in the detection engine group `DE_GROUP_1`, which is on the sensor `Sensor13` which is in the sensor group `NOC4_Sensors` connected to Defense Center `Alpha` in the Defense Center group `NOC_DCs`.

Keep in mind that you do not have to use all the parameters. For example:

```
Active:*/*:NOC_DCs
```

searches for events detected by detection engines named called `Active` that exist on sensors connected to Defense Centers in the `NOC_DCs` group.

Keysight, EX-1004
Page 1560

Chapter 36
Using Custom Tables

As 3D Sensors with RNA collect information about your network and send the information to the Defense Center that manages them, the Defense Center stores the information 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 Services by Count workflow are taken from the fields in the RNA Services 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 RNA Flow Data table and then examine flow data in a new context.

IMPORTANT! You must use a Defense Center with an RNA host license to work with custom tables. You cannot view custom tables on a Master Defense Center.

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 1618.

The following sections describe how to create and use your own custom tables.

- Understanding Custom Tables on page 1562
- Creating a Custom Table on page 1565
- Modifying a Custom Table on page 1568
- Deleting a Custom Table on page 1568

Keysight, EX-1004
Page 1561

- Viewing a Workflow Based on a Custom Table on page 1569
- Searching Custom Tables on page 1569

Understanding Custom Tables

Platform: DC

Component: RNA

Custom tables contain fields from two or more of the 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 the 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 RNA host data so that, if your Defense Center is managing sensors with both RNA and IPS, 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
Intrusion Events with Destination Criticality (Requires IPS + RNA)	Includes fields from the Intrusion Events table and the RNA Hosts table, providing you with information on the intrusion events generated by IPS, 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 (Requires IPS + RNA)	Includes fields from the Intrusion Events table and the RNA Hosts table, providing you with information on the intrusion events generated by IPS, as well as 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.
Hosts with Services	Includes fields from the RNA Hosts and RNA Services tables, providing you with information about the detected services running on your network, as well as basic operating system information about the hosts running those services.

Keysight, EX-1004
Page 1562

Understanding Possible Table Combinations

Platform: DC

Component: RNA

When creating 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...
Compliance Events	• Host Attributes • RNA Hosts
Intrusion Events (Requires: IPS + RNA)	• Host Attributes • RNA Hosts • RNA Services
Flow Summary Data	• Host Attributes • RNA Hosts • RNA Services
Host Attributes	• Compliance Events • Intrusion Events (Requires: IPS + RNA) • Flow Summary Data • RNA Client Applications • RNA Events • Flow Data • RNA Hosts • RNA Services • White List Events
RNA Client Applications	• Host Attributes • RNA Hosts
RNA Events	• Host Attributes • RNA Hosts

Keysight, EX-1004
Page 1563

Custom Table Combinations (Continued)

You can combine fields from...	With fields from...
Flow Data	• Host Attributes • RNA Hosts • RNA Services
RNA Hosts	• Compliance Events • Intrusion Events (Requires: IPS + RNA) • Flow Summary Data • Host Attributes • RNA Client Applications • RNA Events • Flow Data • RNA Services • White List Events
RNA Services	• Intrusion Events (Requires: IPS + RNA) • Host Attributes • RNA Flow Data • RNA Hosts
White List Events	• Host Attributes • RNA 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 RNA 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 RNA Hosts table each represent a single host with a single IP address. Therefore, when you create a custom table based on the Intrusion Events table and the RNA Hosts table, you must choose whether the data you are displaying from the RNA Hosts table applies to the host with the source IP address or the 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.

Keysight, EX-1004
Page 1564

For more information on creating custom tables, see:

- Creating a Custom Table on page 1565
- Modifying a Custom Table on page 1568
- Deleting a Custom Table on page 1568
- Viewing a Workflow Based on a Custom Table on page 1569
- Searching Custom Tables on page 1569

Creating a Custom Table

Platform: DC

Component: RNA

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.

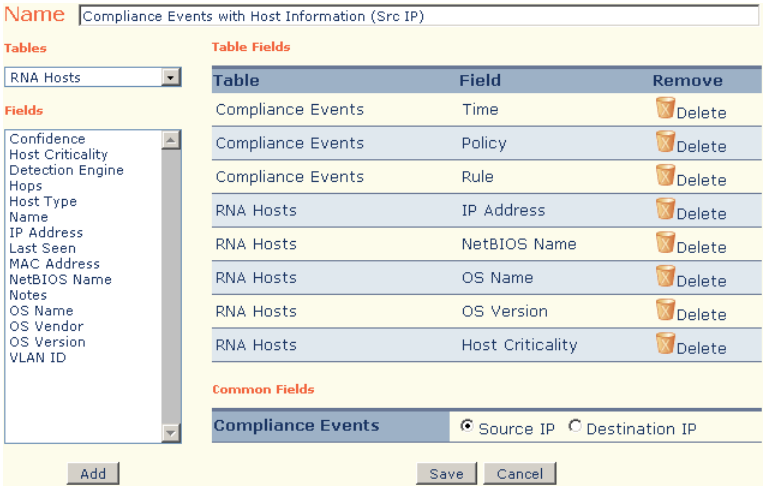
TIP! Instead of creating a new custom table, you can export a custom table from another Defense Center and then import it onto your Defense Center. You can then edit the imported custom table to suit your needs. For more information, see Importing and Exporting Objects on page 1997.

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.

For example, consider a custom table that combines fields from the Compliance Events table and the RNA Hosts table. You can use this custom table to get detailed information about the hosts involved in violations to any of your compliance policies. Note that you need to decide whether to display data from

Keysight, EX-1004
Page 1565

the RNA Hosts table that matches the source IP address or the destination IP address in the Compliance Events table.



If you view the table view of events for this custom table, it displays compliance events, one per row. The following information is included:

- the date and time the event was generated
- the name of the compliance policy that was violated
- name of the rule that triggered the violation
- the IP address of the source, or initiating, host involved in the compliance event
- the source host's NetBIOS name
- the operating system and version the source host is running
- the source host criticality

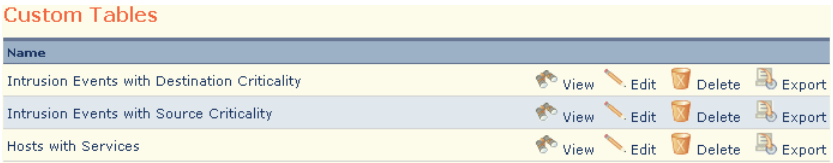
TIP! You could create a similar custom table that displays the same information for destination, or responding, hosts.

To build the custom table in the previous example:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Custom Tables**.

   The Custom Tables page appears.



Keysight, EX-1004
Page 1566

2. Click **Create Custom Table**.

   The Create Custom Table page appears.



3. In the **Name** field, type a name for the custom table, such as `Compliance Events with Host Information (Src IP)`.
4. From the **Tables** drop-down list, select **Compliance Events**.

   The fields in the Compliance Events table appear in the **Fields** list.
5. Under **Fields**, select **Time** and click **Add** to add the date and time a compliance 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 **RNA Hosts**.

   The fields in the Hosts table appear in the **Fields** list. For more information on these fields, see the Host Fields table on page 1023.
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 **Compliance 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 compliance events.

TIP! You could create a custom table that displays detailed host information for the destination, or responding, hosts involved in a compliance event by following this procedure but selecting **Destination IP** instead of **Source IP**.

Keysight, EX-1004
Page 1567

10. Click **Save**.

    The custom table is saved.

Modifying a Custom Table

Platform: DC

Component: RNA

You can add or delete fields in a custom table as your needs change.

To modify a custom table:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Custom Tables**.

   The Custom Tables page appears.

2. Click **Edit** next to the table you want to edit.

   The Create Custom Table page appears. See Creating a Custom Table on page 1565 for information on the various configurations you can change.

3. Optionally, remove fields from the table by clicking **Delete** 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

Platform: DC

Component: RNA

You can delete a custom table that you no longer need.

To delete a custom table:

Access: Any RNA/Admin

1. Select **Analysis & Reporting > Custom Tables**.

   The Custom Tables page appears.

2. Click **Delete** next to the custom table you want to delete.

   The table is deleted.

Keysight, EX-1004
Page 1568

Viewing a Workflow Based on a Custom Table

Platform: DC

Component: RNA

When you create a custom table, RNA automatically creates a default workflow for it. The first page of this workflow displays a table view of events. If your Defense Center is managing a 3D Sensor with IPS and 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 45.

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 1593 for more information.

To view a workflow based on a custom table:

Access: Any RNA/Admin

1. Select **Analysis & Reports > Custom Tables**.

   The Custom Tables page appears.

2. Click **View** next to the custom table upon 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. 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 1597.

Searching Custom Tables

Platform: DC

Component: RNA

You can create and save searches for a custom table. When you create a search you give it a name and specify whether the search is available to you alone or to all users.

Keysight, EX-1004
Page 1569

To perform a search on a custom table:

Access: Any RNA/Admin

1. Select **Analysis & Reports > Custom Tables**.

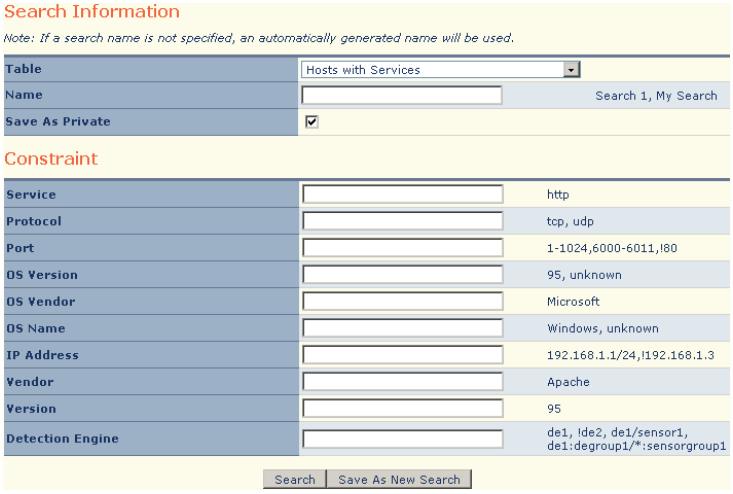
   The Custom Tables page appears.

2. Click **View** 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. 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 1597.

3. On the toolbar, click **Search**.

   The custom table's search page appears. The following graphic shows the search page for the Hosts with Services predefined custom table.



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 automatically created when you save the search.

Keysight, EX-1004
Page 1570

5. Enter your search criteria in the appropriate fields.

   If you enter multiple criteria, the search returns only the records that match all the 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 1978
Client Applications	Searching for Client Applications on page 1054
Compliance Events	Searching for Compliance Events on page 1276
Flow Data	Searching for Flow Data on page 1094
Hosts	Searching for Hosts on page 1027
Host Attributes	Searching for Host Attributes on page 1038
Hosts with Services	Searching for Hosts on page 1027 and Searching for Services on page 1045
Intrusion Events	Searching for Intrusion Events on page 309
Intrusion Events with Destination Criticality	Searching for Intrusion Events on page 309 and Searching for Hosts on page 1027
Intrusion Events with Source Criticality	Searching for Intrusion Events on page 309 and Searching for Hosts on page 1027
Remediation Status Events	Searching for Remediation Status Events on page 1349
RNA Events	Searching for RNA Network Discovery Events on page 1016
RUA Events	Searching for RUA Events on page 1506
SEU Import Log	Searching the SEU Import Log on page 1865
Services	Searching for Services on page 1045
Users	Searching for RUA Users on page 1497

Keysight, EX-1004
Page 1571

For search criteria for....	See...
Vulnerabilities	Searching for Sourcefire Vulnerabilities on page 1062
White List Events	Searching for Compliance White List Events on page 1201
White List Violations	Searching for White List Violations on page 1208

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 save a search as a restriction for restricted event analyst users, 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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.

Keysight, EX-1004
Page 1572

Chapter 37
Understanding and Using Workflows

A *workflow* is a series of pages displaying sensor data that analysts use to evaluate events generated on a sensor. The Sourcefire 3D System provides three types of workflows:

- *Predefined workflows*, which are installed on the system that you cannot modify or delete.
- *Saved custom workflows*, which are predefined custom workflows installed on the Defense Center that you can modify or delete.
- *Custom workflows*, which are workflows you create and customize for your specific needs that you can also modify or delete.

For example, when you are analyzing intrusion events, you can choose between predefined workflows for intrusion events. As another example, there are also predefined workflows for several types of analysis that you can perform for RNA data: event-based, host-based, flow-based, and service-based.

You cannot modify or delete predefined workflows. However, if the predefined workflows do not meet your needs, you can create your own custom workflows. Saved custom workflows are also included on the Defense Center; you can modify these workflows or use them as examples for creating your own custom workflows.

Keysight, EX-1004
Page 1573

See the following sections for more information about using predefined and custom workflows:

- Components of a Workflow on page 1574
- Using Workflows on page 1589
- Using Custom Workflows on page 1618

TIP! You can also use custom workflows as the basis for reports. See Working with Event Reports on page 1531 for more information.

Components of a Workflow

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 RNA events (Defense Center only) includes the Time, Event, IP Address, User, MAC Address, Port, Description, Confidence, and Detection Engine columns.

By contrast, the table view of intrusion events (Master Defense Center, Defense Center, or 3D Sensor with IPS only) includes the Time, Priority, Impact Flag, Inline Result, Detection Engine, Protocol, Source IP, Destination IP, Source User, Destination User, Source Port/ICMP Type, Destination Port/ICMP Code, Generator, and Message 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 RNA 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 only 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 flow data can include graph pages, also called *flow graphs*.

For example, a flow graph might display a line graph that shows the number of flows detected by RNA over time. Generally, flow graphs are, like drill-down

Keysight, EX-1004
Page 1574

pages, intermediate pages that you use to narrow your investigation. For more information, see Understanding Flow Data Graphs on page 1078.

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 either RNA events or third-party vulnerabilities. For more information, see Using Host Profiles on page 944.
- The user detail is the final page for workflows based on RUA events and RUA users. For more information, see Understanding User Details and Host History on page 1495.
- The vulnerability detail view is the final page for workflows based on Sourcefire vulnerabilities. For more information, see Viewing Vulnerability Details on page 980.
- The packet view is the final page for workflows based on intrusion events. For more information, see Using the Packet View on page 290.

Workflows based on other kinds of events (for example, audit log events) do not have final pages.

See the following sections for more information on workflows:

- Comparing Predefined and Custom Workflows on page 1575
- Comparing Workflows for Predefined and Custom Tables on page 1576
- Predefined Intrusion Event Workflows on page 1576
- Predefined RNA Workflows on page 1580
- Predefined Third-Party Vulnerabilities Workflows on page 1582
- Predefined Flow Data Workflows on page 1582
- Predefined RUA Workflows on page 1584
- Predefined Compliance and White List Workflows on page 1585
- Predefined System Workflows on page 1586
- Saved Custom Workflows on page 1587

Comparing Predefined and Custom Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

Your 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.

Keysight, EX-1004
Page 1575

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.

Note that you **cannot** create or view custom workflows on a 3D Sensor that does not have an IPS license.

Comparing Workflows for Predefined and Custom Tables

Platform: DC

Component: RNA

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 RNA data to allow simple searches for events that affect critical systems. See Using Custom Tables on page 1561 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 RNA 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

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Predefined Intrusion Event Workflows table describes the predefined intrusion event workflows included with the Sourcefire 3D System. For

Keysight, EX-1004
Page 1576

information on accessing these workflows, see Viewing Intrusion Events on page 274 and Reviewing Intrusion Events on page 278.

Predefined Intrusion Event Workflows

Workflow Name	Description
Destination Port	Because destination ports are usually tied to a service, this workflow can help you detect services that are experiencing an uncommonly high volume of alerts. The Destination Port column can also help you identify services that should not be present on your network. 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.
Event Specific	This workflow provides two useful features. Events with that occur frequently may indicate: • false positives • a worm • a badly misconfigured network Events that occur infrequently are most likely evidence of a targeted attack and warrant special attention. 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.
Events to Destinations	This workflow provides a high level view of which hosts are being attacked and the nature of the attack. 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.
IP-Specific	This workflow shows which hosts 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 might not belong on the network. 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.

Keysight, EX-1004
Page 1577

Predefined Intrusion Event Workflows (Continued)

Workflow Name	Description
Impact and Priority	Requires: DC/MDC + IPS + RNA 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. 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.
Impact and Source	Requires: DC/MDC + IPS + RNA 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. 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.

Keysight, EX-1004
Page 1578

Predefined Intrusion Event Workflows (Continued)

Workflow Name	Description
Impact to Destination	Requires: DC/MDC + IPS + RNA 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. 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.
Source Port	This workflow indicates which services are generating the most alerts. You can use this information to identify areas that require tuning, and to decide which services require attention. 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.
Source and Destination	This workflow identifies hosts 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. 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.

Keysight, EX-1004
Page 1579

Predefined RNA Workflows

Platform: DC

Component: RNA

There is a predefined workflow for each type of data collected by RNA. For information on accessing these workflows, see Understanding RNA Event Workflows on page 1002.

Predefined RNA Workflows

Workflow Name	Description
Client Applications	This workflow contains a table view of client applications, followed by the host view. See Viewing Client Applications on page 1051 for more information.
Client Application Summaries	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 Client Applications on page 1051 for more information.
Events	This workflow contains a table view of RNA events followed by the host view. See Viewing RNA Network Discovery and Host Input Events on page 1012 for more information.
Flow Data	The Sourcefire 3D System includes several predefined flow data workflows. See Predefined Flow Data Workflows on page 1582 for more information.
Host Attributes	This workflow contains a table view of host attributes and their assigned values for each detected host followed by the host view. See Viewing Host Attributes on page 1034 for more information.
Hosts	This workflow contains a table view of hosts followed by the host view. See Viewing Hosts on page 1020 for more information.

Keysight, EX-1004
Page 1580

Predefined RNA Workflows (Continued)

Workflow Name	Description
Network Services by Count	You can use this workflow to analyze the most frequently used services on your network. This workflow contains a series of pages that show services with a count of hosts where each service occurs, then add the vendor and version of each service. The workflow then concludes with a table view listing the services per host, followed by the host view. See Viewing Services on page 1041 for more information.
Network Services by Hit	You can use this workflow to analyze the most active services on your network. This workflow contains a series of pages that show services with a count of how often each service is accessed, then add the vendor and version information for each service. The workflow finishes with a page containing a table view listing the services per host, followed by the host view. See Viewing Services on page 1041 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, then view 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 1020 for more information.
Services	This workflow contains a table view of services followed by the host view. See Viewing Services on page 1041 for more information.
Vulnerabilities	This workflow contains 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. See Viewing Sourcefire Vulnerabilities on page 1057 for more information.
Third-Party Vulnerabilities	The Sourcefire 3D System includes several predefined third-party vulnerability workflows. See Predefined Third-Party Vulnerabilities Workflows on page 1582 for more information.

Keysight, EX-1004
Page 1581

Predefined Third-Party Vulnerabilities Workflows

Platform: DC

Component: RNA

The Predefined Third-Party Vulnerabilities Workflows table describes the predefined third-party vulnerabilities workflows included on the Defense Center. For information on accessing third-party vulnerabilities workflows, see Viewing Third-Party Vulnerabilities on page 1065.

Predefined Third-Party Vulnerabilities Workflows

Workflow Name	Description
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.
Vulnerabilities by IP Address	You can use this workflow to see quickly how many third-party vulnerabilities you have detected per host on your monitored network. The workflow concludes with a table view of third-party vulnerabilities, followed by the host view.

Predefined Flow Data Workflows

Platform: DC

Component: RNA

The Predefined Flow Data Workflows table describes the predefined flow data workflows included on the Defense Center. All the flow data workflows, with the exception of RNA Flows, include a table view of flow summary data and a table

Keysight, EX-1004
Page 1582

view of flow data. For information on accessing flow data workflows, see Viewing Flow Data on page 1093.

Predefined Flow Data Workflows

Workflow Name	Description
Flow Summaries	This workflow contains a table view of flow summary data and the table view of flow data events.
Flows by Detection Engine	This workflow contains a graph of the 10 most active detection engines on the monitored network segment based on the number of flows detected by each engine.
Flows by Initiator	This workflow contains a graph of the 10 most active hosts on the monitored network segment based on the number of flows where the host initiated the flow transaction.
Flows by Port	This workflow contains a graph of the 10 most active ports on the monitored network segment based on the number of detected flows.
Flows by Responder	This workflow contains a graph of the 10 most active hosts on the monitored network segment based on the number of flows where the host was the responder in the flow transaction.
Flows by Service	This workflow contains a graph of the 10 most active services on the monitored network segment based on the number of detected flows.
Flows over Time	This workflow contains a graph of the total number of flows on the monitored network segment over time.
RNA Flows	This workflow contains only the table view of flow data events.
Traffic by Detection Engine	This workflow contains a graph of the 10 most active detection engines on the monitored network segment based on the total number of kilobytes in the flows detected by each engine.
Traffic by Initiator	This workflow contains a graph of the 10 most active hosts on the monitored network segment based on the total number of kilobytes transmitted by the host.
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.

Keysight, EX-1004
Page 1583

Predefined Flow Data Workflows (Continued)

Workflow Name	Description
Traffic by Responder	This workflow contains a graph of the 10 most active hosts on the monitored network segment based on the total number of kilobytes received by the host.
Traffic by Service	This workflow contains a graph of the 10 most active services on the monitored network segment based on the number of kilobytes transmitted.
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 hosts on the monitored network segment, based on the number of unique initiators that contacted each host.
Unique Responders by Initiator	This workflow contains a graph of the 10 most active initiating hosts on the monitored network segment, based on the number of unique responders that the hosts contacted.

Predefined RUA Workflows

Platform: DC

Component: RUA

The Predefined RUA Workflows table describes the predefined RUA workflows included on the Defense Center.

Predefined RUA Workflows

Workflow Name	Description
Users	This workflow provides a list of user information collected from RUA events or from the LDAP server connection. For details about the users workflow refer to Viewing RUA Users on page 1490.
RUA Events	This workflow provides a list of Real-time User Awareness events. For details about the RUA events refer to Viewing RUA Events on page 1501.

Keysight, EX-1004
Page 1584

Predefined Compliance and White List Workflows

Platform: DC/MDC

Component: RNA

There is a predefined workflow for each type of compliance data, including compliance and white list events, white list violations, and remediation status events.

Predefined Compliance Workflows

Workflow Name	Description
Compliance Events	This workflow, which is also available on the Master Defense Center, contains a table view of compliance events. See Working with Compliance Events on page 1271 for more information.
White List Events	This workflow, which is also available on the Master Defense Center, contains a table view of white list events. See Working with White List Events on page 1197 for more information.
Host Violation Count	This 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 most 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. See Working with White List Violations on page 1204 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 1204 for more information.
Remediation 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 1343 for more information.

Keysight, EX-1004
Page 1585

Predefined System Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

The Sourcefire 3D System is delivered with some additional workflows, including those for system events such as audit events and health events, as well as workflows that list results from SEU imports and active scans.

Additional Predefined Workflows

Workflow Name	Description	Requires
Audit Log	This workflow contains a table view of the audit log listing audit events. See Viewing Audit Records on page 1969 for more information.	None
Health Events	This workflow displays events triggered by the health monitoring policy. See Working with the Health Events Table View on page 1961 for more information.	DC/MDC
SEU Import Log	This workflow contains a table view listing information about both successful and failed SEU imports. For more information, see Importing SEUs and Rule Files on page 1848.	DC/MDC/ Standalone Sensor
Scan Results	The workflow contains a table view listing each completed scan. For more information, see Working with Active Scan Results on page 1443.	DC + RNA

Keysight, EX-1004
Page 1586

Saved Custom Workflows

Platform: DC

Component: IPS/RNA

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 1569.

Saved Custom Workflows

Workflow Name	Description
Events by Impact, Priority, and Host Criticality	Requires DC with IPS + RNA 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. 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 1562.
Events with Destination, Impact, and Host Criticality	Requires DC with IPS + RNA You can use this workflow to find the most recent attacks on hosts that are important to your network and currently vulnerable. 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 1562.
Hosts with Services Default Workflow	Requires DC with RNA You can use this workflow to quickly view the basic information in the Hosts with Services custom table. By default, this workflow begins with a table view of hosts with services, followed by the host view. This workflow is based on the Hosts with Services custom table. For more information, see Understanding Custom Tables on page 1562.

Keysight, EX-1004
Page 1587

Saved Custom Workflows (Continued)

Workflow Name	Description
Intrusion Events with Destination Criticality Default Workflow	Requires DC + IPS + RNA You can use this workflow to quickly view the basic information in the Intrusion Events with Destination Criticality custom table. 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 1562.
Intrusion Events with Source Criticality Default Workflow	Requires DC + IPS + RNA You can use this workflow to quickly view the basic information in the Intrusion Events with Source Criticality custom table. By default, this workflow begins with a table view of hosts with services, followed by the host view. This workflow is based on the Intrusion Events with Source Criticality custom table. For more information, see Understanding Custom Tables on page 1562.
Service and Host Details	Requires DC + RNA You can use this workflow to determine what services are most frequently used on your network and which hosts are running those services. By default, this workflow begins with a summary of services with the frequency of access for each service. The next page lists services by operating system vendor and version. The workflow concludes with a table view of hosts with services, followed by the host view. This workflow is based on the Hosts with Services custom table. For more information, see Understanding Custom Tables on page 1562.

Keysight, EX-1004
Page 1588

Using Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

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 1590 describes the workflow selection page and how to select a workflow to use.
- Understanding the Workflow Toolbar on page 1592 describes the toolbar options available in workflows.
- Using Workflow Pages on page 1593 describes the features that appear on all workflow pages and explains how to use them.
- Setting Event Time Constraints on page 1597 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 1606 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 1609 explains how compound constraints can be used and provides examples.
- Sorting Drill-down Workflow Pages on page 1612 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 1612 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 1613 describes how to open other workflows using the constraints, including any selected events, from the current workflow.
- Navigating between Workflows on page 1614 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 1550 provides information about the feature used to search event data.
- Using Bookmarks on page 1615 describes how to create, manage, and use bookmarks.

Keysight, EX-1004
Page 1589

Selecting Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

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	Requires
Intrusion events	Analysis & Reporting > IPS	Events Reviewed Events Clipboard	DC/MDC/ Standalone Sensor + IPS
RNA-related events (including flow data)	Analysis & Reporting > RNA	Events Hosts Host Attributes Services Client Applications Flow Data Vulnerabilities Third-Party Vulnerabilities	DC + RNA
RUA-related events	Analysis & Reporting > RUA	Users RUA Events	DC + RUA
Compliance events	Analysis & Reporting > Compliance	Compliance Events (also available on the Master Defense Center) White List Events (also available on the Master Defense Center) White List Violations	DC + RNA
Remediation status events	Policy & Response > Responses > Remediations	Status	DC + RNA
Audit events	Operations > Monitoring	Audit	DC/MDC/ 3D Sensor

Keysight, EX-1004
Page 1590

Features Using Workflows (Continued)

Feature	Menu Path	Option	Requires
Health monitoring events	Operations > Monitoring > Health	Health Events	DC/MDC
SEU Import Log	Policy & Response > IPS	SEU	DC/MDC/ Standalone Sensor + IPS
Scan Results	Operations > Tools	Scan Results	DC + RNA

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.

Note that if you are managing a 3D Sensor with a Defense Center and you elected to store data only on the Defense Center and not on the sensor, you will **not** be able to view data in workflows using the sensor's web interface, with the exception of the audit log.

Also note that workflow access depends on your user role (see Configuring User Roles on page 1682), as follows:

- Administrator users can access any workflow, and are the only users who can access the audit log.
- Intrusion Event Analyst users can access intrusion event-related, compliance-related, and RUA-related workflows.
- RNA Event Analyst users can access RNA-related, compliance-related, and RUA-related workflows.
- Policy & Responses Administrator users can access the remediation status workflow and the SEU import log.
- Maintenance users can access health events.
- Anyone can access scan results.

To view the data using a workflow other than the default:

Access: Any except Database

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 45.

Keysight, EX-1004
Page 1591

2. Optionally, use a different workflow. You have two options:
   - From the **Workflows** menu on the toolbar, select the workflow you want to use.
   - On the toolbar, click **Workflows**, then click the name of the workflow that you want to use on the Select Workflow page that appears.

     For example, the following graphic shows the workflows you can use to view intrusion events.



Select Workflow

Name	Description
Destination Port	Drill Down of Destination Port -> Drill Down of Events -> Table View of Events -> Packets
Event-Specific	Drill Down of Events -> Drill Down of Source IPs, or Destination IPs -> Table View of Events -> Packets
Events By Priority and Classification	Drilldown of Event, Priority, and Classification -> Table View Of Events -> Packets
Events to Destinations	Drill Down of Event and Destination -> Table View of Events -> Packets
IP-Specific	Drill Down of Source IPs, or Destination IPs -> Drill Down of Events -> Table View of Events -> Packets
Impact and Priority	Impact Based Event Summary -> Drill Down of Source and Destination IPs -> Table View of Events -> Packets
Impact and Source	Drill Down of Impact and Source -> Drill Down of Source and Destination IPs -> Table View of Events -> Packets
Impact to Destination	Drill Down of Impact and Destination -> Drill Down of Source and Destination IPs -> Table View of Events -> Packets
Source Port	Drill Down of Source Port -> Drill Down of Events -> Table View of Events -> Packets
Source and Destination	Drill Down of Source and Destination IPs -> Drill Down of Events -> Table View of Events -> Packets

Understanding the Workflow Toolbar

Platform: DC/MDC/Standalone Sensor

Component: None

Each page in a workflow includes a toolbar that offers quick access to related features. The Workflow Toolbar Links table table describes each of the links on the toolbar.



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 1615 for information about creating bookmarks.
Report Designer	Opens the report designer with the currently constrained workflow as the selection criteria. See Generating a Basic Report on page 1532 for information about creating reports.

Keysight, EX-1004
Page 1592

Workflow Toolbar Links (Continued)

Feature	Description
 Workflows	Displays the Select Workflow page where you can select a different workflow using the same constraints (if applicable). You can also click the down-arrow icon and select a different workflow from the drop-down menu.
 View Bookmarks	Displays a list of saved bookmarks from which you can select. See Using Bookmarks on page 1615 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 1550 for information about searching workflows.

Using Workflow Pages

Platform: DC/MDC/Standalone Sensor

Component: None

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 constraint 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 1594
- Using Table View Pages on page 1596
- Using Drill-down Pages on page 1596
- Using the Host View, Packet View, or Vulnerability Detail Pages on page 1597

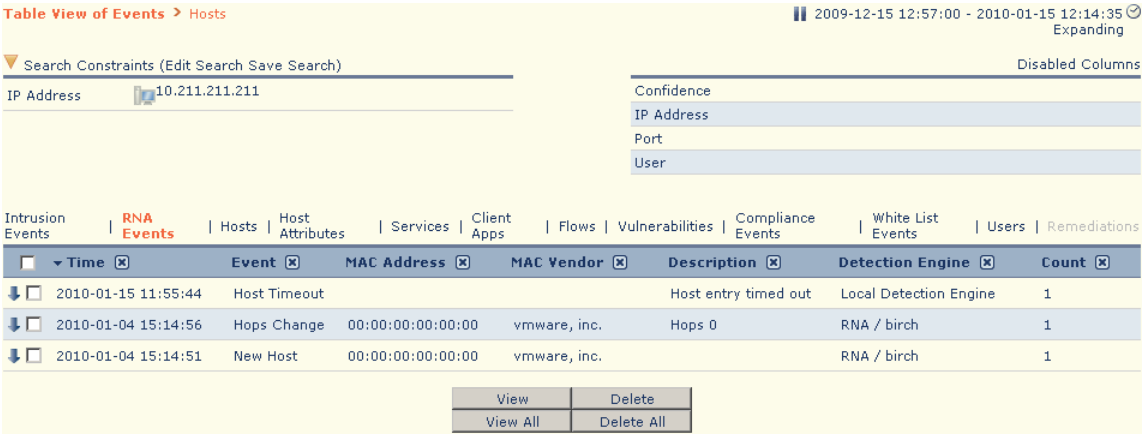
Keysight, EX-1004
Page 1593

Using Common Table View or Drill-down Page Functionality

Platform: DC/MDC/Standalone Sensor

Component: None

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. The Defense Center version of a sample page is shown below.



The features are described in the Table View and Drill-down Page Features table.

Table View and Drill-down Page Features

Feature	Description
	Click this icon to display the corresponding row in the next page of the workflow.
	Click the host profile icon, which appears in IP address columns, to display the host profile for the host. The host profile appears in a pop-up window. For more information, see Using Host Profiles on page 944.
	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 1495.
	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 980.

Keysight, EX-1004
Page 1594

Table View and Drill-down Page Features (Continued)

Feature	Description
Check boxes	Select the check box in 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.
Search Constraints	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. To remove a single constraint, click it. To remove a compound constraint, click **Compound Constraints**. Click **Edit Search** or **Save Search** to open a search page pre-populated with the current single constraints. See Constraining Events on page 1606 for more information. **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.
Time Range	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 1597 for more information. 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.
Workflow Page Links	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.
Workflow Name	The name of the workflow appears at the top of the page.

Keysight, EX-1004
Page 1595

Using Table View Pages

Platform: DC/MDC/Standalone Sensor

Component: None

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 and the page refreshes with the column for that value disabled. 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
	Click this 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**.
Disabled Columns list	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. 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. See Sorting Drill-down Workflow Pages on page 1612 for more information.

Using Drill-down Pages

Platform: DC/MDC/Standalone Sensor

Component: None

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 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 in on events matching your selected values. Clicking a value in a drill-

Keysight, EX-1004
Page 1596

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 1594.

Using the Host View, Packet View, or Vulnerability Detail Pages

Platform: DC/MDC/Standalone Sensor

Component: None

The final page in an RNA event, host, host attributes, services, client applications, or flow 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.

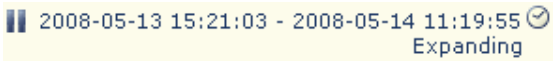
Setting Event Time Constraints

Platform: DC/MDC/Standalone Sensor

Component: None

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 that are 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.



By default, workflows on Sourcefire appliances use an expanding time window set to the past hour. For example, if you log in at 11:30AM you will see events that occurred between 10:30AM and 11:30AM. As time moves forward, the time window expands. At 12:30PM, you will see events that occurred between 10:30AM and 12:30PM.

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)

Keysight, EX-1004
Page 1597

For general information on the default time window, see Default Time Windows on page 47.

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 1598
- Changing the Default Time Window for Your Event Type on page 1603
- Pausing the Time Window on page 1606

Changing the Time Window

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 can 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 RNA hosts, host attributes, services, client applications, vulnerabilities, RUA users, or white list violations.

Keysight, EX-1004
Page 1598

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:

- Requires: IPS or DC/MDC **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
- Requires: DC/MDC **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

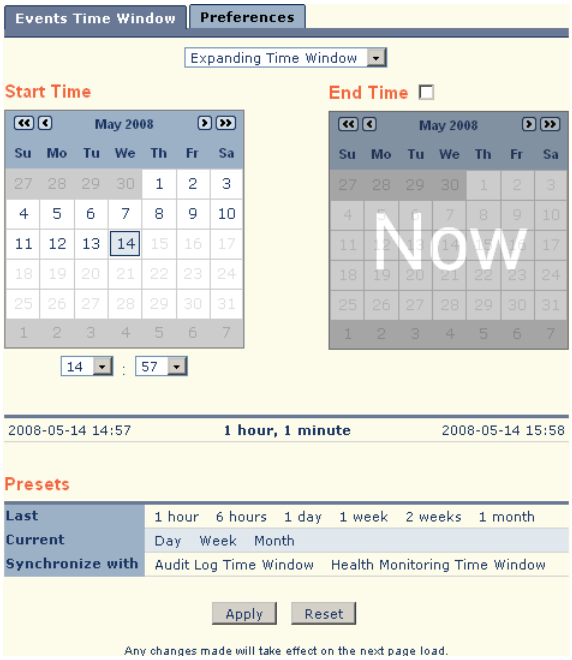
The first decision you must make when configuring a time window is the type of time window you want to use:

- 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).

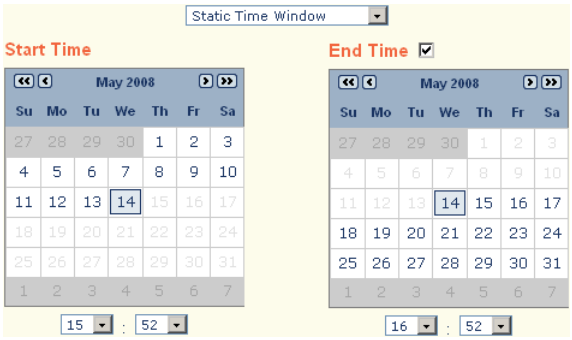
Depending on what type you select, the Date/Time window changes to give you different configuration options. The following graphic shows the Defense Center version of the Date/Time window, specifying that you want to use an expanding

Keysight, EX-1004
Page 1599

time window. With expanding time windows, the End Time calendar is grayed out and specifies that the end time is "Now."

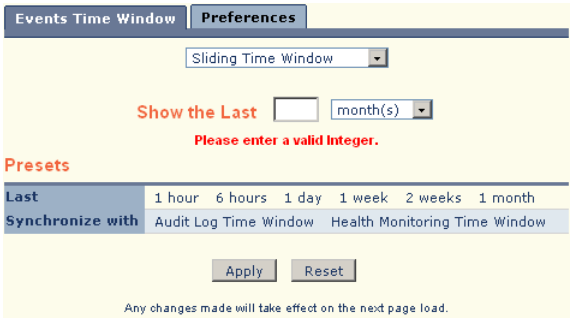


If you use a static time window, you can set an end time.



Keysight, EX-1004
Page 1600

If you choose to use a sliding time window, your options change further.



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 50 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 can 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.

Keysight, EX-1004
Page 1601

Time Window Settings (Continued)

Setting	Time Window Type	Description
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.
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: • 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: • Requires: IPS or DC/MDC **Events Time Window** to synchronize the current time window with the events time window • Requires: DC/MDC **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

Keysight, EX-1004
Page 1602

To change the time window during your event analysis:

Access: Any except Database

1. On a workflow constrained by time, click the time range link.

   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 1601.

   TIP! Click **Reset** to change the time window back to the default settings.

3. Click **Apply**.

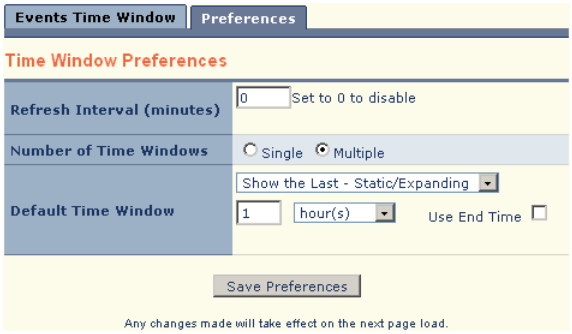
   The window closes and the event view page displays events from the new time range.

Changing the Default Time Window for Your Event Type

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 47).

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.



Keysight, EX-1004
Page 1603

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 you use: • 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	This setting allows you to 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.
Default Time Window: Show the Last - Static/ Expanding	This setting 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.

Keysight, EX-1004
Page 1604

Time Window Preferences (Continued)

Preference	Description
Default Time Window: Current Day - Static/ Expanding	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. 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.
Default Time Window: Current Week - Static/ Expanding	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. 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.

To change time window preferences during your event analysis:

Access: Any except Database

1. On a workflow constrained by time, click the time range link.

   The Date/Time window appears.

2. Select the Preferences tab and change your preferences, as described in the Time Window Preferences table on page 1604.

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**.

Keysight, EX-1004
Page 1605

Pausing the Time Window

Platform: DC/MDC/Standalone Sensor

Component: None

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 can 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 1613). 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: Any except Database

- On the time range control, click the pause icon (❚❚).

  The time window is paused until you unpause it.

To unpause the time window:

Access: Any except Database

- 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

Platform: DC/MDC/Standalone Sensor

Component: None

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. For more information, see Setting Event Time Constraints on page 1597.

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

Keysight, EX-1004
Page 1606

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 1609.

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 **10.14.12.29** in the Source IP column on a page with the following events:

Time	Source IP	SRC Port/ICMP Type
2006-01-10 19:29:15	10.14.12.13	36030/udp
2006-01-10 19:29:15	10.14.12.13	36030/udp
2006-01-10 19:28:59	10.14.12.29	35825/udp
2006-01-10 19:28:59	10.14.12.29	35825/udp
2006-01-10 19:28:54	10.14.10.37	33766/udp
2006-01-10 19:28:54	10.14.10.37	33766/udp

Then the constrained page includes only the events with that IP address:

Search Constraints (Edit Search Save Search)

Source IP	10.14.12.29

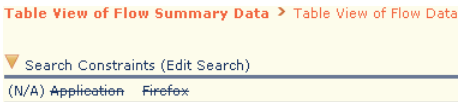
Time	SRC Port/ICMP Type
2006-01-10 19:28:59	35825/udp
2006-01-10 19:28:59	35825/udp
2006-01-10 19:13:59	35824/udp
2006-01-10 19:13:59	35824/udp
2006-01-10 18:58:59	35824/udp
2006-01-10 18:58:59	35824/udp
2006-01-10 18:43:59	35821/udp

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 other to other workflows, unless they are compound constraints (see Navigating between Workflows on page 1614).

Note that because you use the same page in the Defense Center web interface to search both flow summary data and individual flows, you must pay careful attention to whether your search constraints apply to the table you are searching. For example, client application data is not available in flow summaries. If you search for flow events based on the detected client application in the flow and

Keysight, EX-1004
Page 1607

then view the results in a flow summary event view, the Defense Center displays flow data as if you had not constrained it at all. Search constraints that do not apply to the current event view are labeled as not applicable (N/A) and are marked with a strikethrough, as shown in the following graphic.



For your convenience, the flow data search page marks fields that constrain flow summary event views and flow graphs (flow graphs are built using flow summary data) with an asterisk (*).

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. Note that because all of the remaining events share that value, in a table view page the column where you clicked is removed and appears in the Disabled Columns list. When you remove the first column from a page, the count column is added. If you remove a column on a table view page, the column remains disabled until you enable it again. For example, with RNA, if you are viewing a list of services and want to constrain the list to only web services, click **http** in the **Service** column. With IPS, 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 1609.
remove a constraint	the name of the constraint in the Search Constraints box.

Keysight, EX-1004
Page 1608

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 1609.
use the same constraints with another event view	the link at the top of the table that corresponds to the event view, on a Defense Center with RNA. See Navigating between Workflows on page 1614 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 1609.
toggle the display of constraints	the expand icon. This is useful when the list of constraints is large and takes up most of the screen.

Using Compound Constraints

Platform: DC/MDC/Standalone Sensor

Component: None

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

Keysight, EX-1004
Page 1609

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 the Workflows toolbar button 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**.

Keysight, EX-1004
Page 1610

Sorting Table View Pages and Changing Their Layout

Platform: DC/MDC/Standalone Sensor

Component: None

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 1618 for more information.

Sorting and Layout Functions

To...	Click...
sort a column	the column title. Click the column title again to reverse the sort order. **TIP!** The direction icon (▼) indicates which column the data is sorted by, and whether the sort is ascending (upwards-pointing icon) or descending (downwards-pointing icon).
remove a column from a table view	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. You cannot disable the Count column. **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, use the Expand arrow (▶) to expand the search constraints, then click the column name under Disabled Columns.
add a disabled column back to the view	the column name under Disabled Columns. 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.

Keysight, EX-1004
Page 1611

Sorting Drill-down Workflow Pages

Platform: DC/MDC/Standalone Sensor

Component: None

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. The direction icon (▼) indicates which column the data is sorted by, and whether the sort is ascending (upwards-pointing icon) or descending (downwards-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 1618 for more information.

To sort a column:

▶ Click the column title.

To reverse the sort order:

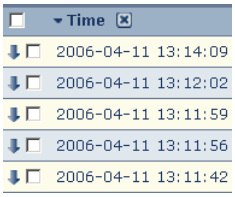
▶ Click the column title again.

Selecting Rows on a Workflow Page

Platform: DC/MDC/Standalone Sensor

Component: None

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.

Keysight, EX-1004
Page 1612

- 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 ( ).

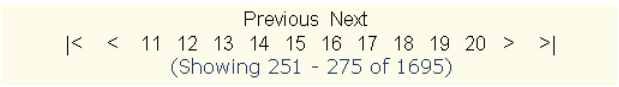
IMPORTANT! You cannot select rows from multiple pages at once.

Navigating to Other Pages in the Workflow

Platform: DC/MDC/Standalone Sensor

Component: None

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.



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 1597.

The Navigating Pages table describes how to use the navigation links.

Navigating Pages

To...	Click...
view the previous page	**Previous**
view the next page	**Next**
view a different page	the page number
jump to the next set of 10 pages	**>**
jump to the previous set of 10 pages	**<**
jump to the last page	**>|**
jump to the first page	**|<**

Keysight, EX-1004
Page 1613

Navigating between Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

You can navigate to other workflows using the links arrayed above the column headings on a workflow page.

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.

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 1597.



The links provide quick access to workflow pages for the following:

- Intrusion Events
- RNA Events
- Hosts
- Host Attributes
- Services
- Client Applications
- Flows
- Vulnerabilities
- Third-Party Vulnerabilities
- Compliance Events
- White List Events
- RUA Users
- Remediation status events (if there are any available)

This feature enhances your ability to investigate suspicious activity. For example, if you are viewing flow 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 **Services**. The services workflow will use the responder IP and port as IP Address and Port constraints and display additional information about the service, such as what kind of service

Keysight, EX-1004
Page 1614

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 service, you can click **Flow Data** to return to the flow data workflow, remove the Responder IP from the constraints, add the Initiator IP to constraints, and click **Client Applications** to see what application the user on the initiating host used when transferring data to the remote host. Note that the Port constraint is not transferred to the Client Application 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 **Compliance 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

Platform: DC/MDC/Standalone Sensor

Component: None

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 unrestricted data 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 event analyst.

IMPORTANT! If the events that appear in a bookmark are deleted (either directly by an event analyst or by automatic database clean-up), the bookmark no longer displays the original set of events.

See these sections for more information about using bookmarks:

- Creating Bookmarks on page 1616 describes how to create a new bookmark.

Keysight, EX-1004
Page 1615

- Viewing Bookmarks on page 1616 describes how to view and use existing bookmarks.
- Deleting Bookmarks on page 1617 describes how to delete bookmarks.

Creating Bookmarks

Platform: DC/MDC/Standalone Sensor

Component: None

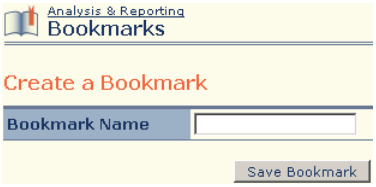
Use the following procedure to create a new bookmark.

To create a bookmark:

Access: Any except Database

1. During an event analysis, with the events of interest displayed, click **Bookmark This Page**.

   The Bookmarks page appears.



2. Type a name (up to 80 alphanumeric characters and spaces) for the bookmark and click **Save Bookmark**.

   The bookmark is saved and the event page you bookmarked appears again.

Viewing Bookmarks

Platform: DC/MDC/Standalone Sensor

Component: None

Use the following procedure to view and use existing bookmarks.

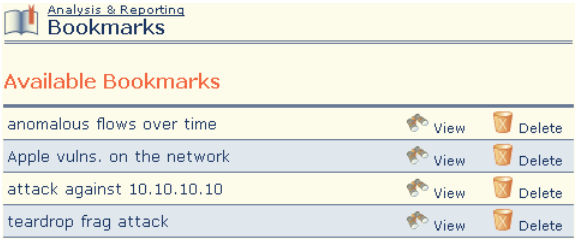
Keysight, EX-1004
Page 1616

To view a bookmark:

Access: Any except Database

1. From any event view, click **View Bookmarks**, or select **Analysis & Reporting > Bookmarks**.

   The Bookmarks page appears. The Defense Center version of the page is displayed below.



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 an event analyst or by automatic database clean-up), the bookmark no longer displays the original set of events.

Deleting Bookmarks

Platform: DC/MDC/Standalone Sensor

Component: None

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: Any except Database

1. From any event view, click **View Bookmarks**, or select **Analysis & Reporting > Bookmarks**.

   The Bookmarks page appears.

2. Click **Delete** next to the bookmark you want to remove.

   The bookmark is deleted.

Keysight, EX-1004
Page 1617

Using Custom Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

If the predefined and Sourcefire-provided custom workflows do not meet your needs, you can create custom workflows. Note that you **cannot** create or view custom workflows on a 3D Sensor that does not have an IPS license.

For more information, see:

- Creating Custom Workflows on page 1618 for a procedure to create custom workflows
- Creating Custom Flow Data Workflows on page 1621 for a procedure to create a custom workflow based on flow data
- Viewing Custom Workflows on page 1623 for procedures for viewing custom workflows based on event and custom tables
- Editing Custom Workflows on page 1624 for a procedure for editing custom workflows
- Deleting Custom Workflows on page 1625 for a procedure for deleting custom workflows

Creating Custom Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

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 1997.

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

Keysight, EX-1004
Page 1618

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, or 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...	Requires
RNA event tables	host profile	DC + RNA
vulnerabilities	vulnerability detail	DC + RNA
third-party vulnerabilities	host profile	DC + RNA
RUA users or RUA events	user detail	DC + RUA
intrusion events	packet view	IPS or DC/MDC

The appliance does not add a final page to custom workflows based on other kinds of events (for example, audit log events).

IMPORTANT! The procedure for creating a custom workflow based on flow data is slightly different. For more information, see the next section, Creating Custom Flow Data Workflows.

To create a custom workflow:

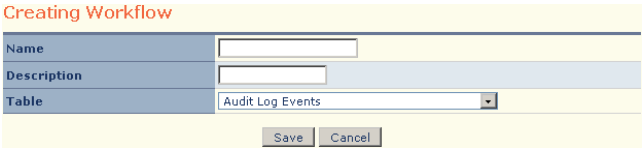
Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Custom Workflows**.

   The Custom Workflows page appears.

2. Click **Create Custom Workflow**.

   The Edit Custom Workflow page appears.



3. Type a name for the workflow in the **Name** field.

   You can use up to 60 alphanumeric characters and spaces in the name.

Keysight, EX-1004
Page 1619

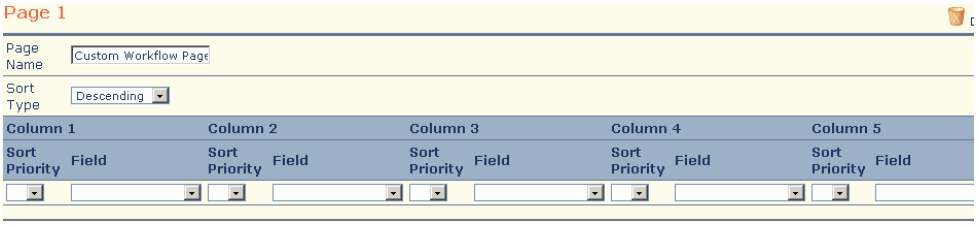
4. Optionally, type a description for the workflow in the **Description** field.

   You can use up to 80 alphanumeric characters and spaces.

5. Select the table you want to include from the **Table** drop-down list.

6. Optionally, click **Add Page** to add one or more drill-down pages to the workflow.

   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. 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! On a Defense Center with RNA, 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 1062.

7. 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.

8. Click **Save**.

   The new workflow is saved and added to the list of custom workflows.

Keysight, EX-1004
Page 1620

Creating Custom Flow Data Workflows

Platform: DC

Component: RNA

Custom workflows based on flow data are like other custom workflows, except you can include flow 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 flow 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 flow data, including flow summaries, flow graphs, and datasets, see Understanding Flow Data on page 1072.

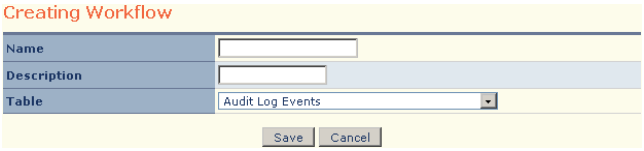
TIP! Instead of creating a new custom workflow, you can export a custom workflow from another appliance and then import it onto your appliance. You can then edit the imported workflow to suit your needs. For more information, see Importing and Exporting Objects on page 1997.

To create a custom workflow based on flow data:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Custom Workflow**.
2. Click **Create Custom Workflow**.

   The Edit Custom Workflow page appears.



3. Type a name for the workflow in the **Name** field.

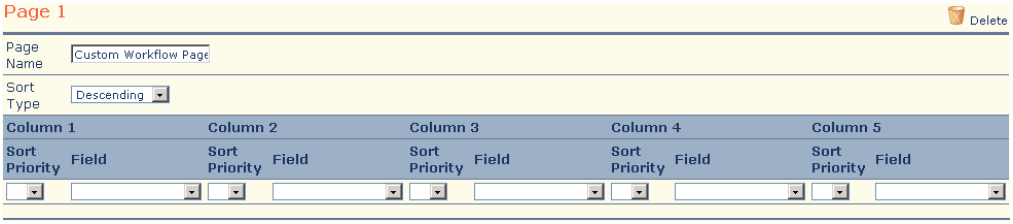
   You can use up to 60 alphanumeric characters and spaces in the name.
4. Optionally, type a description for the workflow in the **Description** field.

   You can use up to 80 alphanumeric characters and spaces.
5. From the **Table** drop-down list, select **Flow Data**.

Keysight, EX-1004
Page 1621

6. Optionally, add one or more drill-down pages to the workflow.
   - To add a drill-down page that contains data on individual flows, click **Add Page**.
   - To add a drill-down page that contains flow 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.

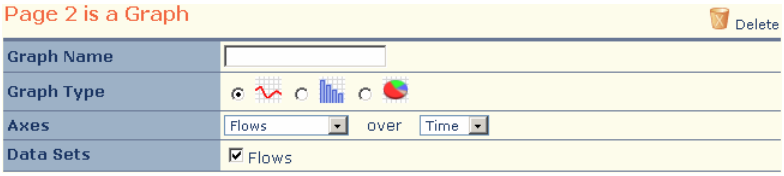
   Under **Column 1**, select a sort priority and a table column. This column will appear in the leftmost column of the page.

   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.

Keysight, EX-1004
Page 1622

8. Optionally, add a table view to the workflow.
   - To add a table view of individual flows, click **Add Table View**.
   - To add a table view of flow summaries, click **Add Summary Table View**.
9. Click **Save**.

   The new workflow is saved and added to the list of custom workflows.

Viewing Custom Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

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 use the Defense Center to access a custom workflow based on the RNA Hosts table, select **Analysis & Reporting > RNA > Hosts**. If, on the other hand, you are using a Defense Center and 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 45.

For more information, see:

- Viewing Custom Workflows for Predefined Tables on page 1623
- Viewing Custom Workflows for Custom Tables on page 1624

Viewing Custom Workflows for Predefined Tables

Platform: DC/MDC/Standalone Sensor

Component: None

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 1590.

Keysight, EX-1004
Page 1623

To view a custom workflow based on a predefined table:

Access: Any except Database

- 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. 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 1597.

Viewing Custom Workflows for Custom Tables

Platform: DC

Component: RNA

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: Any Analyst/Admin

1. Select **Analysis & Reporting > Custom Tables**.

   The Custom Tables page appears, listing the available custom tables.

2. Click **View** next to the custom table you want to view.

   The first page of the default workflow for that table appears. To use a different workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. 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 1597.

Editing Custom Workflows

Platform: DC/MDC/Standalone Sensor

Component: None

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: Any Analyst/Admin

1. Select **Analysis & Reporting > Custom Workflows**.

   The Custom Workflows page appears, listing the existing custom workflows.

Keysight, EX-1004
Page 1624

2. Click **Edit** 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

Platform: DC/MDC/Standalone Sensor

Component: None

The following procedure explains how to delete a custom workflow that you no longer need.

To delete a custom workflow:

Access: Any Analyst/Admin

1. Select **Analysis & Reporting > Custom Workflows**.
   The Custom Workflows page appears, listing the available custom workflows.
2. Click **Delete** next to the name of the workflow that you want to delete.
   The workflow is deleted.

Keysight, EX-1004
Page 1625

Chapter 38
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 3D Sensor. 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 1626
- Managing Authentication Objects on page 1631
- Managing User Accounts on page 1675

Understanding Sourcefire User Authentication

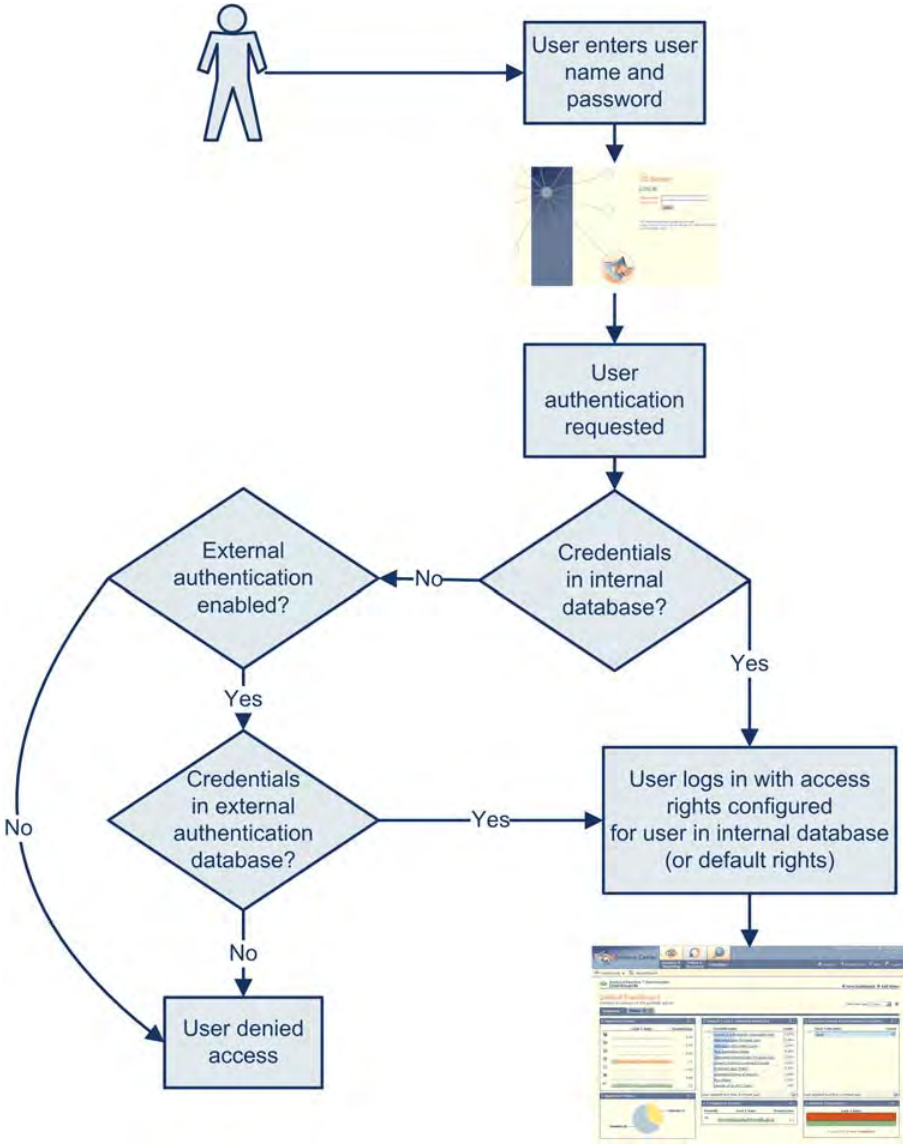
Platform: DC/MDC/3D Sensor

Component: None

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

Keysight, EX-1004
Page 1626

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.

Keysight, EX-1004
Page 1627

For more information, see the following sections:

- Understanding Internal Authentication on page 1628
- Understanding External Authentication on page 1628
- Understanding User Privileges on page 1629

Understanding Internal Authentication

Platform: DC/MDC/3D Sensor

Component: None

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. Once an internally authenticated user converts to an externally authenticated user, you cannot revert to internal authentication for that user.

Understanding External Authentication

Platform: DC

Component: None

External authentication occurs when the Defense Center or managed sensor 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.

Keysight, EX-1004
Page 1628

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 3D Sensors.

You can push a system policy to a managed 3D Sensor to enable external authentication on that sensor, but you cannot control the authentication object from the sensor's web interface. The only configuration of external authentication on the sensor occurs when you select the type of authentication for a new user. If you want to disable external authentication on a managed 3D Sensor, disable it in the system policy on the managing Defense Center and re-apply the policy to the sensor. If you apply a local system policy (created on the sensor) to the sensor itself, external authentication is also disabled.

IMPORTANT! Sourcefire does not support external authentication for RNA Software for Red Hat Linux, Intrusion Agents, 3Dx800 sensors, or Crossbeam-based software sensors.

For more information on specific types of external authentication, see the following sections:

- Understanding LDAP Authentication on page 1631
- Understanding RADIUS Authentication on page 1662

Understanding User Privileges

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 Intrusion Event Analyst and RNA Event Analyst access privileges for analysts and reserve the Administrator role for the network administrator managing the Sourcefire 3D System.

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

Keysight, EX-1004
Page 1629

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 user roles, listed in order of precedence, depending on the features you have licensed:

- *Administrators* can set up the appliance's network configuration, manage user accounts, configure system policies and system settings. Users with the Administrator role also have Intrusion Event Analyst, RNA Event Analyst, Policy & Response (P&R) Administrator, Maintenance, and External Database access rights.
- *Intrusion Event Analysts* can view, analyze, review, and delete intrusion events and compliance and RUA events. They can also create incidents, generate reports, and view (but not delete or modify) health events.
- *Intrusion Event Analysts (Read Only)* have all the same rights as Intrusion Event Analysts, except that they cannot delete events.
- *RNA Event Analysts* can view, analyze, and delete network change events, hosts, host attributes, services, vulnerabilities, client applications, compliance events, and RUA events. RNA analysts can also generate reports and view (but not delete or modify) health events.
- *RNA Event Analysts (Read Only)* have all the same rights as RNA Event Analysts, except that they cannot delete events.
- *Restricted Event Analysts* have the combined privileges of Intrusion Event Analysts and RNA Event Analysts, but users are limited to subsets of that data. Restricted analysts can also be assigned the Policy & Response Administrator or Maintenance User roles, but cannot be assigned the Intrusion Event Analyst or RNA Event Analyst roles.
- Note that on the Defense Center you cannot select Restricted Event Analyst as the default user role in the system policy, but you can modify a user's settings via the User Management page to grant this level of access.
Policy & Response Administrators can manage intrusion rules, policies, and responses, as well as compliance rules, policies, and responses.

Keysight, EX-1004
Page 1630

- *Maintenance Administrators* 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 administrators do not have access to the functions in the Policy & Response menu and can only access the dashboard from the Analysis & Reporting menu.
- *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.

Managing Authentication Objects

Platform: DC/MDC

Component: None

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 1631
- Preparing to Create an LDAP Authentication Object on page 1636
- Quick Start to LDAP Authentication on page 1637
- Tuning Your LDAP Authentication Connection on page 1640
- Creating Advanced LDAP Authentication Objects on page 1642
- LDAP Authentication Object Examples on page 1657
- Editing LDAP Authentication Objects on page 1661
- Creating RADIUS Authentication Objects on page 1662
- RADIUS Authentication Object Examples on page 1671
- Editing RADIUS Authentication Objects on page 1674
- Deleting Authentication Objects on page 1674

Understanding LDAP Authentication

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.

Keysight, EX-1004
Page 1631

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, Sun Directory Services 5.2 P4 on Windows 2003 and XP, or OpenLDAP on Linux. However, Sourcefire does not support external authentication for RNA Software for Red Hat Linux, Intrusion Agents, 3Dx800 sensors, or Crossbeam-based software sensors.

For more information, see the following sections:

- Setting Defaults on page 1632
- Setting a Base DN on page 1633
- Setting a Base Filter on page 1633
- Selecting an Impersonation Account on page 1633
- Encrypting Your LDAP Connection on page 1633
- Setting the User Name Template on page 1634
- Setting a Connection Timeout on page 1634
- Using Attributes to Manage Access on page 1634
- Using Group Membership to Manage Access on page 1635
- Setting up Shell Access on page 1635
- Testing the Connection on page 1636

Setting Defaults

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

Keysight, EX-1004
Page 1632

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

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`.

Once 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

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 1656.

Selecting an Impersonation Account

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

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.

Keysight, EX-1004
Page 1633

Setting the User Name Template

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

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

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, 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

Keysight, EX-1004
Page 1634

are unique, can include underscores (_), periods (.), hyphens (-), and alphanumeric characters.

Using Group Membership to Manage Access

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

You can use the LDAP server to authenticate accounts for shell access on a 3D Sensor 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 1741.

With the exception of the admin account, shell access is controlled entirely though 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/password` 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.

Keysight, EX-1004
Page 1635

Shell users should log in using user names with all lowercase letters.

WARNING! On Series 2 appliances and 8000 Series Defense Centers, all shell users have `sudoers` privileges. Make sure that you restrict the list of users with shell access appropriately. On 8000 Series sensors and Virtual 3D Sensors, shell access granted to externally authenticated users defaults to the **Configuration** level of command line access, which also grants `sudoers` privileges.

Testing the Connection

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

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 1631.

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.

Keysight, EX-1004
Page 1636

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

Platform: DC/MDC

Component: None

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

Keysight, EX-1004
Page 1637

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 1642 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 1631.

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 **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.
2. Click **Create Authentication Object**.
3. Select **LDAP** from the **Authentication Method** drop-down list.
4. Type a name and description for the authentication server in the **Name** and **Description** fields.
5. Set the server defaults for the server. You have the following options:
   - Select **MS Active Directory** and click **Set Defaults** if you are connecting to a Microsoft Active Directory server.
   - Select **Sun Directory** and click **Set Defaults** if you are connecting to a Sun Java Systems Directory Server or Oracle Directory Server.

Keysight, EX-1004
Page 1638

- Select **OpenLDAP** and click **Set Defaults** if you are connecting to an OpenLDAP server.
- Select **Other** and click **Set Defaults** if you are connecting to a server other than those listed above and want to clear default settings.

6. 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.

7. To fetch a list of all base DNs, click **Fetch DNs** 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`.

8. 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)`.

9. 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.`

10. Re-type the password in the **Confirm Password** field.

11. 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.

Keysight, EX-1004
Page 1639

12. 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 you can retrieve the `JSmith` user credentials at the Example company, type `JSmith.`

13. 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.

14. 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 1741 and Applying a System Policy on page 1736.

    - 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

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.

Keysight, EX-1004
Page 1640

  - 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 hostname 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 1644.
- 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 1633 and Setting up Shell Access on page 1635.
- 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.

Keysight, EX-1004
Page 1641

- 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 1633
- Setting a Base Filter on page 1633
- Configuring LDAP-Specific Parameters on page 1645

Creating Advanced LDAP Authentication Objects

Platform: DC/MDC

Component: None

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

Keysight, EX-1004
Page 1642

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:

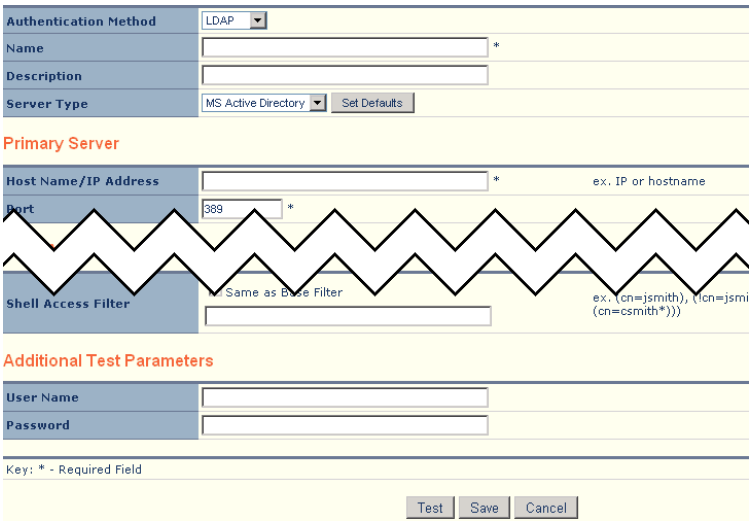
Access: Admin

1. Select **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.

2. Click **Create Authentication Object**.

   The Create Authentication Object page appears.



3. 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 1644.

4. 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 1645.

5. 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 1651.

Keysight, EX-1004
Page 1643

6. Optionally, configure authentication settings for shell access. For more information, see Configuring Administrative Shell Access on page 1654.

7. 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 1656.

   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 1741 and Applying a System Policy on page 1736.

Identifying the LDAP Authentication Server

Platform: DC/MDC

Component: None

When you create an authentication object, you first specify the primary and backup server and server port where you want the 3D Sensor or Defense Center to connect for authentication.

To identify an LDAP authentication server:

Access: Admin

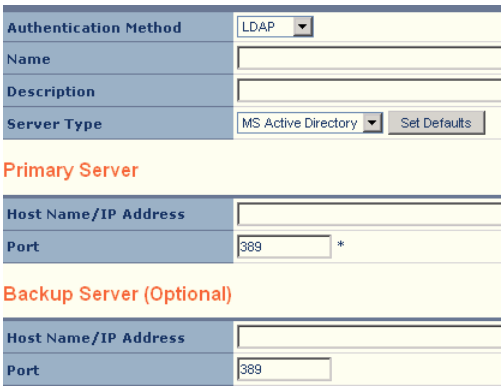
1. Select **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.

2. Click **Create Authentication Object**.

   The Create Authentication Object page appears.

3. Select **LDAP** from the **Authentication Method** drop-down list.



4. Type a name and description for the authentication server in the **Name** and **Description** fields.

Keysight, EX-1004
Page 1644

5. Optionally, 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 **Sun 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**.
6. 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.

7. Optionally, modify the port used by the primary authentication server in the **Primary Server Port** field.
8. 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.
9. 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

Platform: DC/MDC

Component: None

The settings in the LDAP-specific parameters section determine the area of the LDAP directory where the 3D Sensor or Defense Center searches for user names, and control details of how the 3D Sensor or Defense Center connects to the LDAP server.

When configuring these settings, note that valid user names are unique, 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.

Keysight, EX-1004
Page 1645

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	Supplies the base distinguished name of the directory where the appliance searches for user information on the LDAP server. Typically, the base DN has a basic structure indicating the company domain and operational unit. Note that once you identify a primary server, you can automatically retrieve a list of available base DNs from the server and select the appropriate base DN.	The Security organization of the Example company might have a base DN of `ou=security, dc=example,dc=com`
Base Filter	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. To test your base filter more specifically by entering a test user name and password, see Testing User Authentication on page 1656.	To filter for only users with a common name starting with F, use the filter `(cn=F*)`.
User Name/ Password	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.	The user name for the `admin` user in the Security organization of the Example company might have a user name of `cn=admin, ou=security, dc=example,dc=com`

Keysight, EX-1004
Page 1646

LDAP-Specific Parameters (Continued)

Setting	Description	Example
Encryption	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. If you change the encryption method after specifying the port, the port resets to the default value for the selected server type.	If you enter `10.10.10.250` in the authentication profile and `computer1.example.com` in the certificate, the connection fails, even if `computer1.example.com` has an IP address of `10.10.10.250`. Changing the name of the server in the authentication profile to `computer1.example.com` causes the connection to succeed.
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`.
Timeout	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. 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.	If the primary server has LDAP disabled, the appliance queries the backup server.

Keysight, EX-1004
Page 1647

LDAP-Specific Parameters (Continued)

Setting	Description	Example
UI Access Attribute	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. Selecting a server type and setting defaults prepopulates the **UI Access Attribute** with a value typically appropriate for that type of server. 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.	`sAMAccountName`
Shell Access Attribute	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. If you leave this field blank, the user distinguished name is used for shell access authentication. Note that selecting a server type and setting defaults prepopulates this field with an attribute typically appropriate for that type of server.	`sAMAccountName`

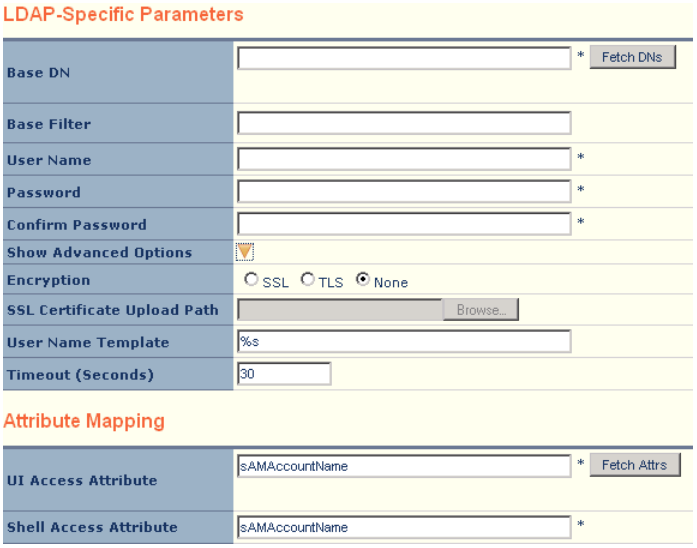
Keysight, EX-1004
Page 1648

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`.



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`.

Keysight, EX-1004
Page 1649

4. Re-type 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 1651.
   - If you are not using LDAP groups for authentication, continue with Configuring Administrative Shell Access on page 1654.
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 re-apply 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.

Keysight, EX-1004
Page 1650

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 1654.

Configuring Access Settings by Group

Platform: DC/MDC

Component: None

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.

Keysight, EX-1004
Page 1651

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.

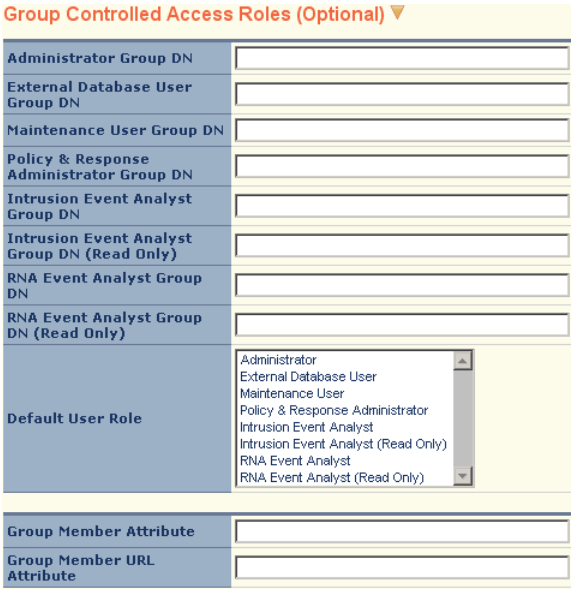
Keysight, EX-1004
Page 1652

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.



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 Group DN** 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 1676.

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.

Keysight, EX-1004
Page 1653

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 Policy & Response Administrator 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

Platform: DC/MDC

Component: None

You can also use the LDAP server to authenticate accounts for shell access on your 3D Sensor 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 1741.

IMPORTANT! Sourcefire does not support external authentication for RNA Software for Red Hat Linux, Intrusion Agents, 3Dx800 sensors, or Crossbeam-based software sensors. In addition, IPv6 is not supported for shell access authentication.

With the exception of the admin account, shell access is controlled entirely though 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/password` to disable the shell. If the user then is re-enabled, the shell is reset, using the same home directory.

Keysight, EX-1004
Page 1654

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.

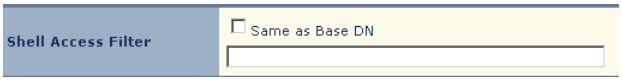
Shell users should log in using user names with all lowercase letters.

WARNING! On Series 2 appliances and 8000 Series Defense Centers, all shell users have `sudoers` privileges. Make sure that you restrict the list of users with shell access appropriately. On 8000 Series sensors and Virtual 3D Sensors, 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.



   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.

Keysight, EX-1004
Page 1655

Testing User Authentication

Platform: DC/MDC

Component: None

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, 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 UI page size limitations.

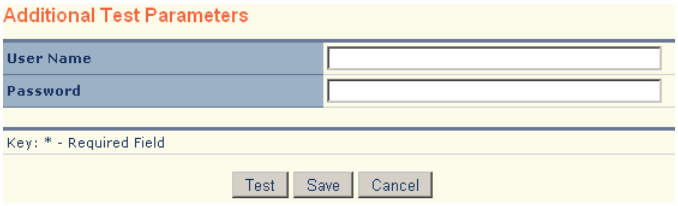
TIP! If you mistype the name or password of the test user, the test fails even if the server configuration is correct. Test the server configuration without the additional test parameters first. If that succeeds supply a user name and password to test with the specific user.

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 you can retrieve the `JSmith` user credentials at the Example company, type `JSmith.`



Keysight, EX-1004
Page 1656

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 1741 and Applying a System Policy on page 1736.
   - If the test fails, see Tuning Your LDAP Authentication Connection on page 1640 for suggestions for troubleshooting the connection. Note that the error message that appears indicates what caused the connection to fail.

LDAP Authentication Object Examples

Platform: DC/MDC

Component: None

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 1657
- Example: Advanced LDAP Configuration on page 1659

Example: Basic LDAP Configuration

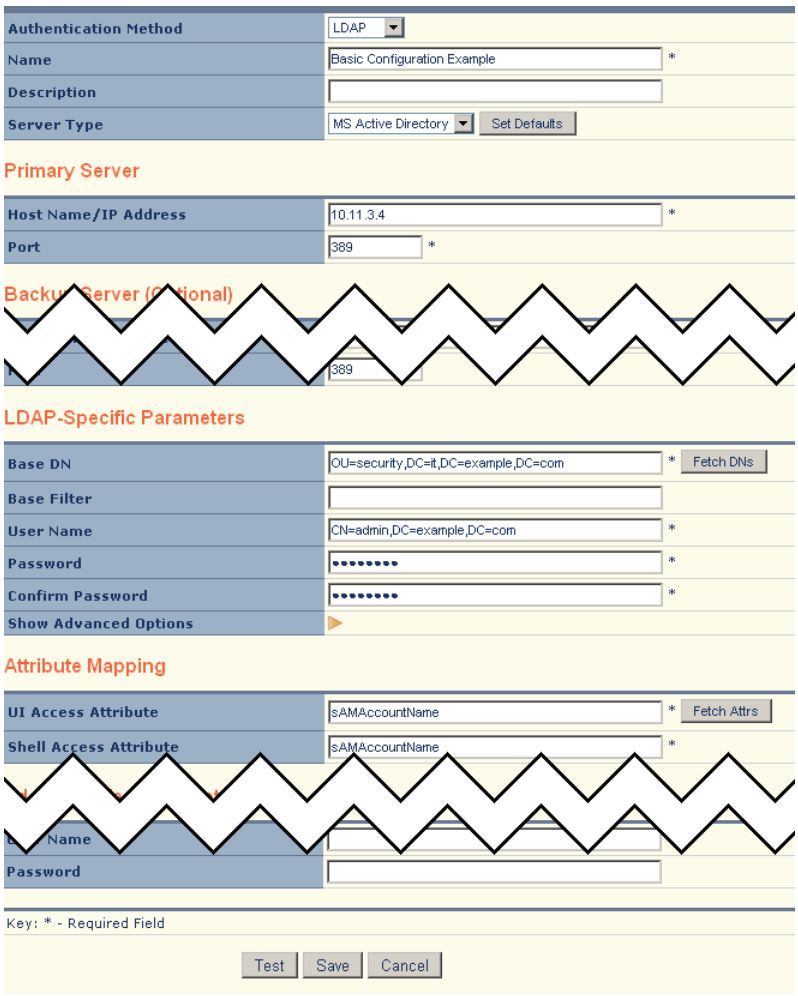
Platform: DC/MDC

Component: None

The following figure illustrates a basic configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in

Keysight, EX-1004
Page 1657

this example has an IP address of 10.11.3.4. The connection uses port 389 for access.



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.

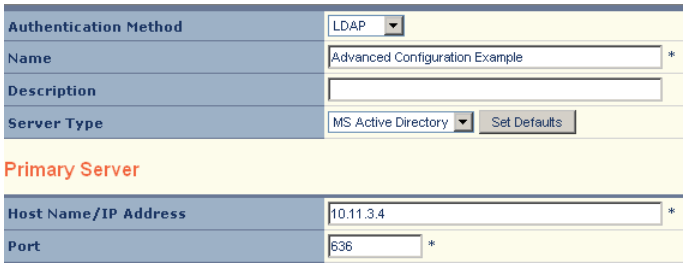
Keysight, EX-1004
Page 1658

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.

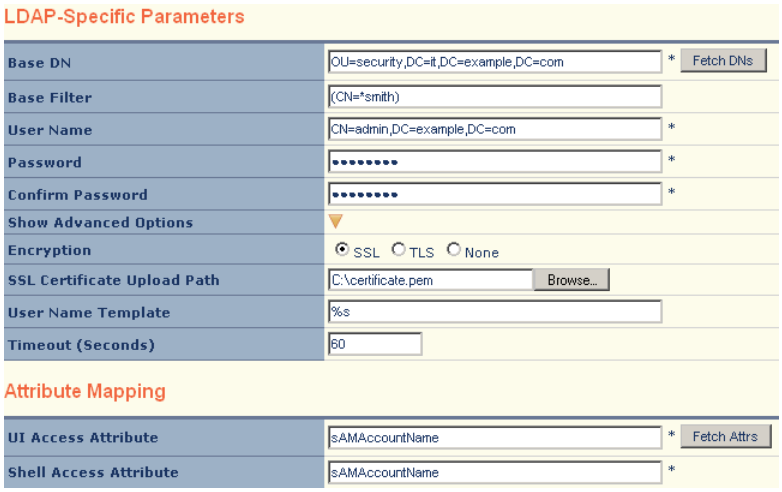
Note that because no base filter is applied to this server, the Sourcefire 3D System checks attributes for all objects in the directory indicated by the base distinguished name. Connections to the server time out after the default time period (or the timeout period set on the LDAP server).

Example: Advanced LDAP Configuration

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.



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`.



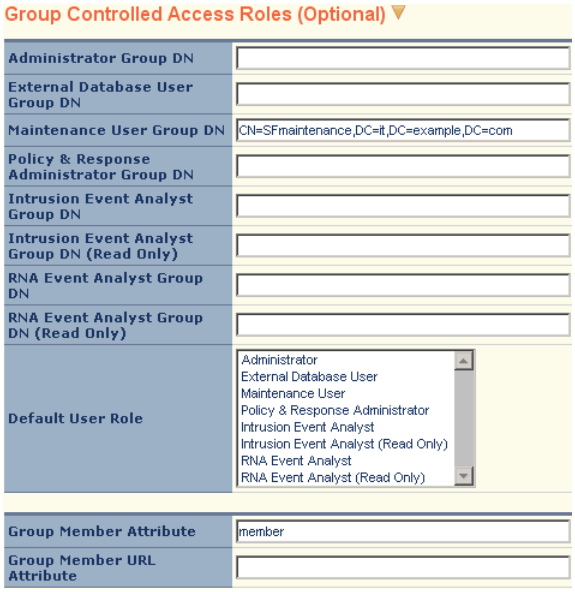
Keysight, EX-1004
Page 1659

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.

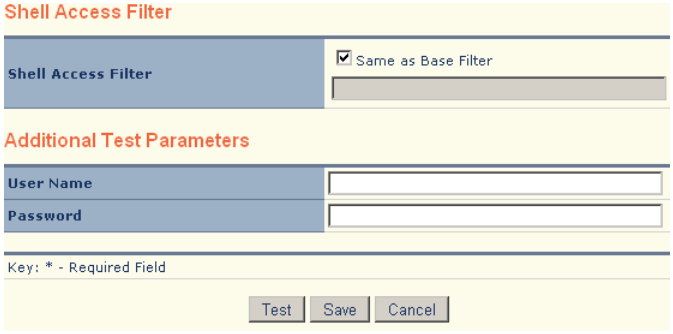
In addition, a **Shell Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a shell account on the appliance.

This example also has group settings in place. The maintenance 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`.



Keysight, EX-1004
Page 1660

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.



Editing LDAP Authentication Objects

Platform: DC/MDC

Component: None

You can edit an existing authentication object. Your changes do not take effect until you re-apply the policy.

To edit an authentication object:

Access: Admin

1. Select **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.

2. Click **Edit** next to the object you want to edit.

   The Create Authentication Object page appears.

3. Modify the object settings as needed.

   For more information, see the following topics:

   - Quick Start to LDAP Authentication on page 1637
   - Creating Advanced LDAP Authentication Objects on page 1642
   - Identifying the LDAP Authentication Server on page 1644
   - Configuring LDAP-Specific Parameters on page 1645
   - Configuring Access Settings by Group on page 1651
   - Configuring Administrative Shell Access on page 1654
   - Testing User Authentication on page 1656

Keysight, EX-1004
Page 1661

4. 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 1640 for suggestions for troubleshooting the connection. Note that the error message that appears indicates what caused the connection to fail.

5. 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 1741 and Applying a System Policy on page 1736.

Understanding RADIUS Authentication

Platform: DC/MDC

Component: None

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.

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

Platform: DC/MDC

Component: None

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

Keysight, EX-1004
Page 1662

default users. If your RADIUS server returns custom attributes for any users you plan to authenticate, you need to 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 **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.

2. Click **Create Authentication Object**.

   The Create Authentication Object page appears.

3. 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 1663.

4. 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 1665.

5. Optionally, configure administrative shell access. For more information, see Configuring Administrative Shell Access on page 1668.

6. 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 1669.

7. 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 1670.

   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 1741 and Applying a System Policy on page 1736.

Configuring RADIUS Connection Settings

Platform: DC/MDC

Component: None

When you create a RADIUS authentication object, you first specify the primary and backup server and server port where you want the local appliance (3D Sensor or Defense Center) to connect for authentication.

Keysight, EX-1004
Page 1663

For FreeRADIUS to function correctly, you need to open both ports 1812 and 1813 on your firewall and on the FreeRADIUS server.

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.

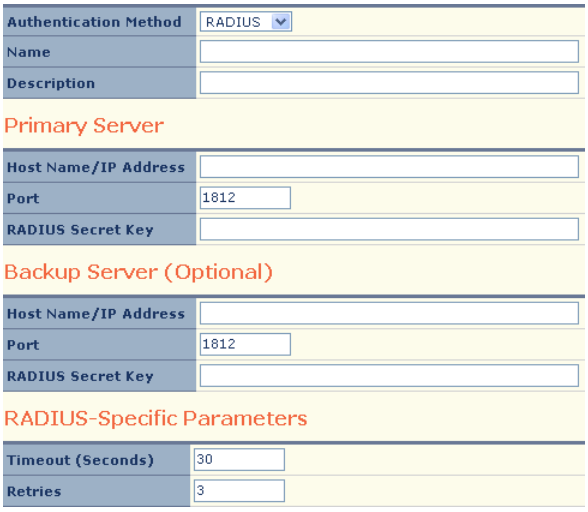
If, for example, the primary server has RADIUS disabled, the appliance queries the backup server. If RADIUS is running on the port of the primary RADIUS server and for some reason refuses to service the request (due to misconfiguration or other issues), however, the failover to the backup server does not occur.

To identify a RADIUS authentication server:

Access: Admin

1. Select **RADIUS** from the **Authentication Method** drop-down list.

   RADIUS options appear.



2. Type a name and description for the authentication server in the **Name** and **Description** fields.

3. 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.

4. Optionally, modify the port used by the primary RADIUS authentication server in the **Primary Server Port** field.
5. Type the secret key for the primary RADIUS authentication server in the **RADIUS Secret Key** field.
6. 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.
7. Optionally, modify the port used by the backup RADIUS authentication server in the **Backup Server Port** field.
8. Type the secret key for the backup RADIUS authentication server in the **RADIUS Secret Key** field.
9. Type the number of seconds that should elapse before retrying the connection in the **Timeout** field.
10. Type the number of times the primary server connection should be tried before rolling over to the backup connection in the **Retries** field.
11. Continue with Configuring RADIUS User Roles.

Configuring RADIUS User Roles

Platform: DC/MDC

Component: None

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.

Keysight, EX-1004
Page 1665

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 RNA Analysts have the value `Analyst` for their `User-Category` attribute, you can type `User-Category=Analyst` in the RNA 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 1669.

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 1682.

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.

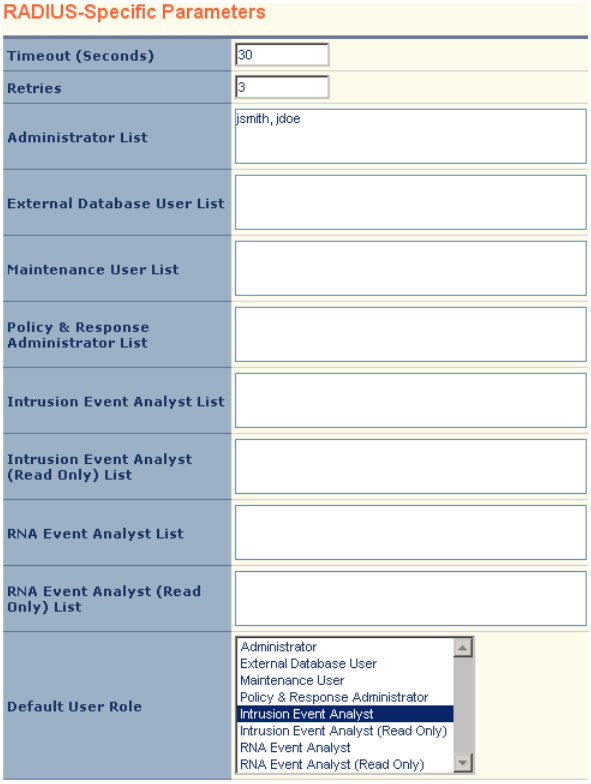
Keysight, EX-1004
Page 1666

To base access on user lists:

Access: Admin

1. In the **List** 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 List** field.



   As another example, to grant the Maintenance role to all users with a `User-Category` value of `Maintenance`, type `User-Category=Maintenance` in the **Maintenance List** field.

   For more information on user access roles, see Configuring User Roles on page 1682.

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.

Keysight, EX-1004
Page 1667

Configuring Administrative Shell Access

Platform: DC/MDC

Component: None

You can also use the RADIUS server to authenticate accounts for shell access on your local appliance (3D Sensor 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 1741.

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 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/password` 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 2 appliances and 8000 Series Defense Centers, all shell users have `sudoers` privileges. Make sure that you restrict the list of users with shell access appropriately. On 8000 Series sensors and Virtual 3D Sensors, 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 1669.

Keysight, EX-1004
Page 1668

Defining Custom RADIUS Attributes

Platform: DC/MDC

Component: None

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 **Intrusion Event Analyst (Read Only)** field to grant read-only intrusion event 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.



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.

Keysight, EX-1004
Page 1669

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 1670.

Testing User Authentication

Platform: DC/MDC

Component: None

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.

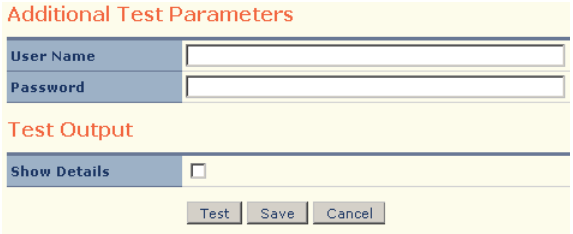
> **TIP!** If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click Test without entering user information in the Additional Test Parameters first. If that succeeds supply a user name and 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 you can retrieve the `jsmith` user credentials at our example company, type `jsmith`.



Keysight, EX-1004
Page 1670

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 1741 and Applying a System Policy on page 1736.

RADIUS Authentication Object Examples

Platform: DC/MDC

Component: None

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 1671
- Authenticating a User with Custom Attributes on page 1672

Authenticating a User using RADIUS

Platform: DC/MDC

Component: None

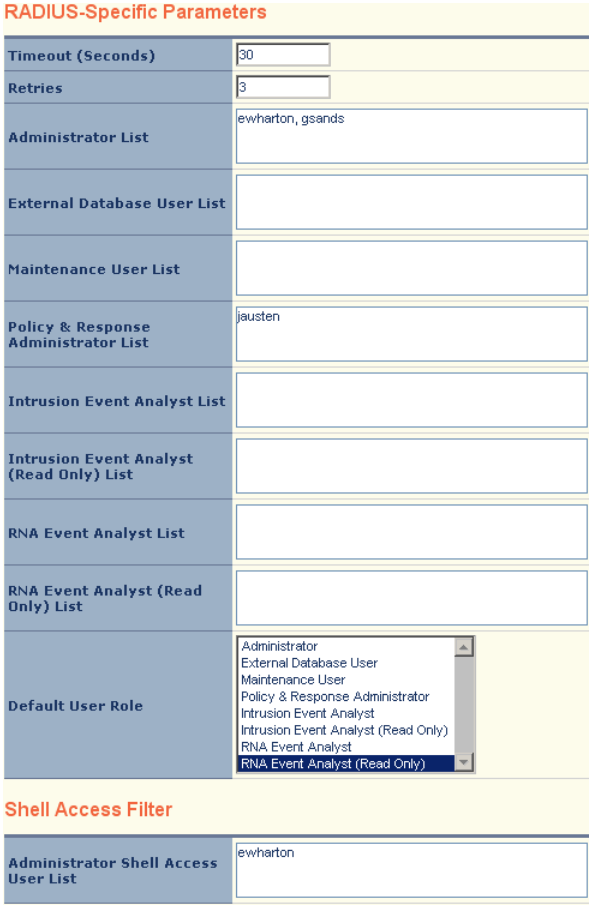
The following figure illustrates a sample RADIUS login authentication object for a server running FreeRADIUS with an IP address of 10.10.10.98. Note that the connection uses port 1812 for access and that connections to the server time out after 30 seconds of disuse and then retry three times before attempting to connect to a backup authentication server.

This example illustrates important aspects of RADIUS user role configuration:

- Users `ewharton` and `gsands` are granted administrative access to Sourcefire 3D System appliances where this authentication object is enabled.
- The user `jausten` is granted Intrusion Event Analyst access to Sourcefire 3D System appliances where this authentication object is enabled.
- The user `cbronte` is granted RNA Event 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.

Keysight, EX-1004
Page 1671

The following graphic depicts the role configuration for the example:



Authenticating a User with Custom Attributes

Platform: DC/MDC

Component: None

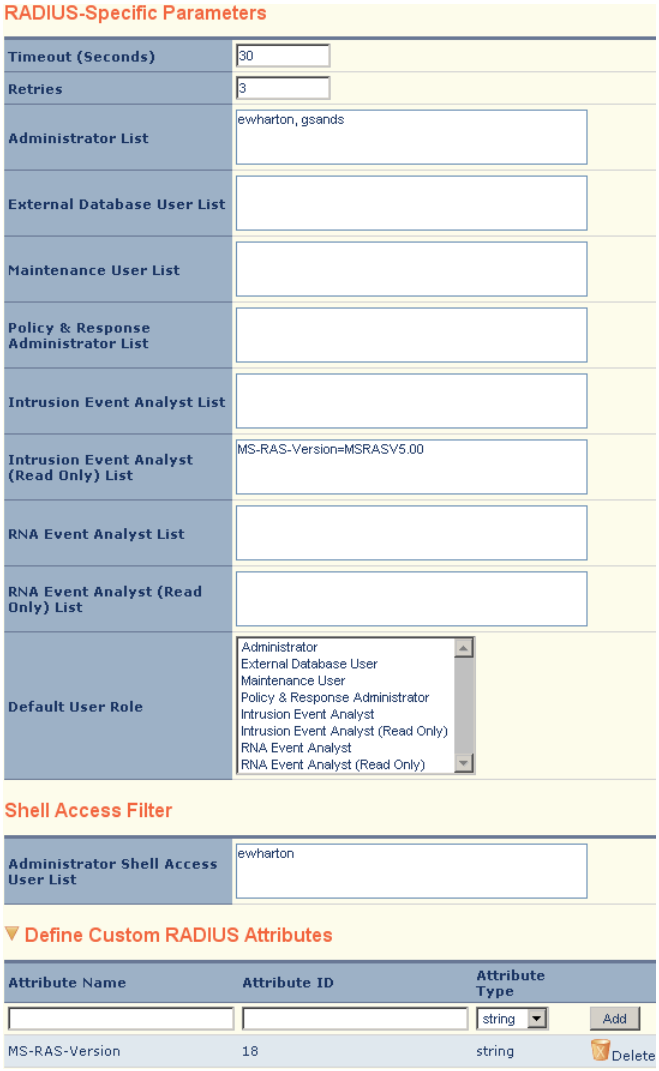
You can use an attribute-value pair to identify users who should receive a particular user role. If the attribute you use is a custom attribute, you must define the custom attribute.

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

Keysight, EX-1004
Page 1672

logging in to RADIUS through a Microsoft v. 5.00 remote access server should receive the Intrusion Event Analyst (Read Only role), so you type the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Intrusion Event Analyst (Read Only)** field.



Keysight, EX-1004
Page 1673

Editing RADIUS Authentication Objects

Platform: DC/MDC

Component: None

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 re-apply the policy.

To edit an authentication object:

Access: Admin

1. Select **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.

2. Click **Edit** next to the object you want to edit.

   The Create Authentication Object page appears.

3. Modify the object settings as needed.

   For more information, see the following topics:

   - Creating RADIUS Authentication Objects on page 1662
   - Configuring RADIUS Connection Settings on page 1663
   - Configuring RADIUS User Roles on page 1665
   - Configuring Administrative Shell Access on page 1668
   - Testing User Authentication on page 1670

4. Click **Save**.

   Your changes are saved and the Login Authentication page re-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 1741 and Applying a System Policy on page 1736.

Deleting Authentication Objects

Platform: DC/MDC

Component: None

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 **Operations > Configuration > Login Authentication**.

   The Login Authentication page appears.

Keysight, EX-1004
Page 1674

2. Click **Delete** next to the object you want to delete.

   The object is deleted and the Login Authentication page appears.

Managing User Accounts

If you have Admin access, you can use the web interface to view and manage user accounts on a Defense Center or a 3D Sensor, including adding, modifying, and deleting accounts. User accounts without Admin 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 1675 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 1676 describes the different options you can use when you add a new user account.
- Managing Externally Authenticated User Accounts on page 1680 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 1684 explains how to access and modify an existing user account.
- Modifying Restricted Event Analyst Access Properties on page 1685 explains how to restrict the data available to a user account with restricted data access.
- Deleting User Accounts on page 1690 explains how to delete user accounts.
- User Account Privileges on page 1690 contains tables that list the menus and options each type of user account can access.

Viewing User Accounts

Platform: DC/MDC/3D Sensor

Component: None

From the User Management page, you can view, edit, and delete existing accounts. You can determine 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 Action column allows you to 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)**.

Keysight, EX-1004
Page 1675

To access the User Management page:

Access: Admin

▶ Select **Operations > 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 1676
- Configuring User Roles on page 1682
- Modifying User Privileges and Options on page 1684
- Modifying Restricted Event Analyst Access Properties on page 1685
- Modifying User Passwords on page 1689
- Deleting User Accounts on page 1690

Adding New User Accounts

Platform: DC/MDC/3D Sensor

Component: None

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 an 8000 Series sensor, you can also configure the level of command line access the user will have.

To add a new user:

Access: Admin

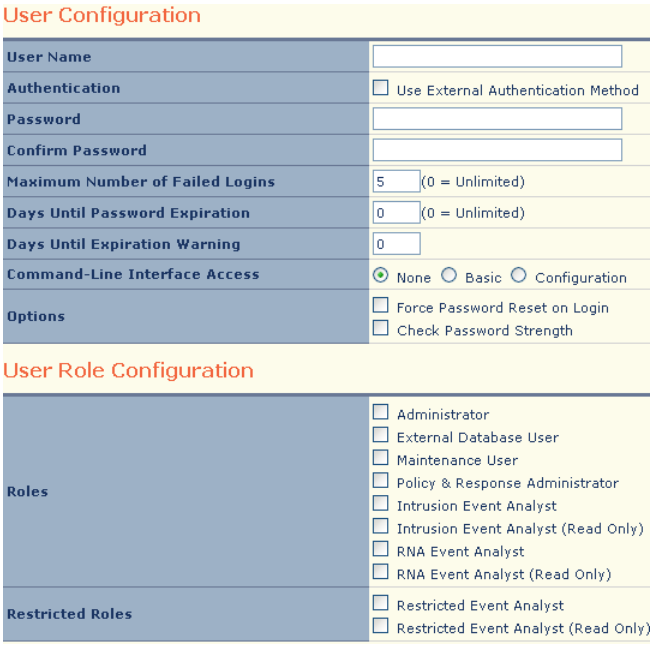
1. Select **Operations > User Management**.

   The User Management page appears.

Keysight, EX-1004
Page 1676

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.

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 1631 and Configuring Authentication Profiles on page 1741.

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. It cannot be a word that appears in a dictionary or include consecutive repeating characters.

Keysight, EX-1004
Page 1677

6. Configure the remaining password user account options.

   For more information, see the User Account Password Options table on page 1681.

7. If you are creating a local user through the web interface of an 8000 Series sensor, 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 8000 Series Sensors on page 1678.

8. Select access roles to grant to the user.

   IMPORTANT! For Series 2 sensors without an IPS license and all 8000 Series sensors, the Administrator access role is granted by default and cannot be changed. Therefore, the User Role Configuration options will not appear. Skip to step 10 to continue creating the user account.

   For more information, see Configuring User Roles on page 1682.

9. Optionally, for users with event analyst roles, select **Restrict Deletion Rights - User Cannot Delete Bookmarks, Searches, Reports, Report Profiles, Custom Workflows or Custom Tables Created by Other Users** to restrict the user to deletion of reports, report profiles, searches, bookmarks, custom tables, and custom workflows created by the user.

10. Click **Save**.

    The user is created and the User Management page appears again.

    TIP! Click **Deactivate** next to the name of an internally authenticated user on the User Management page to disable that user account without deleting it. To reactivate a user, click **Activate** next to the user name.

Managing Command Line Access on 8000 Series Sensors

Platform: 8000 Series Sensor

Component: None

On a Series 3 sensor, you can assign command line interface access to local sensor users.

Keysight, EX-1004
Page 1678

Note that you can also assign command line access for users on a Virtual 3D Sensor, but you use commands from the command line interface. For more information, see user on page 2066.

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	show hyperthreading
end	show interfaces
exit	show interface-sets
help	show link-state
history	show managers
logout	show memory
? (question mark)	show model
?? (double question mark)	show network
show alarms	show network-modules
show clustering	show ntp
show cpu	show portstats
show detection-engines	show processes
show disk	show reserve-cpu
show dns	show serial-number
show failopen	show summary
show gui	show time
show hostname	show version

Keysight, EX-1004
Page 1679

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 1635, Configuring Administrative Shell Access on page 1654, and Configuring Administrative Shell Access on page 1668.

Managing Externally Authenticated User Accounts

Platform: DC/MDC/3D Sensor

Component: None

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 1741 and Understanding User Privileges on page 1629.

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 sensors if you want to use external authentication on them.

For more information on modifying user access, see Modifying User Privileges and Options on page 1684. 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**.

Keysight, EX-1004
Page 1680

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**.

WARNING! On Series 2 appliances and 8000 Series Defense Centers, all shell users have `sudoers` privileges. Make sure that you restrict the list of users with shell access appropriately. On 8000 Series sensors and Virtual 3D Sensors, 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 1635, Configuring Administrative Shell Access on page 1654, and Configuring Administrative Shell Access on page 1668.

Managing User Password Settings

You can also control how and when the password for each user account is changed, as well as when user accounts are disabled. The User Account Password Options table describes some of the options you can use to regulate passwords and account access.

IMPORTANT! After you enable **Use External Authentication Method**, password options no longer appear. Use the external authentication server to manage password settings.

User Account Password Options

Option	Description
Use External Authentication Method	Select this option if you want this user's credentials to be externally authenticated. **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.
Maximum Number of Failed Logins	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.
Days Until Password Expiration	Enter the number of days after which the user's password expires. The default setting is 0, which indicates that the password never expires.

Keysight, EX-1004
Page 1681

User Account Password Options (Continued)

Option	Description
Days Until Expiration Warning	Enter the number of warning days users have to change their password before their password actually expires. The default setting is 0 days. **WARNING!** The number of warning days must be **less than** the number of days before the password expires
Force Password Reset on Login	Select this option to force the user to change his password the first time the user logs in.
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. It cannot be a word that appears in a dictionary or include consecutive repeating characters.

Configuring User Roles

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 grant Intrusion Event and RNA Event Analyst access to analysts while reserving the Administrator role for the network administrator managing the Sourcefire 3D System.

The 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 1690.

The menus and other options in the web interface that users can access depend on their roles, which are briefly described in the following table. For a list of the

Keysight, EX-1004
Page 1682

menus and options available to each role, see User Account Privileges on page 1690.

User Roles

User Role	Privileges
Administrator	Provides access to analysis and reporting features, rule and policy configuration, system management, and all maintenance features. Administrator users see the main toolbar as well as all the menu options. **IMPORTANT!** For Series 2 sensors without an IPS license and all 8000 Series sensors, the Administrator access role is granted by default and cannot be changed. Note that you should limit use of the Administrator role for security reasons.
External Database	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 1799. On the web interface, External Database users see the main toolbar and online help-related options on the **Operations** menu.
Maintenance	Provides access to monitoring and maintenance features. Maintenance users see the main toolbar and maintenance-related options on the **Operations** menu.
RNA Event Analyst	Provides access to RNA analysis features, including event views, network maps, host profiles, services, vulnerabilities, client applications, and reports. RNA Event Analysts see the main toolbar and RNA analysis-related options on the **Analysis & Reporting** and **Operations** menus.
RNA Event Analyst (Read Only)	Provides read-only access to analysis features, including event views, network maps, host profiles, services, vulnerabilities, client applications, incidents, and reports. RNA Event Analysts see the main toolbar and analysis-related options on the **Analysis & Reporting** and **Operations** menus.
Intrusion Event Analyst	Provides access to IPS analysis features, including intrusion event views, incidents, and reports. Intrusion Event Analysts see the main toolbar and IPS analysis-related options on the **Analysis & Reporting** and **Operations** menus.
Intrusion Event Analyst (Read Only)	Provides read-only access to IPS analysis features, including intrusion event views, incidents, and reports. Intrusion Event Analysts see the main toolbar and IPS analysis-related options on the **Analysis & Reporting** and **Operations** menus.

Keysight, EX-1004
Page 1683

User Roles (Continued)

User Role	Privileges
Restricted Event Analyst	Provides access to the same features as Intrusion Event Analyst or RNA Event Analyst access. You can restrict access by allowing access to only for those events that match specified search criteria or you can turn off access for an entire category of events, as described in Modifying Restricted Event Analyst Access Properties on page 1685. Restricted event analyst users see only the main toolbar and analysis-related options on the **Analysis & Reporting** and **Operations** menus.
Policy & Response Administrator	Provides access to rules and policy configuration. Policy & Response Administrators have access to the main toolbar and rule and policy-related options on the **Policy & Response** and **Operations** 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 1676.

Some user roles are mutually exclusive. For example, an External Database user can also have Administrator access, but cannot have any other user roles. However, a single user can have both Intrusion Event Analyst and RNA Event Analyst access.

Note that externally authenticated users 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.

Modifying User Privileges and Options

Platform: DC/MDC/3D Sensor

Component: None

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.

Keysight, EX-1004
Page 1684

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 **Operations > User Management**.

   The User Management page appears.

2. Click **Edit** 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 1680 for a description of how users can be authenticated through external servers.
   - See Managing User Password Settings on page 1681 for information on changing password settings for internally authenticated users.
   - See Configuring User Roles on page 1682 for more information on configuring roles to grant access for Sourcefire 3D System functions.
   - Optionally, for users with event analyst roles, select or clear the **Only delete items created by user** option to manage the user's ability to delete of items not created by that user.

Modifying Restricted Event Analyst Access Properties

Platform: DC/MDC/3D Sensor

Component: None

User accounts with Restricted Event Analyst access use saved searches to specify which events a user can view. You can specify this information only after the user is added. See Adding New User Accounts on page 1676 for information about adding new user accounts.

Keysight, EX-1004
Page 1685

Restricted event analyst users have access to only a few sections of the web interface. The Restricted Event Analyst Settings table shows the correlation between platform and access requirements for the restricted event analyst.

Restricted Event Analyst Settings

To allow the restricted event analyst to...	When these platforms are present...	Set this data set or data sets to Show All or to a specific search
view the network map	DC + RNA	One or more of the following: • Host Attributes Data • RNA Client Applications Data • RNA Hosts Data • RNA Services Data • Vulnerabilities Data • Third-Party Vulnerabilities Data
view network discovery events	DC + RNA	RNA Events Data
view hosts	DC + RNA	RNA Hosts Data
view host attributes	DC + RNA	Host Attributes Data
view services	DC + RNA	RNA Services Data
view vulnerabilities	DC + RNA	Vulnerabilities Data
view third-party vulnerabilities	DC + RNA	Third-Party Vulnerabilities Data
view client applications	DC + RNA	RNA Client Applications Data
view flow data	DC + RNA	Flow Data
view compliance events	DC + RNA	Compliance Events Data
view white list events	DC + RNA	White List Events Data
view white list violations	DC + RNA	White List Violations Data

Keysight, EX-1004
Page 1686

Restricted Event Analyst Settings (Continued)

To allow the restricted event analyst to...	When these platforms are present...	Set this data set or data sets to Show All or to a specific search
view users or user events	DC + RUA	Users Data
view intrusion events	IPS	Intrusion Events Data
use the clipboard	IPS	N/A - included in the base set of rights for the restricted analyst role
generate (but not view) reports	IPS	All data sets for which the user can generate reports
create (but not modify) incident reports	IPS	All data sets for which the user can create incident reports
change user-specific preferences such as the account password, time zone, and event view settings	DC/MDC or 3D Sensor	N/A - included in the base set of rights for the restricted analyst role
create custom workflows and, on the Defense Center, custom tables	DC/MDC or 3D Sensor	All data sets for which the user can create custom workflows
create and manage bookmarks	DC/MDC or 3D Sensor	All data sets for which the user can create or access bookmarks
view events from a custom table	Platforms required to view custom table	All data sets for the applicable custom tables

If you want to ensure that a user only sees data for a specific subnet, create multiple private saved searches, one for each of the event types, and then apply each saved search to the account as described in the following procedure.

Note that you must have saved private searches available before you can add restricted event analyst values to a user account. Searches **must** be private. If they are saved as public, restricted event analyst users could delete the searches and enhance their access privileges. See Searching for Events on page 1550 for more information.

Keysight, EX-1004
Page 1687

To restrict event analyst access to events:

Access: Admin

1. Select **Operations > User Management**.

   The User Management page appears.

2. Click **Edit** next to the user to whom you want to grant restricted event analyst rights.

3. If the user you want to modify does not already have the Restricted Event Analyst option enabled, select **Restricted Event Analyst**.

IMPORTANT! You cannot select **Restricted Event Analyst** if **Administrator**, **Intrusion Event Analyst**, **Intrusion Event Analyst (Read Only)**, **RNA Event Analyst**, or **RNA Event Analyst (Read Only)** access is enabled.

The Restrictions section of the page appears. The Defense Center version of the page is shown below.



IMPORTANT! If you created any custom tables on the Defense Center, they appear on this page.

---

4. For each row, you have three choices:
   - To grant access to all events for a category, select **Show All Data**.
   - To grant access to events that match a specific saved search, select the search that you want to use to restrict the user account.
   - To deny access to all events in a category, select **Hide Data**.

Keysight, EX-1004
Page 1688

5. Click **Save** to save your changes and return to the User Management page.

Modifying User Passwords

Platform: DC/MDC/3D Sensor

Component: None

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.

TIP! If you want to force a user to change the password on the next log-in, click **Reset Password** next to the user account on the User Management page.

To change a user's password:

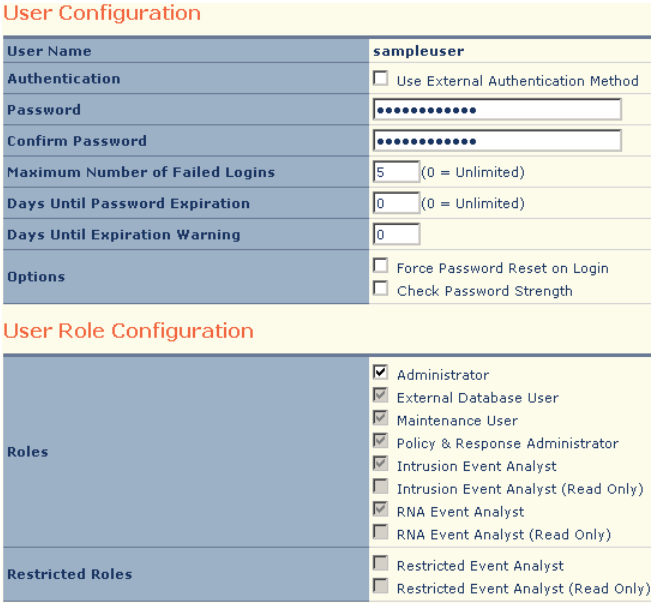
Access: Admin

1. Select **Operations > User Management**.

   The User Management page appears.

2. Next to the user name, click **Edit**.

   The Edit User page appears.



3. In the **Password** field, type the new password (up to 32 alphanumeric characters).

Keysight, EX-1004
Page 1689

4. In the **Confirm Password** field, re-type 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. 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 Password Settings on page 1681.
   - For more information on user roles, see Configuring User Roles on page 1682.

6. Click **Save**.

   The password is changed and any other changes saved.

Deleting User Accounts

Platform: DC/MDC/3D Sensor

Component: None

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 **Operations > User Management**.

   The User Management page appears.

2. Next to the user whose account you want delete, click **Delete**.

   The account is deleted.

User Account Privileges

Platform: DC/MDC/3D Sensor

Component: None

The following sections provide a list of the menus and toolbar options in Sourcefire 3D System and the user account privileges required to access them.

For more information on the access notations used in the tables that follow and throughout this documentation, see Access Conventions on page 57.

- Analysis & Reporting Menu on page 1691
- Policy & Response Menu on page 1694
- Operations Menu on page 1695
- Toolbar Options on page 1697

Keysight, EX-1004
Page 1690

Analysis & Reporting Menu

Platform: DC/MDC/3D Sensor

Component: None

The Analysis & Reporting Menu table lists the user account privileges required to access each option on the Analysis & Reporting menu. An X indicates that the user can access the option. Users with External Database cannot see the Analysis & Reporting menu at all.

Analysis & Reporting Menu

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Dashboards	X		X	X	X		X
Event Summary	X			X	X	X	
Intrusion Event Statistics	X				X		
Event Graphs	X				X		
RNA Statistics	X			X		X	
Flow Summary	X			X		X	
IPS	X			X	X	X	
Events	X			X	X	X	
Reviewed Events	X			X	X	X	
Clipboard	X			X	X	X	
Incidents	X				X		
RNA	X			X		X	
Network Map | Hosts	X			X		X	
Network Map | Network Devices	X			X		X	
Network Map | Services	X			X		X	
Network Map | Vulnerabilities	X			X		X	

Keysight, EX-1004
Page 1691

Analysis & Reporting Menu (Continued)

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Network Map | Host Attributes	X			X		X	
RNA Events	X			X		X	
Hosts	X			X		X	
Host Attributes	X			X		X	
Services	X			X		X	
Client Applications	X			X		X	
Flow Data	X			X		X	
Vulnerabilities	X			X		X	
Third-Party Vulnerabilities	X			X		X	
RUA	X			X	X	X	
Users	X			X	X	X	
RUA Events	X			X	X	X	
Compliance	X			X	X	X	
Compliance Events	X			X	X	X	
White List Events	X			X	X	X	
White List Violations	X			X	X	X	
Custom Tables	X			X		X	
Searches	X			X	X	X	
Audit Log	X						
Client Applications	X			X		X	
Compliance Events	X			X	X	X	

Keysight, EX-1004
Page 1692

Analysis & Reporting Menu (Continued)

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Flow Data	X			X		X	
Health Events	X		X		X		
Host Attributes	X			X		X	
Hosts	X			X		X	
Intrusion Events	X			X	X	X	
Remediation Status	X						
RNA Events	X			X	X	X	
RUA Events	X			X	X	X	
Scan Results	X						
Services	X			X		X	
SEU Import Log	X						
Third-Party Vulnerabilities	X			X		X	
Users	X			X	X	X	
Vulnerabilities	X			X		X	
White List Events	X			X	X	X	
White List Violations	X			X	X	X	
Custom Workflows	X			X	X	X	
Bookmarks	X			X	X	X	
Report Profiles	X			X	X	X	

Keysight, EX-1004
Page 1693

Policy & Response Menu

Platform: DC/MDC/3D Sensor

Component: None

The Policy & Response Menu table lists the user account privileges required to access each option on the Policy & Response menu. An X indicates that the user can access the option. Users with External Database, Intrusion Event Analyst, RNA Event Analyst, or Maintenance access cannot see the Policy & Response menu at all.

Policy & Response Menu

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
IPS	X						X
Intrusion Policy	X						X
SEU	X						X
Rule Editor	X						X
Email	X						X
OPSEC	X						X
RNA	X						X
Detection Policy	X						X
Host Attributes	X						X
RNA Detectors	X						X
Custom Fingerprinting	X						X
Custom Product Mappings	X						
User Third-Party Mappings	X						
Network Map | Custom Topology	X						
Compliance	X						X

Keysight, EX-1004
Page 1694

Policy & Response Menu (Continued)

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Policy Management	X						X
Rule Management	X						X
White List	X						X
Traffic Profiles	X						X
Responses	X						X
Alerts	X						X
Impact Flag Alerts	X						X
RNA Event Alerts	X						X
Remediations	X						X
Groups	X						X

Operations Menu

Platform: DC/MDC/3D Sensor

Component: None

The Operations Menu table lists the user account privileges required to access each option on the Operations menu. An X indicates that the user can access the option. All users can access at least some options on the Operations menu.

Operations Menu

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Configuration	X						
RNA/RUA Event Purge	X					X	
Detection Engines	X						

Keysight, EX-1004
Page 1695

Operations Menu (Continued)

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
High Availability	X						
eStreamer	X						
Login Authentication	X						
RUA	X						
Sensors	X						
User Management	X						
System Settings	X						
System Policy	X						
Update	X						
Monitoring	X		X				X
Statistics	X		X				
Performance | IPS	X		X				
Performance | RNA	X		X				
Audit	X						
Task Status	X		X				X
Syslog	X		X				
Health	X		X		X		
Tools	X	X	X	X	X	X	X
Scheduling	X		X				
Backup/Restore	X						
Import/Export	X						

Keysight, EX-1004
Page 1696

Operations Menu (Continued)

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Whois	X	X	X	X	X	X	X
Scan Results	X		X				X
Scanners	X						X
Help	X		X	X	X	X	X
About	X	X	X	X	X	X	X
Online	X	X	X	X	X	X	X
Email Support	X	X	X	X	X	X	X
Support Site	X	X	X	X	X	X	X

Toolbar Options

Platform: DC/MDC/3D Sensor

Component: None

The Toolbar Options table lists the user account privileges required to access each option on the toolbar and its sub-menus. An X indicates that the user can

Keysight, EX-1004
Page 1697

access the option. All users can access at least some of the options on the toolbar.

Toolbar Options

Menu	Admin	External DB	Maint	RNA/ RNA-RO Event Analyst	IPS/ IPS-RO Event Analyst	Restr. Event Analyst	P&R Admin
Health	X		X	X	X		
Preferences	X	X	X	X	X	X	X
Preferences | Home Page	X	X	X	X	X	X	X
Preferences | Event View Settings	X			X	X	X	
Preferences | Change Password	X	X	X	X	X	X	X
Preferences | Time Zone Settings	X	X	X	X	X	X	X
Help	X	X	X	X	X	X	X
Logout	X	X	X	X	X	X	X

Keysight, EX-1004
Page 1698

Chapter 39
Scheduling Tasks

You can schedule many different types of administrative tasks to run at scheduled times, including:

- running backups
- applying intrusion policies
- generating reports
- running Nmap scans
- using RNA rule recommendations
- importing Security Enhancement Updates (SEUs)
- pushing (copying) downloaded updates to managed sensors
- downloading and installing software updates
- downloading and installing vulnerability database (VDB) updates

You can schedule tasks to run once or on a recurring schedule.

IMPORTANT! Some tasks (such as those involving automated software and SEU updates and those that require pushing updates or intrusion policies to managed sensors) can place a significant load on networks with low bandwidths. You should always schedule tasks like these to run during periods of low network use.

Keysight, EX-1004
Page 1699

See the following sections for more information:

- Configuring a Recurring Task on page 1700 explains how to set up a scheduled task so that it runs at regular intervals.
- Automating Backup Jobs on page 1702 provides procedures for scheduling backup jobs.
- Automating Software Updates on page 1704 provides procedures for scheduling the download, push, and installation of software updates.
- Automating Vulnerability Database Updates on page 1710 provides procedures for scheduling the download, push, and installation of software updates.
- Automating SEU Imports on page 1716 provides procedures for scheduling rule updates.
- Automating Intrusion Policy Applications on page 1719 provides procedures for scheduling intrusion policy applications.
- Automating Reports on page 1721 provides procedures for scheduling reports.
- Automating Nmap Scans on page 1722 provides procedures for scheduling Nmap scans.
- Automating Recommended Rule State Generation on page 1725 provides procedures for scheduling automatic update of intrusion rule state recommendations based on RNA data.
- Viewing Tasks on page 1727 describes how to view and manage tasks after they are scheduled.
- Editing Scheduled Tasks on page 1730 describes how to edit an existing task.
- Deleting Scheduled Tasks on page 1730 describes how to delete one-time tasks and all instances of recurring tasks.

Configuring a Recurring Task

Platform: DC/MDC/Standalone Sensor

Component: None

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.

Keysight, EX-1004
Page 1700

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 system settings. Further, the Defense Center or 3D Sensor with IPS 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 2am during standard time, it will run at 3am during DST. Similarly, if you create a task scheduled for 2am during DST, it will run at 1am during standard time.

To configure a recurring task:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

   The Scheduling page appears.

2. Click **Add Task**.

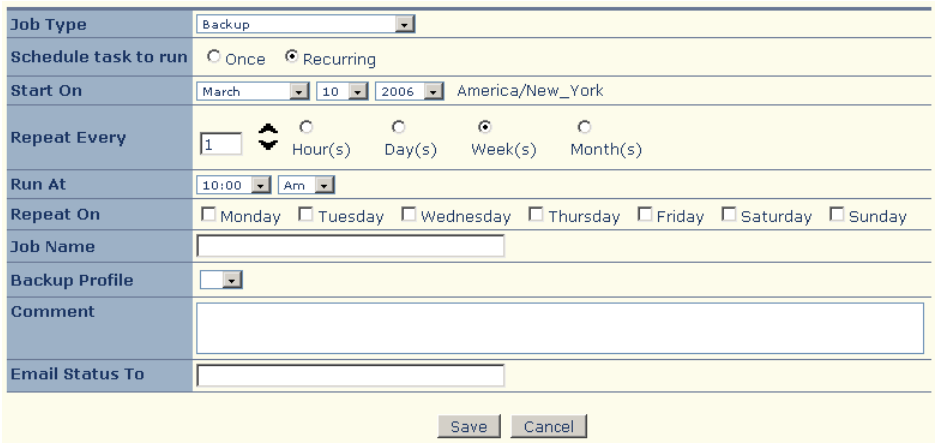
   The Add Task page appears.

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.



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 use the arrow buttons to specify the interval. For example, type 2 and select **Day(s)** to run the task every two days.

Keysight, EX-1004
Page 1701

7. In the **Run At** field, specify the time when you want to start your recurring task.
8. If you selected Week(s) in the **Repeat Every** field, 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 Month(s) in the **Repeat Every** field, 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 Add Task page are determined by the task you are creating. See the following sections for more information:

   - Automating Backup Jobs on page 1702
   - Automating Software Updates on page 1704
   - Automating Vulnerability Database Updates on page 1710
   - Automating SEU Imports on page 1716
   - Automating Intrusion Policy Applications on page 1719
   - Automating Reports on page 1721
   - Automating Nmap Scans on page 1722
   - Automating Recommended Rule State Generation on page 1725

Automating Backup Jobs

Platform: DC/MDC/Standalone Sensor

Component: None

You can use the scheduler to automate system backups of a Defense Center or a 3D Sensor with IPS.

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 1990.

To automate backup tasks:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

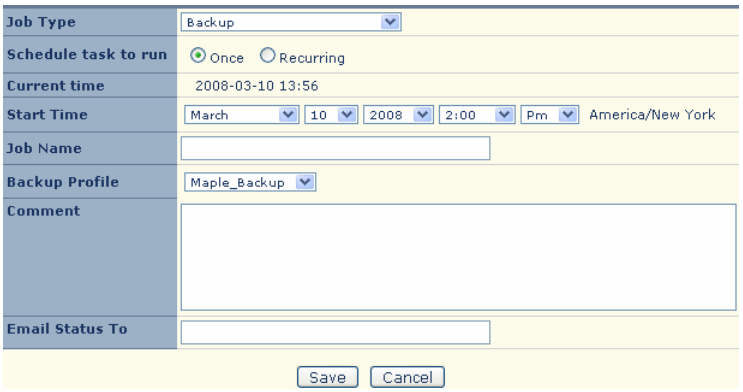
   The Scheduling page appears.
2. Click **Add Task**.

   The Add Task page appears.

Keysight, EX-1004
Page 1702

3. From the **Job Type** list, select **Backup**.

   The page reloads to show the backup options.



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 1700 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 1990.
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 1751 for more information about configuring a relay host.
9. Click **Save**.

   The backup task is created.

Keysight, EX-1004
Page 1703

Automating Software Updates

The tasks you must schedule to install software updates vary depending on whether you are updating an appliance directly or are using a Defense Center to update managed sensors. Sourcefire **strongly** recommends that you use your Defense Centers to update the sensors they manage.

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 appliances that cannot access the Sourcefire Support site. For information on manually updating the Sourcefire 3D System, see Updating System Software on page 1831.

If you want to update an appliance directly, you can schedule the software installation using the Install Latest Update task. To use a Defense Center to automate software updates for its managed 3D Sensors, you must schedule two tasks:

1. Push (copy) the update to managed sensors using the Push Latest Update task.
2. Install the update on managed sensors 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 sensors, you must first push the update to the sensor before you can install it. (Note that during the manual update process you do not have to push an update to managed sensors before you install it. For more information, see Updating Managed Sensors on page 1838.)

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 sensor, 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 1705
- Automating Software Pushes on page 1706
- Automating Software Installs on page 1708

Keysight, EX-1004
Page 1704

Automating Software Downloads

Platform: DC/MDC/Standalone Sensor

Component: None

You can create a scheduled task that automatically downloads the latest software updates from Sourcefire. On the Defense Center, you can also automate VDB updates. You can use this task to schedule download of updates you plan to install manually.

To automate software updates:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

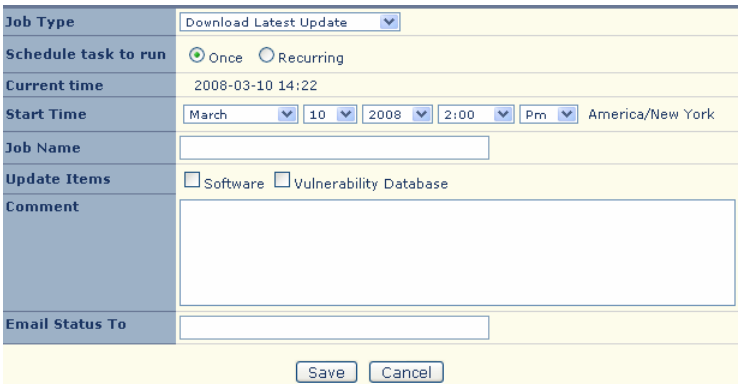
   The Scheduling page appears.

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Download Latest Update**.

   The Add Task page reloads to show the update options. The Defense Center version of the page is shown below.



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 1700 for details.

Keysight, EX-1004
Page 1705

5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

   If your appliance is not directly connected to the Internet, you should set up a proxy as described in Configuring Network Settings on page 1800 to allow it to download updates from the Sourcefire Support site (https://support.sourcefire.com/).

6. In the **Update Items** section, specify which updates you want to download.
   - Select **Software** to download the most recent software update.
   - On a Defense Center only, select **Vulnerability Database** to download the most recent VDB update.

   Both options are selected by default.

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 1751 for more information about configuring a relay host.

9. Click **Save**.

   The task is created.

Automating Software Pushes

Platform: DC

Component: None

If you want to automate the installation of software updates on managed 3D Sensors, you must push the updates to the managed sensors before installing.

IMPORTANT! You cannot schedule pushes from appliances that cannot access the Sourcefire Support site. For information on manually updating the Sourcefire 3D System, see Updating System Software on page 1831.

When you push updates to managed sensors, information about the push process status is reported on the Tasks page. See Viewing the Status of Long-Running Tasks on page 2016 for more information.

Keysight, EX-1004
Page 1706

When you create the task to push software updates to managed sensors, make sure you allow enough time between the push task and a scheduled install task for the updates to be copied to the sensor.

To push software updates to managed sensors:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

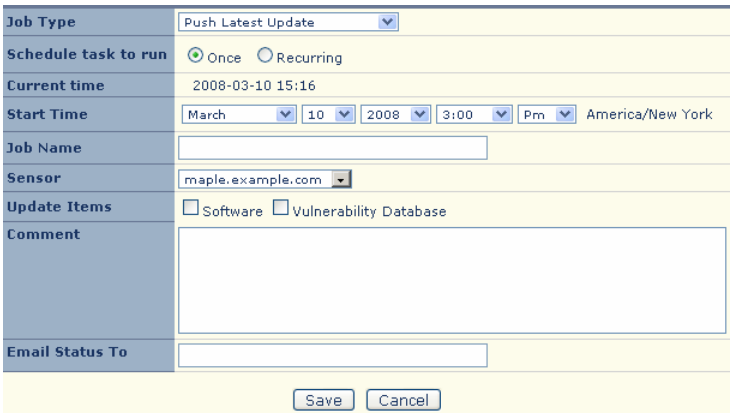
   The Scheduling page appears.

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Push Latest Update**.

   The page reloads to show the options for pushing updates.



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 1700 for details.

5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

6. From the **Sensor** list, select the sensor that you want to receive updates.

7. In the **Update Items** section, specify which updates you want to push to your managed sensors.
   - Select **Software** to push the software update.
   - On a Defense Center only, select **Vulnerability Database** to push the VDB update.

   Both options are selected by default.

Keysight, EX-1004
Page 1707

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 1751 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 2016 for more information.

Automating Software Installs

Platform: DC/MDC/Standalone Sensor

Component: None

If you are using a Defense Center to create a task to install a software update on a managed sensor, make sure you allow enough time between the task that pushes the update to the sensor and the task that installs the update. See Automating Software Pushes on page 1706 for information about pushing updates to managed sensors.

You **cannot** schedule updates for appliances that cannot access the Sourcefire Support site. For information on manually updating the Sourcefire 3D System, see Updating System Software on page 1831.

WARNING! Depending on the update being installed, the appliance may reboot after the software is installed.

To schedule a software installation task:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

   The Scheduling page appears.

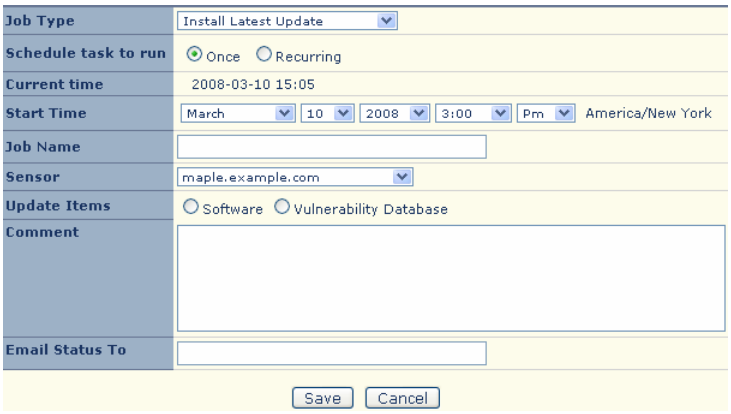
2. Click **Add Task**.

   The Add Task page appears.

Keysight, EX-1004
Page 1708

3. From the **Job Type** list, select **Install Latest Update**.

   The page reloads to show the options for installing updates. The Defense Center version of the page is shown below.



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 1700 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. If you are using a Defense Center, from the **Sensor** list, you have the following options:
   - Select the sensor 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** to install the software update.
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 1751 for more information about configuring a relay host.

Keysight, EX-1004
Page 1709

10. Click **Save**.

    The scheduled software installation 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 2016 for more information.

Automating Vulnerability Database Updates

Sourcefire uses vulnerability database (VDB) updates to distribute new operating system, client application, service, and payload fingerprints as we expand the list of network assets that RNA recognizes. VDB updates also include new vulnerabilities discovered by the Sourcefire Vulnerability Research Team (VRT). You can use the scheduling feature to download and install the latest VDB updates, thereby ensuring that RNA is using the most up-to-date information to evaluate the hosts on your network.

If your Sourcefire 3D System deployment includes IPS and RNA monitoring the same network segments, you should download and install VDB updates and SEUs on a regular basis. This ensures that your Defense Center is correctly setting the impact flag on the intrusion events generated by the traffic on your network.

IMPORTANT! You cannot schedule updates for appliances that cannot access the Sourcefire Support site. For information on manually updating the Sourcefire 3D System, see Updating System Software on page 1831.

When automating VDB updates for your Defense Center, you must automate two separate steps:

1. Downloading the VDB update.
2. Installing the VDB update.

When automating VDB updates for managed sensors with RNA, you must schedule three tasks in this order:

1. Download the VDB update on your Defense Center.
2. Push the VDB update to your managed 3D Sensors that are using the RNA component.
3. Install the VDB update on the Defense Center and on those managed sensors.

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 it runs the next day.

Keysight, EX-1004
Page 1710

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.

See the following sections for more information:

- Automating VDB Update Downloads on page 1711
- Automating VDB Update Pushes on page 1712
- Automating VDB Update Installs on page 1714

Automating VDB Update Downloads

Platform: DC/MDC

Component: RNA

You can create a scheduled task that automatically downloads the latest VDB updates from Sourcefire. Because you should install the same version of the VDB on all the appliances in your deployment, use your Defense Centers to push and install the VDB on all managed 3D Sensors with RNA.

To automate VDB updates:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

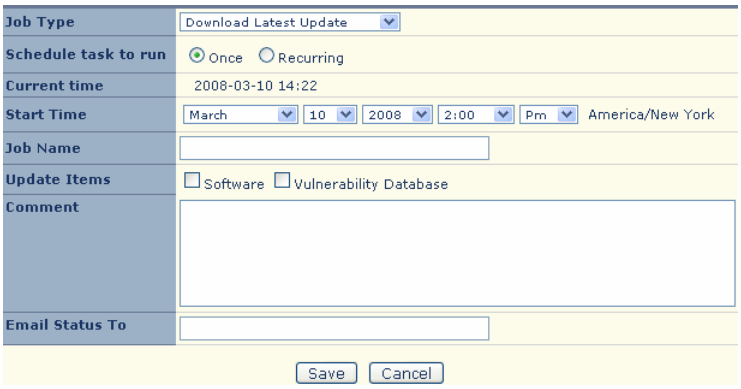
   The Scheduling page appears.

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Download Latest Update**.

   The Add Task page reloads to show the update options.



Keysight, EX-1004
Page 1711

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 1700 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

   If your appliance is not directly connected to the Internet, you should set up a proxy as described in Configuring Network Settings on page 1800 to allow it to download updates from the Sourcefire Support site (https://support.sourcefire.com/).
6. In the **Update Items** section, make sure **Vulnerability Database** is selected.

   Both the **Software** and **Vulnerability Database** options are selected by default.
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 1751 for more information about configuring a relay host.
9. Click **Save**.

   The task is created.

Automating VDB Update Pushes

Platform: DC/MDC

Component: RNA

If you are installing VDB updates on managed 3D Sensors with RNA, you must download the VDB update, then push the update to the managed sensors before installing. (Note that during the manual update process you do not have to push a

Keysight, EX-1004
Page 1712

VDB update to managed sensors before you install it. For more information, see Updating the Vulnerability Database on page 1846.)

IMPORTANT! You cannot schedule pushes from appliances that cannot access the Sourcefire Support site. For information on manually updating the VDB, see Updating the Vulnerability Database on page 1846.

When you push VDB updates to managed sensors, information about the process status is reported on the Tasks page. See Viewing the Status of Long-Running Tasks on page 2016 for more information.

When you create the task to push VDB updates to managed sensors, make sure you allow enough time between the push task and a scheduled install task for the updates to be copied to the sensor.

To push VDB updates to managed 3D Sensors with RNA:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

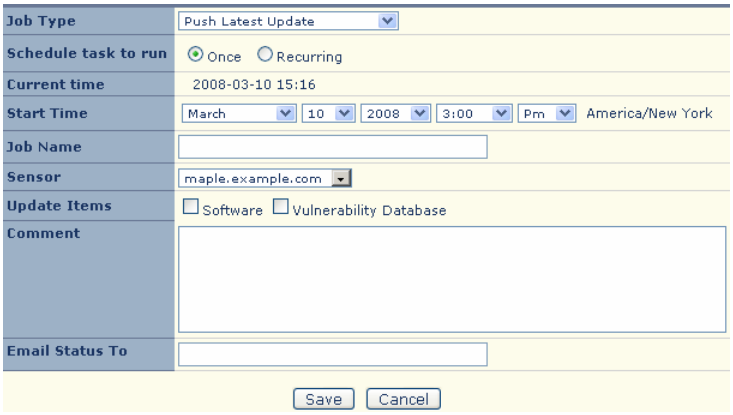
   The Scheduling page appears.

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Push Latest Update**.

   The page reloads to show the options for pushing updates.



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 1700 for details.

Keysight, EX-1004
Page 1713

5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. From the **Sensor** list, select the sensor that you want to receive updates.
7. In the **Update Items** section, make sure **Vulnerability Database** is selected.

   Both the **Software** and **Vulnerability Database** options are selected by default.
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 1751 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 2016 for more information.

Automating VDB Update Installs

Platform: DC/MDC

Component: RNA

You should allow enough time between the tasks that download and push the VDB update to the sensor and the task that installs the update; see Automating VDB Update Downloads on page 1711 and Automating VDB Update Pushes on page 1712 for information.

IMPORTANT! You cannot schedule updates for appliances that cannot access the Sourcefire Support site. For information on manually updating the VDB, see Updating the Vulnerability Database on page 1846.

To schedule a VDB update:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

   The Scheduling page appears.

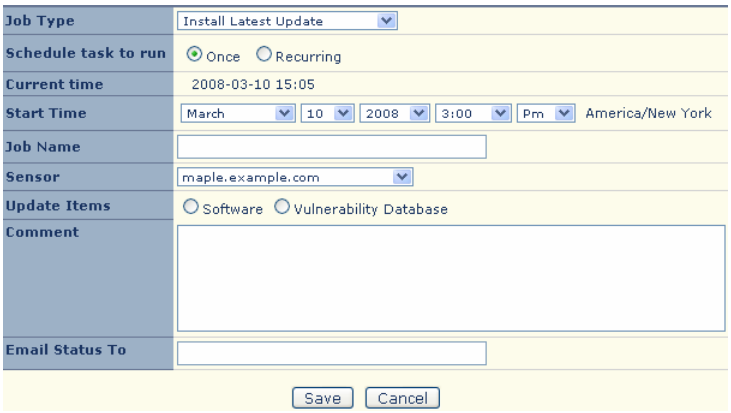
Keysight, EX-1004
Page 1714

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Install Latest Update**.

   The page reloads to show the options for installing updates.



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 1700 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. From the **Sensor** list, you have the following options:
   - If you want to install the update on a managed sensor, select the name of the sensor from the drop-down list.
   - If you want to install the update on the Defense Center, select the name of the Defense Center from the drop-down list.
7. In the **Update Items** section, select **Vulnerability Database** to install the VDB update.
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.

Keysight, EX-1004
Page 1715

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 1751 for more information about configuring a relay host.

10. Click **Save**.

    The scheduled VDB installation 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 2016 for more information.

Automating SEU Imports

Platform: DC/MDC/Standalone Sensor

Component: IPS

As new vulnerabilities are identified, the Sourcefire Vulnerability Research Team (VRT) releases Security Enhancement Updates (SEUs). An SEU contains new and updated standard text rules and shared object rules and may contain updated versions of Snort® and features such as preprocessors and decoders. You can automatically download and install SEUs.

The Import SEU task allows you to schedule the following subtasks separately or to combine them into one scheduled task:

1. Download the latest SEU.
2. Import the SEU.
3. Re-apply your intrusion policy so that the new SEU takes effect.

The selected subtasks present in the Import SEU task occur in the following order: download, install, rule state update, and policy re-apply. Once one subtask completes, the next configured subtask begins. Note that you can only re-apply policies applied from the appliance where the scheduled task is configured.

In addition to configuring SEU imports on the Scheduling page, you can also use the recurring SEU import feature on the Import SEU page. For more information on the recurring SEU import feature and a comparison of the two methods of setting up recurring imports, see Importing SEUs and Rule Files on page 1848.

IMPORTANT! SEUs may contain new binaries. Make sure your process for downloading and importing SEUs complies with your security policies. In addition, SEUs can be quite large, so make sure you schedule downloads during periods of low network use.

Keysight, EX-1004
Page 1716

Note the following when automating SEU imports:

- On the Defense Center, you also must re-apply your intrusion policies on your managed 3D Sensors with IPS. Applying an intrusion policy from a Defense Center to a managed sensor after you import an SEU does not apply the SEU to the sensor. However, any new rules or features provided by the SEU that are enabled in the policy you apply to the sensor are also enabled on the sensor by that policy.
- When you have previously applied a non-filtered VLAN or subnetwork policy as well as one or more filtered policies, and you do not automatically re-apply your intrusion policies after importing an SEU, you can use **Reapply All** or **Reapply Intrusion Policies** on the Detection Engines page (**Operations > Configuration > Detection Engines**) to manually and simultaneously re-apply your non-filtered policy and all filtered policies. See Defining VLANs and Subnetworks on page 393 for more information.
- If your Sourcefire 3D System deployment includes two Defense Centers configured as a high availability pair, you only need to import the SEU on one of the Defense Centers. The second Defense Center receives the SEU as part of the regular synchronization process.
- If you enable **Update when a new SEU is installed** for the base policy of an existing policy and the SEU contains changes to the default rule states for existing rules in that base policy, those changes are also imported. Note, however, that if you changed a rule state, the SEU does not override your change.

  VRT sometimes uses an SEU to change the default state of one or more rules in a default policy. If you allow SEUs to update your base policy, you also allow the SEU 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 SEU will not override your change.

To schedule an Import SEU task:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

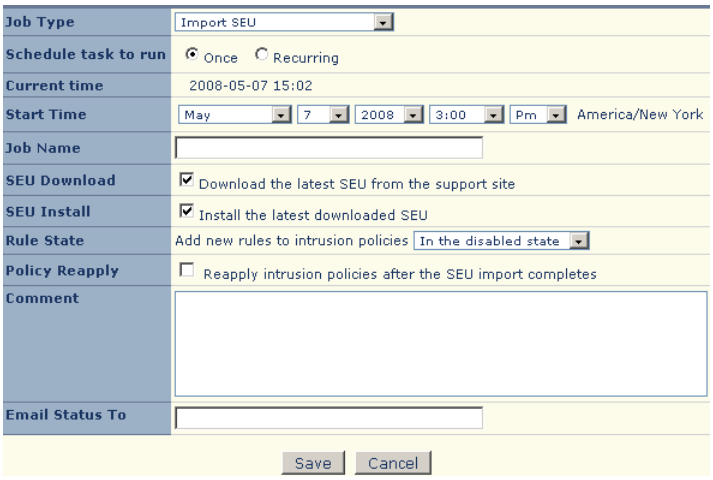
   The Scheduling page appears.
2. Click **Add Task**.

   The Add Task page appears.

Keysight, EX-1004
Page 1717

3. From the **Job Type** list, select **Import SEU**.

   The page reloads to show the options for importing SEUs.



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 1700 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. To use this task to download the latest SEU, select **Download the latest SEU from the support site**.
7. To use this task to install the latest downloaded SEU, select **Install the latest downloaded SEU.**
8. To re-apply intrusion policies after installing an SEU, select **Reapply intrusion policies after the SEU import completes**.
9. 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.

Keysight, EX-1004
Page 1718

10. 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 1751 for more information about configuring a relay host.

11. Click **Save**.

    The task is created.

Automating Intrusion Policy Applications

Platform: DC/Standalone Sensor

Component: IPS

You can automatically apply intrusion policies at scheduled intervals. This feature is useful if you need to use different policies during different times of the day.

IMPORTANT! You cannot apply intrusion policies to a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). See Using Clustered 3D Sensors on page 250 for more information.

To automate intrusion policy application:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

   The Scheduling page appears.

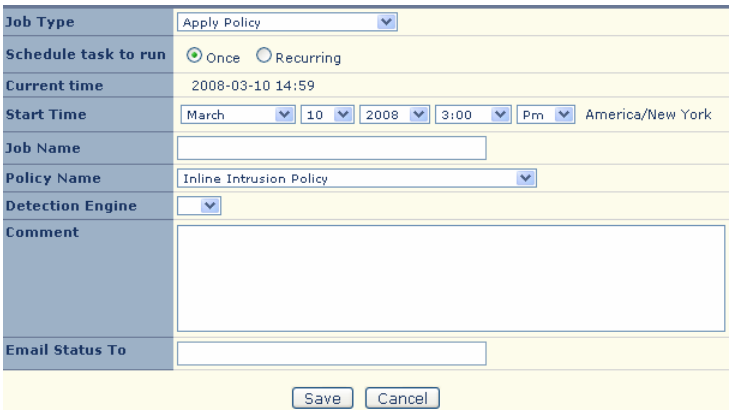
2. Click **Add Task**.

   The Add Task page appears.

Keysight, EX-1004
Page 1719

3. From the **Job Type** list, select **Apply Policy**.

   The page reloads to show the options for applying an intrusion policy.



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 1700 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. In the **Policy Name** field, select the intrusion policy you want to apply from the drop-down list or select **Policy Default** to apply the policy to each detection engine targeted in the policy.
7. In the **Detection Engine** field, select the detection engine where you want to apply the policy.
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 1751 for more information about configuring a relay host.

Keysight, EX-1004
Page 1720

10. Click **Save**.

    The task is created.

    You can check the status of a running task on the Task Status page. See Viewing the Status of Long-Running Tasks on page 2016 for more information.

Automating Reports

Platform: DC/MDC/Standalone Sensor

Component: None

You can automate reports so that they run at regular intervals. However, you must design a profile for your report before you can configure it as a scheduled task. See Creating a Report Profile on page 1542 for more information about using the report designer to create a report profile.

To automate a report:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

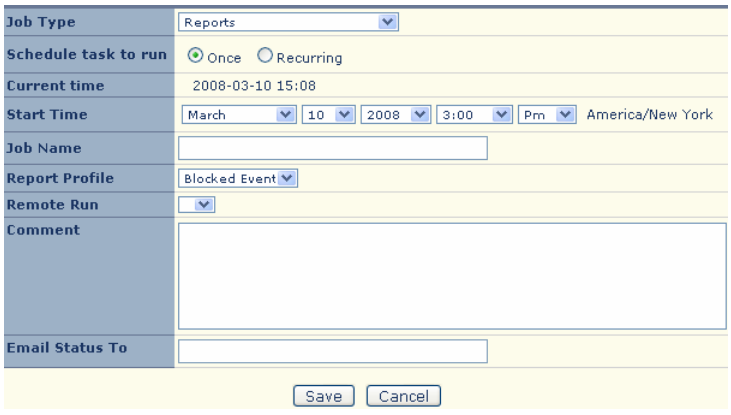
   The Scheduling page appears.

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Reports**.

   The page reloads to show the options for setting up a report to run automatically. The Defense Center version of the page is displayed below.



Keysight, EX-1004
Page 1721

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 1700 for details.
5. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
6. In the **Report Profile** field, select the report profile that you want to use from the drop-down list.

   IMPORTANT! You **cannot** run remote reports on Crossbeam-based software sensors.

7. On a Defense Center, if you want to run the report on a managed sensor, in the **Remote Run** field, select the name of the sensor from the drop-down list.
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 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 1751 for more information about configuring a relay host.
10. Click **Save**.

    The task is created.

Automating Nmap Scans

You can schedule regular Nmap scans of targets on your network. Automated scans allow you to refresh operating system and service information previously supplied by an Nmap scan. Because RNA 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 services on hosts in your network. See the following sections for more information:

- Preparing Your System for an Nmap Scan
- Scheduling an Nmap Scan

Keysight, EX-1004
Page 1722

Note that a Policy & Response Administrator can also use an Nmap scan as a remediation. For example, when an operating system conflict occurs on a host, that conflict can 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 1335.

Preparing Your System for an Nmap Scan

If you have not used the Nmap scanning capability before, you must complete several Nmap configuration steps prior to defining a scheduled scan.

1. Create a scan instance to define the Nmap server to be used by your scan.

   For more information on setting up a Nmap server connection profile, see Creating an Nmap Scan Instance on page 1428.

   IMPORTANT! Make note of the name of the scan instance you create. You need to select this name when prompted for the Nmap Configuration name when setting up the scheduled scan.

2. Create a scan target to define the target hosts and host ports to scan.

   For more information on setting up a scan target, see Creating an Nmap Scan Target on page 1430.

3. Create a remediation definition to define what plugins and Nmap scan settings should be used when the scheduled scan runs.

   For more information on setting up a remediation definition, see Creating an Nmap Remediation on page 1431.

4. Continue with Scheduling an Nmap Scan.

Scheduling an Nmap Scan

Platform: DC

Component: RNA

You can schedule a scan of a host or hosts on your network using the Nmap utility.

Once Nmap replaces a host's operating system or services detected by RNA with the results from an Nmap scan, RNA 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 system and services up to date. If the host is deleted from the network map and re-added, any Nmap scan results are discarded and RNA resumes monitoring of all operating system and service data for the host.

Keysight, EX-1004
Page 1723

To schedule Nmap scanning:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

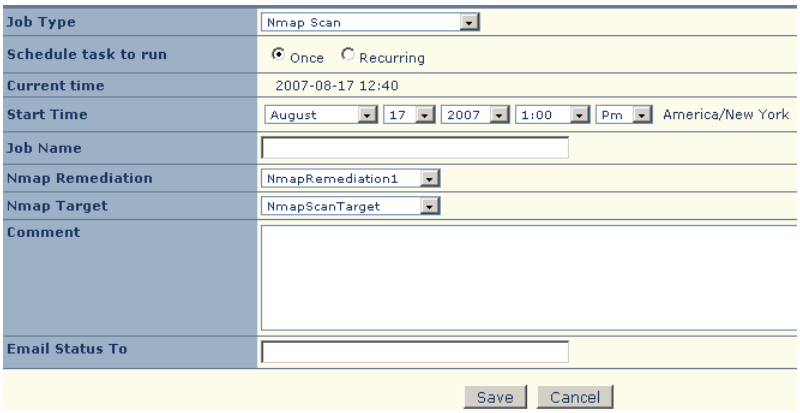
   The Scheduling page appears.

2. Click **Add Task**.

   The Add Task page appears.

3. From the **Job Type** list, select **Nmap Scan**.

   The page reloads to show the options for automating Nmap scans.



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 1700 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.

Keysight, EX-1004
Page 1724

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 1751 for more information about configuring a relay host.

10. Click **Save**.

    The task is created.

Automating Recommended Rule State Generation

Platform: DC

Component: IPS + RNA

You can automatically generate rule state recommendations based on RNA 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 345 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 RNA Rule State Recommendations on page 468. Modified rule states take effect the next time you apply your intrusion policy. See Using RNA Recommendations on page 472 for more information.

To generate recommendations:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

   The Scheduling page appears.

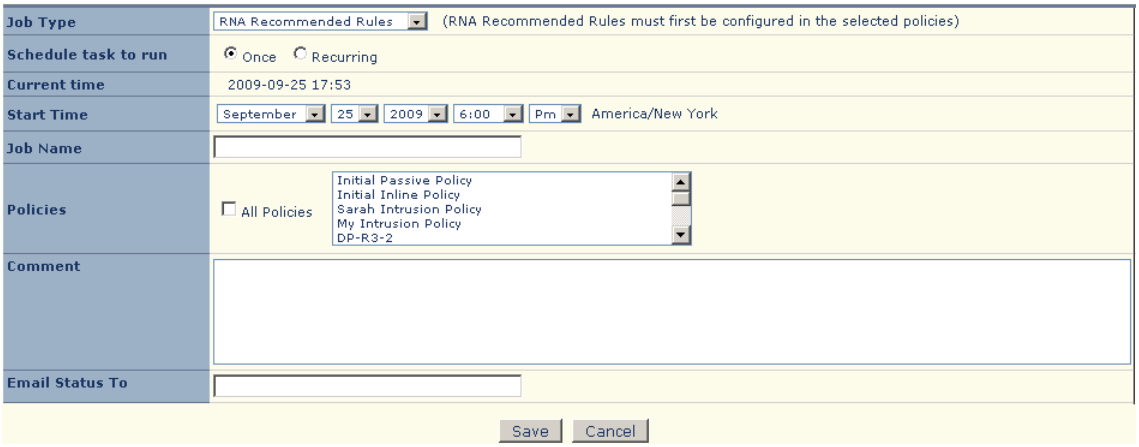
2. Click **Add Task**.

   The Add Task page appears.

Keysight, EX-1004
Page 1725

3. From the **Job Type** list, select **RNA Recommended Rules**.

   The page reloads to show the options for generating RNA-recommended rule states.



4. Optionally, click the **policies** link in the **Job Type** field to display the Detection & Prevention page, where you can configure RNA Recommended Rules in a policy. See Managing RNA Rule State Recommendations on page 468 for more information.
5. 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 1700 for details.
6. In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.
7. 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.

Keysight, EX-1004
Page 1726

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 1751 for more information about configuring a relay host.

10. Click **Save**.

    The task is created.

Viewing Tasks

After adding scheduled tasks, you can view them and evaluate their status. The View Options section of the page allows you to view scheduled tasks using a calendar and a list of scheduled tasks.

See the following sections for more information:

- Using the Calendar on page 1727
- Using the Task List on page 1729

Using the Calendar

Platform: DC/MDC/Standalone Sensor

Component: None

The Calendar view option allows you to view which scheduled tasks occur on which day.

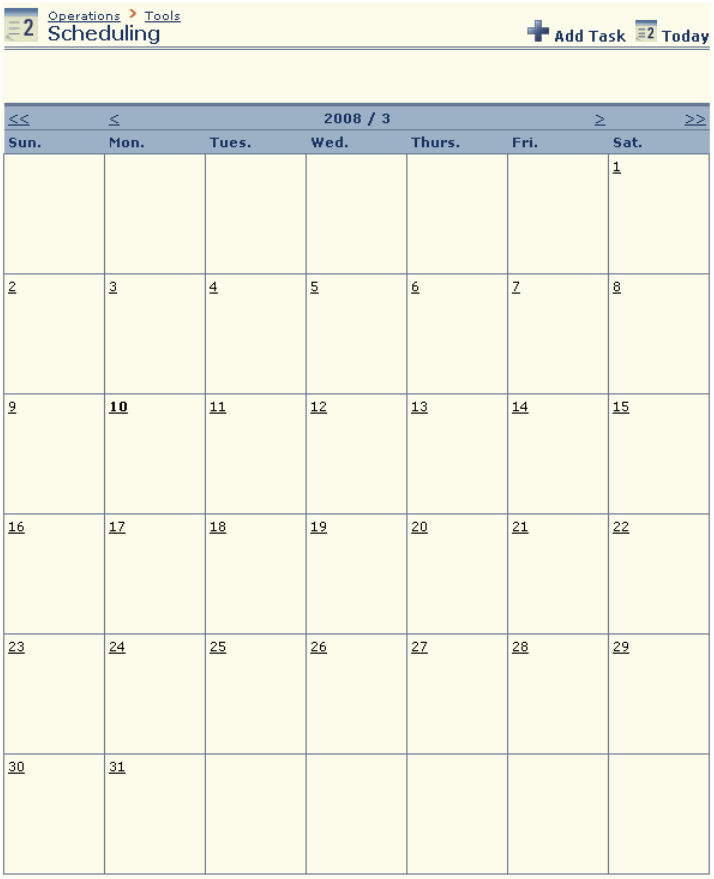
Keysight, EX-1004
Page 1727

To view scheduled tasks using the calendar:

Access: Maint/Admin

1. Select **Operations > Tools > Scheduling**.

   The Scheduling page appears.



2. You can perform the following tasks using the calendar view:
   - Click **<<** to move back one year.
   - Click **<** to move back one month.
   - Click **>** to move forward one month.
   - Click **>>** to move forward one year.
   - Click **Today** to return to the current month and year.
   - Click **Add Task** to schedule a new task.

Keysight, EX-1004
Page 1728

- 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 on page 1729.

Using the Task List

Platform: DC/MDC/Standalone Sensor

Component: None

The Task List shows a list of tasks along with their status. The task list appears at 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 1727 for more information.)

Task List Columns

Column	Description
Name	Displays the name of the scheduled task.
Type	Displays the type of scheduled task.
Start Time	Displays the scheduled start date and time.
Frequency	Displays how often the task is run.
Comment	Displays the comment that accompanies the scheduled task.
Status	Describes the current status for a scheduled task. • A check mark icon indicates that the task ran successfully. • A question mark icon indicates that the task is in an unknown state. • A red ! indicates that the task failed.
Creator	Displays the name of the user that created the scheduled task.
Delete	Deletes the scheduled task.

Keysight, EX-1004
Page 1729

Editing Scheduled Tasks

Platform: DC/MDC/Standalone Sensor

Component: None

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: Maint/Admin

1. Select **Operations > 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 **Edit**.

   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, 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 1702
   - Automating Software Updates on page 1704
   - Automating Vulnerability Database Updates on page 1710
   - Automating SEU Imports on page 1716
   - Automating Intrusion Policy Applications on page 1719
   - Automating Reports on page 1721
   - Automating Nmap Scans on page 1722
   - Automating Recommended Rule State Generation on page 1725

5. Click **Save** to save your edits.

   Your change are saved and the Scheduling page appears again.

Deleting Scheduled Tasks

There are two types of deletions you can perform from the Schedule View page. You can delete a specific one-time task that has not yet run or you can delete every instance of a recurring task. If you delete an instance of a recurring task, all instances of the task are deleted. If you delete a task that is scheduled to run once, only that task is deleted.

Keysight, EX-1004
Page 1730

The following sections describe how to delete tasks:

- To delete all instances of a task, see Deleting a Recurring Task on page 1731.
- To delete a single instance of a task, see Deleting a One-Time Task on page 1731.

Deleting a Recurring Task

Platform: DC/MDC/Standalone Sensor

Component: None

When you delete one instance of a recurring task, you automatically delete all instances of that task.

To delete a recurring task:

Access: Maint/Admin

1. Select **Operations > 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 **Delete**.

   All instances of the recurring task are deleted.

Deleting a One-Time Task

Platform: DC/MDC/Standalone Sensor

Component: None

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: Maint/Admin

1. Select **Operations > 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 **Delete**.

   The instance of the task you selected is deleted.

Keysight, EX-1004
Page 1731

Chapter 40
Managing System Policies

A system policy allows you to manage the following on your Defense Center or 3D Sensor:

- access control lists
- audit log settings
- authentication profiles
- command line interface settings
- 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
- RNA settings
- RUA settings
- SNMP polling settings
- synchronizing time
- serving time from the Defense Center
- mapping vulnerabilities for services

You can use a system policy to control the aspects of your Defense Center that are likely to be similar for other Sourcefire 3D System appliances in your

Keysight, EX-1004
Page 1732

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 sensors that it manages.

You can also benefit from having multiple 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 System Settings on page 1775 for more information.

IMPORTANT! You **cannot** apply system policies to Crossbeam-based software sensors or Intrusion Agents.

See the following sections for more information:

- Creating a System Policy on page 1733
- Editing a System Policy on page 1735
- Applying a System Policy on page 1736
- Deleting System Policies on page 1737

Creating a System Policy

Platform: DC/MDC/Standalone Sensor

Component: None

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 1997.

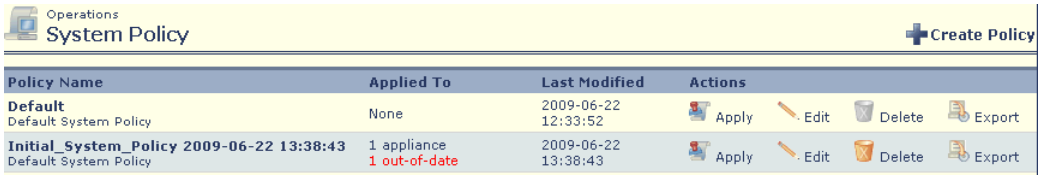
Keysight, EX-1004
Page 1733

To create a system policy:

Access: Admin

1. Select **Operations > System Policy**.

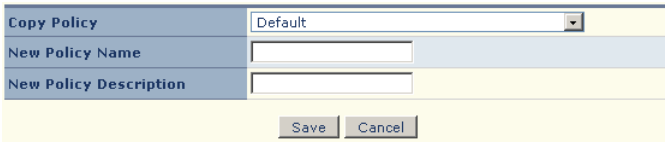
   The System Policy page appears.



   The **Policy Name** column includes its 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 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 and description (up to 40 alphanumeric characters and spaces each) for your new policy.
5. Click **Save**.

   Your system policy is saved and the Access List 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 1738
   - Configuring Audit Log Settings on page 1740
   - Configuring Authentication Profiles on page 1741
   - Configuring Command Line Interface Settings on page 1745
   - Configuring Dashboard Settings on page 1746
   - Configuring Database Event Limits on page 1747
   - Configuring DNS Cache Properties on page 1750
   - Configuring a Mail Relay Host and Notification Address on page 1751
   - Configuring Intrusion Policy Preferences on page 1753
   - Specifying a Different Language on page 1754
   - Adding a Custom Login Banner on page 1755

Keysight, EX-1004
Page 1734

- Configuring RNA Settings on page 1756
- Configuring RUA Settings on page 1764
- Configuring SNMP Polling on page 1766
- Synchronizing Time on page 1768
- Serving Time from the Defense Center on page 1771
- Mapping Vulnerabilities for Services on page 1773

Editing a System Policy

Platform: DC/MDC/Standalone Sensor

Component: None

You can edit a system policy that is currently in use, but remember to re-apply the policy as explained in Applying a System Policy on page 1736.

To edit an existing system policy:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears, including a list of the existing system policies.

2. Click **Edit** next to the system policy that you want to edit.

   With the **Policy Name** and **Policy Description** fields at the top, Access List, the first section of the system policy, appears. You can change the policy name and description. 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 1738
   - Configuring Audit Log Settings on page 1740
   - Configuring Authentication Profiles on page 1741
   - Configuring Command Line Interface Settings on page 1745
   - Configuring Dashboard Settings on page 1746
   - Configuring Database Event Limits on page 1747
   - Configuring DNS Cache Properties on page 1750
   - Configuring a Mail Relay Host and Notification Address on page 1751
   - Configuring Intrusion Policy Preferences on page 1753
   - Specifying a Different Language on page 1754
   - Adding a Custom Login Banner on page 1755
   - Configuring RNA Settings on page 1756
   - Configuring RUA Settings on page 1764

Keysight, EX-1004
Page 1735

- Configuring SNMP Polling on page 1766
- Synchronizing Time on page 1768
- Serving Time from the Defense Center on page 1771
- Mapping Vulnerabilities for Services on page 1773

IMPORTANT! If you are editing the current system policy, make sure you apply the updated policy when you are finished. See Applying a System Policy on page 1736.

Applying a System Policy

Platform: DC/MDC/Standalone Sensor

Component: None

After you create or edit a system policy, your settings do not take effect until you apply it.

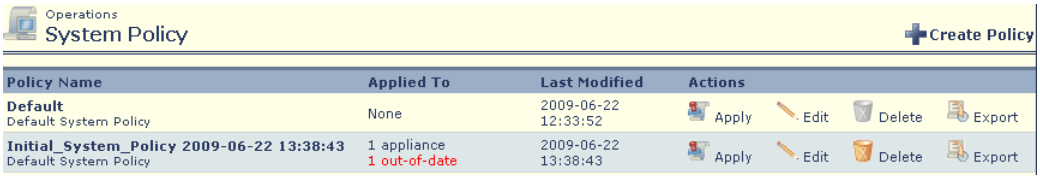
IMPORTANT! You **cannot** apply system policies to Crossbeam-based software sensors or Intrusion Agents.

To apply a system policy:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears, including a list of the existing system policies.



2. Click **Apply** next to the system policy that you want to apply.

   On the 3D Sensor, the system policy is applied.

   On the Defense Center, the Apply page appears. If a policy has been updated since it was applied, the name of the policy appears in italics.

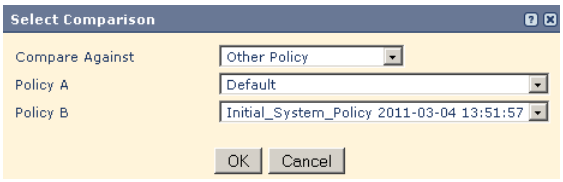
Keysight, EX-1004
Page 1736

3. On the Defense Center, select the sensors, and, if required, the Defense Center itself, where you want to apply the system policy.

   TIP! You can sort the sensors by sensor group, model, type of sensor, or previously applied policy. You can also select an entire group.

4. Click **Apply**.

   A message appears indicating that the task is added to the task queue.



Deleting System Policies

Platform: DC/MDC/Standalone Sensor

Component: None

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 **Operations > System Policy**.

   The System Policy page appears, including a list of the existing system policies.

2. Click **Delete** next to the system policy that you want to delete.

   The policy is deleted.

Keysight, EX-1004
Page 1737

Configuring the Sections of Your System Policy

Platform: DC/MDC/Standalone Sensor

Component: None

You can change various sections of your system policy. 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 1738
- Configuring Audit Log Settings on page 1740
- Configuring Authentication Profiles on page 1741
- Configuring Command Line Interface Settings on page 1745
- Configuring Dashboard Settings on page 1746
- Configuring Database Event Limits on page 1747
- Configuring DNS Cache Properties on page 1750
- Configuring a Mail Relay Host and Notification Address on page 1751
- Configuring Intrusion Policy Preferences on page 1753
- Specifying a Different Language on page 1754
- Adding a Custom Login Banner on page 1755
- Configuring RNA Settings on page 1756
- Configuring RUA Settings on page 1764
- Synchronizing Time on page 1768
- Serving Time from the Defense Center on page 1771
- Mapping Vulnerabilities for Services on page 1773

Configuring the Access List for Your Appliance

Platform: DC/MDC/Standalone Sensor

Component: None

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.

Keysight, EX-1004
Page 1738

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 policy. In either case, the access list does not take effect until you apply the system policy.

To configure the access list:

Access: Admin

1. Select **Operations > System Policy**.

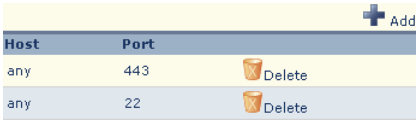
   The System Policy page appears.

2. You have two options:

   - To modify the access list in an existing system policy, click **Edit** 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 1733, and click **Save**.

   In either case, the Access List page appears.



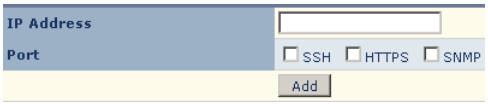
3. Optionally, to delete one of the current settings, click **Delete**.

   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 if 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**.

   The Add IP Address page appears.



5. In the **IP Address** field, use the following syntax 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 59.
   - `any`, to designate any IP address

Keysight, EX-1004
Page 1739

6. Select **SSH**, **HTTPS**, **SNMP**, or a combination of those options to specify which ports you want to enable for these IP addresses, then click **Add**.

   The Access List page appears again, reflecting the changes you made.

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 1736 for more information.

Configuring Audit Log Settings

Platform: DC/MDC/Standalone Sensor

Component: None

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 name of the sending host is part of the sent information and 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.

To configure the audit log settings:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy Page appears.

2. You have two options:

   - To modify the audit log settings in an existing system policy, click **Edit** 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 1733, and click **Save**.

   In either case, the Access List page appears.

3. Click **Audit Log Settings**.

4. Select **Enabled** next to **Send Audit Log to Syslog**.

   The default setting is **Disabled**.

Keysight, EX-1004
Page 1740

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! The computer you configure to receive an audit log must be set up to accept remote messages. Otherwise, the appliance may the send audit log to the host, but it will not be accepted.

6. Label the audit data that you are sending with a facility and severity.

   The default for **Facility** is `USER`. The default for **Severity** is `INFO`. However, you can select any of the standard syslog facility and severity settings.

7. Optionally, insert a reference tag in the **TAG** field.

8. 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 sensors. See Applying a System Policy on page 1736 for more information.

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
```

Configuring Authentication Profiles

Platform: DC/MDC

Component: None

Normally, when a user logs into a Sourcefire 3D System Defense Center or managed sensor, the appliance verifies the user credentials by comparing them to a user account stored in the Defense Center or managed sensor'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 sensor authenticate to that server rather than using the local database.

When you apply a 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

Keysight, EX-1004
Page 1741

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 both the Intrusion Event Analyst role and the RNA Event Analyst 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 1629.

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 1645 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 1684. 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 40.

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 re-apply, any user accounts created before the modification retain the first user role until you modify or delete and recreate them.

The Authentication Profiles page only displays in the system policy on a Defense Center. You can enable authentication in a system policy on your Defense Center and then push that policy to managed sensors. Once you apply the policy to a sensor, eligible externally authenticated users can log into the sensor. However, the system policy on the sensor does not display authentication profile settings, so you cannot manage them on the sensor itself. To make changes to the authentication profile settings, you have to modify the policy on the Defense Center and then push it to the sensor again. To disable authentication on a managed sensor, you can either disable it in a system policy on the Defense Center and push that to the sensor or apply a local system policy (which cannot contain authentication profile settings) on the sensor.

Keysight, EX-1004
Page 1742

Note that you can only enable external authentication on Defense Centers and 3D Sensors. Enabling external authentication by applying a system policy is not supported on the following sensor types:

- managed 3Dx800 sensors
- Crossbeam-based software sensors
- Intrusion Agents
- managed RNA Software for Red Hat Linux

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. On the Defense Center, select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the authentication profile settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the authentication profile settings as part of a new system policy, click **Create Policy**.

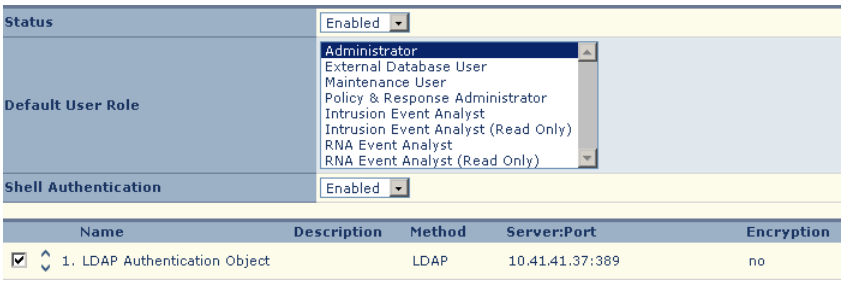
     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

Keysight, EX-1004
Page 1743

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 a user role to define the default permissions you want to grant to users authenticated externally.

   TIP! Press Ctrl before selecting roles to select multiple default user roles. Note that although you can select both an event analyst role and the corresponding read-only event analyst role, only the 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. To enable use of an authentication object, select the check box next to the object.

   IMPORTANT! You must enable at least one authentication object to enable external authentication.

8. Optionally, use the up and down arrows to change the order in which authentication servers are accessed when an authentication request occurs. Remember that shell access users can only authenticate against the server whose authentication object is highest in the profile order.
9. 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 sensors. See Applying a System Policy on page 1736 for more information.

Keysight, EX-1004
Page 1744

Configuring Command Line Interface Settings

Platform: DC/MDC/Standalone Sensor

Component: None

If you are using a Master Defense Center or Defense Center to manage a Virtual 3D Sensor or an 8000 Series sensor or using an 8000 Series sensor in standalone mode, you can use the command line settings in the system policy to set the period of inactivity after which a shell session times out or to restrict command line activity on the appliance to only the commands available through the command line interface.

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.

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.

For more information on command line interface commands, see Command Line Reference on page 2044. For information on setting up users for command line access, see Managing Command Line Access on 8000 Series Sensors on page 1678 and user on page 2059 (for Virtual 3D Sensor CLI user management).

To manage command line access settings:

Access: Admin

1. Select **Operations > System Policy.**

   The System Policy page appears.

2. You have two options:

   - To modify the command line access settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the command line access 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 1733, and click **Save.**

   In either case, the **Access List** page appears.

3. Click **Command-Line Interface**.

   The Command-Line Interface page appears.



Keysight, EX-1004
Page 1745

4. Set a **Shell Timeout** by typing the number of seconds that a shell session can remain inactive before timing out and requiring the user to log back in to continue using the shell.

   Note that the default setting of 0 disables timeout of shell sessions.

5. To permanently prevent any user on the appliance where you apply the policy from entering expert mode on the command line, enable **Permanently Disable Expert Access**.

   WARNING! Remember that you must contact Sourcefire Support to restore the expert mode capability if you enable this option and apply the policy to an appliance.

6. Click **Save Policy and Exit.**

   The system policy is updated. Your changes do not take effect until you apply the system policy. See Deleting System Policies on page 1737 for more information.

Configuring Dashboard Settings

Platform: DC/MDC/Standalone Sensor

Component: None

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 71 for more information on how to use custom widgets.

To enable Custom Analysis widgets:

Access: Admin

1. Select **Operations > System Policy.**

   The System Policy page appears.

2. You have two options:

   - To modify the dashboard settings in an existing system policy, click **Edit** 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 1733, and click **Save.**

   In either case, the **Access List** page appears.

Keysight, EX-1004
Page 1746

3. Click **Dashboard.**

   The Dashboard Settings page appears.



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.

   By default, the Custom Analysis widget use is enabled.

5. Click **Save Policy and Exit.**

   The system policy is updated. Your changes do not take effect until you apply the system policy. See Deleting System Policies on page 1737 for more information.

Configuring Database Event Limits

Platform: DC/MDC/Standalone Sensor

Component: None

You can use the Database page to specify the maximum number of events you want to store on an appliance. To improve performance, you should try to tailor the database event limit to the number of events you regularly work with.

In most cases, the minimum number of records you can store in any database is one record (or, in the case of the compliance violation history database, one day's history). However, for some databases, you can choose not to store any events. These databases include those that store RNA and RUA events, as well as flow events, flow summaries, and health events.

The Database Event Limits on page 1748 below describes the maximum number of records you can store in the databases on your appliance. Note that if you apply a system policy to an appliance that does not support the maximum limit you specify (for example, if you specify 100 million intrusion events and apply that policy to a 3D Sensor), the maximum limit for the appliance is silently enforced.

In addition, database limits that do not apply to a particular appliance are silently ignored. For example, if you use the Defense Center to apply the same system

Keysight, EX-1004
Page 1747

policy to itself and the 3D Sensors it manages, any health alert limits you set in the policy have no effect on the sensors.

Database Event Limits

The...	Is the database that stores...	And can store up to...
Intrusion Event Database (Defense Center or Master Defense Center)	intrusion events on a Defense Center or on a Master Defense Center (which is always a DC3000)	2.5 million events on the DC500 10 million events on the Virtual Defense Center or the DC1000 100 million events on the DC3000
Intrusion Event Database (3D Sensor)	intrusion events on a 3D Sensor	2 million events
RNA Event Database	RNA network discovery events on a Defense Center	10 million events
RNA Flow Database	RNA flows on a Defense Center	10 million events on the DC500, Virtual Defense Center, or DC1000 100 million events on the DC3000
RNA Flow Summary Database	RNA flow summaries (aggregated RNA flows) on a Defense Center	10 million events on the DC500, Virtual Defense Center, or DC1000 100 million events on the DC3000
Compliance & White List Event Database	compliance events and white list events on a Defense Center or Master Defense Center	1 million events
Health Event Database	health events on a Defense Center or Master Defense Center	1 million events
Audit Event Database	audit records	100,000 records
Remediation Status Event Database	remediation status events on a Defense Center	10 million events
White List Violation History Database	the white list violation history of the hosts on your network, on a Defense Center	a 30-day history of violations

Keysight, EX-1004
Page 1748

Database Event Limits (Continued)

The...	Is the database that stores...	And can store up to...
RUA Event Database	RUA events on a Defense Center	10 million events
RUA History Database	RUA storage of user logins on a Defense Center	10 million user login records
SEU Import Log Database	SEU import log records	1 million records

Note that 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 85% of its capacity, unified files are deleted from the system, beginning with the oldest files. See Configuring a Mail Relay Host and Notification Address on page 1751 for information about generating automated email notifications when events are automatically pruned.

For information on manually pruning the RNA and RUA databases, see Purging the RNA and RUA Databases on page 2014.

In addition, you can configure an email address that will receive notifications when intrusion events and audit logs are pruned from the database.

To configure the maximum number of records in the database:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the database settings in an existing system policy, click **Edit** 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 1733, and click **Save**.

   In either case, the Access List page appears.

3. Click **Database**.

   The Database page appears. 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 Database Event Limits on page 1748.

Keysight, EX-1004
Page 1749

4. Optionally, in the **Data Pruning Notification Address** field, enter the email address you want to receive notifications when intrusion events and audit logs are pruned from the appliance's database.
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 1736 for more information.

Configuring DNS Cache Properties

Platform: DC/MDC/Standalone Sensor

Component: None

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 an administrator, 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

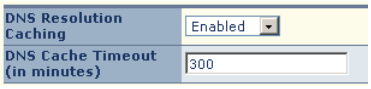
1. Select **Operations > System Policy**.

   The System Policy page appears.
2. You have two options:
   - To modify the DNS cache settings in an existing system policy, click **Edit** 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 1733, and click **Save**.

   In either case, the Access List page appears.
3. Click **DNS Cache**.

   The DNS Cache page appears.



Keysight, EX-1004
Page 1750

4. Next to **DNS Resolution Caching**, select **Enabled** to enable caching or **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 1800. For information about configuring event preferences, see Configuring Event View Settings on page 45.

5. In the **DNS Cache Timeout** 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 1736 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

Platform: DC/MDC/Standalone Sensor

Component: None

You must configure a mail host if you plan to:

- email event-based reports
- email status reports for scheduled tasks
- use email for RNA event, impact flag, and compliance event alerting (Defense Center only - requires RNA)
- use email for intrusion event alerting (Defense Center only - requires IPS)
- use email for health event alerting (Defense Center only)

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.

Keysight, EX-1004
Page 1751

To configure a mail relay host:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

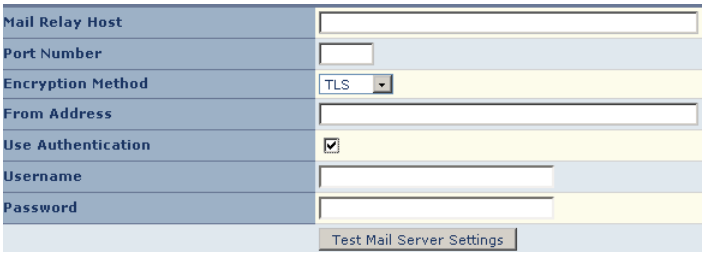
   - To modify the email settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the email settings as part of a new system policy, click **Create Policy**.

     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

3. Click **Email Notification**.

   The Configure Email Notification page appears.



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 and 465 or 587 when using SSLv3 or TLS.

6. Select an encryption method:

   - To encrypt communications between the appliance and the mail server using Transport Layer Security, select **TLS**.
   - To encrypt communications between the appliance and the mail server using Secure Socket Layers, select **SSLv3**.
   - To allow unencrypted communication between the appliance and the mail server, select **None**.

   Note that certificate validation is not required for encrypted communication between the appliance and mail server.

Keysight, EX-1004
Page 1752

7. In the **From Address** field, enter a valid email address 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, click **Use Authentication** and enter a user name and password.
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 1736 for more information.

Configuring Intrusion Policy Preferences

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can allow or require comments to be added to the audit log when an intrusion policy changes. You can also track all changes to intrusion policies in the audit log.

To configure intrusion policy change tracking:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.
2. You have two options:
   - To modify the intrusion policy preferences in an existing system policy, click **Edit** 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 1733, and click **Save**.

   In either case, the Access List page appears.
3. Click **Intrusion Policy Preferences**.

   The Intrusion Policy Preferences page appears.



Keysight, EX-1004
Page 1753

4. Select Disabled, Optional, or Required from the **Comments on policy change** drop-down list.

   If you select Optional or Required, a **Description of Changes** text box appears when you commit your intrusion policy changes.

5. Optionally, if you want to track changes to intrusion policies, 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 1736 for more information.

Specifying a Different Language

Platform: DC/MDC/Standalone Sensor

Component: None

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 **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the language settings in an existing system policy, click **Edit** 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 1733, and click **Save**.

   In either case, the Access List page appears.

3. Click **Language**.

   The Language page appears.



4. Select the language you want to use.

Keysight, EX-1004
Page 1754

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 1736 for more information.

Adding a Custom Login Banner

Platform: DC/MDC/Standalone Sensor

Component: None

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 (>).

Custom login banners are part of the system policy. You can specify the login banner either by creating a new system policy or by editing an existing policy. In either case, the login banner is not used until you apply the system policy.

To add a custom banner:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the login banner in an existing system policy, click **Edit** next to the system policy.
   - To configure the login banner as part of a new system policy, click **Create Policy**.

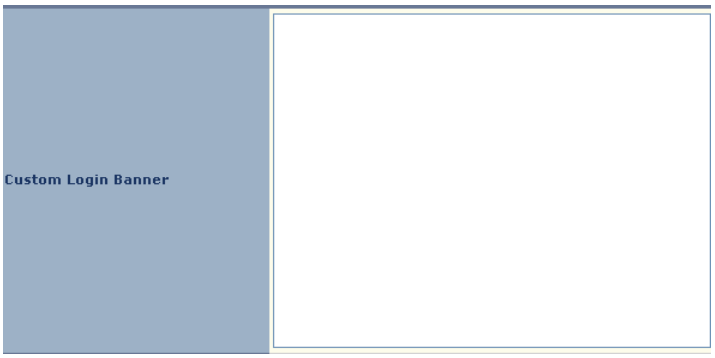
     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

Keysight, EX-1004
Page 1755

3. Click **Login Banner**.

   The Login Banner page appears.



4. In the **Custom Login Banner** field, enter the login banner that 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 1736 for more information.

Configuring RNA Settings

Platform: DC/ MDC

Component: RNA

You can configure several aspects of RNA behavior through the system policy, including various settings for the RNA detection policy and subnet detection, how RNA data is stored, how the Sourcefire 3D System performs impact flag correlation with intrusion events, and how RNA deals with operating system and services identity conflicts.

To configure RNA settings:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

Keysight, EX-1004
Page 1756

2. You have two options:
   - To modify the RNA settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the RNA 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 1733, and click **Save**.

   In either case, the Access List page appears.
3. Click **RNA Settings**.

   The RNA Settings page appears.
4. Specify RNA detection policy preferences.You must decide whether you want to:
   - automatically generate subnet recommendations for your currently applied RNA detection policies on a daily basis
   - automatically update those policies and apply the updated policies to your RNA detection engines
   - require users to confirm their actions when manually applying a detection policy

   See Understanding RNA Detection Policy Settings on page 1757.
5. Specify RNA data storage settings, including timeout values, how you want to treat duplicate flow data, the types of events you want to log, and so on.

   See Understanding RNA Data Storage Settings on page 1759.
6. If your deployment includes IPS, specify how the Sourcefire 3D System performs impact flag correlation with intrusion events.

   See Understanding Vulnerability Impact Assessment Settings on page 1762.
7. Configure multiple fingerprint settings to manage operating system and service source priorities and identity conflict resolution.

   See Understanding Multiple Fingerprint Settings on page 1762.
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 1736 for more information.

Understanding RNA Detection Policy Settings

A key step in deploying RNA is creating and applying an RNA detection policy to your RNA detection engines. As a precaution, you can require users to confirm their actions when manually applying a detection policy.

One of the major decisions you must make when configuring detection policies is whether to use subnet detection. This feature allows RNA to make

Keysight, EX-1004
Page 1757

recommendations about which detection engines you should use to analyze your network traffic, based on which hosts are closest (in hops) to each detection engine.

As RNA continuously monitors your network traffic, it may be able to refine any subnet recommendations it has made for your RNA detection policies, especially if your network configuration has been altered through routing or host changes. As a time-saving and performance-maximizing measure, you can configure RNA to automatically generate subnet recommendations for your currently applied RNA detection policies on a daily basis. Optionally, you can configure the Defense Center to automatically update those policies and apply the updated policies to your RNA detection engines.



The following table describes the settings you can use to enable daily generation and application of subnet recommendations for your currently applied RNA detection policies, as well as the setting to require users to confirm their actions when manually applying a detection policy

RNA Detection Policy Settings

Field	Description
Require confirmation before applying policy?	Select this check box if you want to force users to confirm their actions each time they apply RNA detection policies. If you enable this option, the appliance also warns you if the detection engine has a different policy applied to it than the one you are attempting to apply.
Generate Recommendations Daily At	Select the time when you want RNA to automatically generate daily subnet recommendations for all applied RNA detection policies. To disable daily generation of subnet recommendations, select **Disabled**.

Keysight, EX-1004
Page 1758

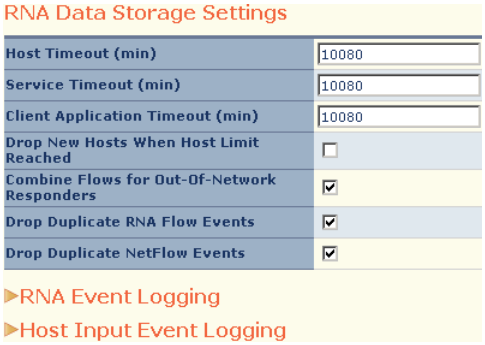
RNA Detection Policy Settings (Continued)

Field	Description
Mail Notifications To	Enter the email address where you want to receive notifications of new subnet recommendations. **TIP!** To receive email notifications, you must configure a valid mail relay host; see Configuring a Mail Relay Host and Notification Address on page 1751.
Automatically Apply Daily Recommendations	Enable this check box to automatically update and apply your RNA detection policies after RNA generates subnet recommendations. This option has no effect unless you enable daily recommendations. Note that you cannot apply an RNA detection policy to the combined detection engines on a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). Auto-apply does not occur for clustered pairs with mismatched versions.

For more information on RNA detection policies and the subnet detection process, see Introduction to Sourcefire RNA on page 878.

Understanding RNA Data Storage Settings

RNA data storage settings control the kinds of RNA data stored in the database, and therefore determine the data that other parts of the Sourcefire 3D System can use. These settings also control how long data is retained in the network map.



The following options comprise the RNA data storage settings.

Keysight, EX-1004
Page 1759

Host Timeout

The amount of time that passes, in minutes, before RNA drops a host from the network map due to inactivity. The default setting is 10080 minutes (7 days).

To avoid premature timeout of hosts, make sure that the host timeout value is longer than the update interval in the RNA detection policy. For more information, see Creating RNA Detection Policies on page 900.

Service Timeout

The amount of time that passes, in minutes, before RNA drops a service from the network map due to inactivity. The default setting is 10080 minutes (7 days).

To avoid premature timeout of services, make sure that the service timeout value is longer than the update interval in the RNA detection policy. For more information, see Creating RNA Detection Policies on page 900.

Client Application Timeout

The amount of time that passes, in minutes, before RNA drops a client application from the network map due to inactivity. The default setting is 10080 minutes(7 days).

You should make sure that the client application timeout value is longer than the update interval in the RNA detection policy. For more information, see Creating RNA Detection Policies on page 900.

Drop New Hosts When Host Limit Reached

Select this check box if you want new hosts rather than old hosts dropped when the Defense Center reaches its host limit 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.

Combine Flows for Out-Of-Network Responders

Select this check box if you want you want to combine flow summaries involving external hosts.

Enabling this option treats flow summary data from IP addresses that are not in your list of monitored networks (as defined by your RNA detection policy) as coming from a single host. Event views, graphs, and reports use `external` to indicate the hosts outside your monitored network, instead of an individual IP address.

The Defense Center will combine flow summaries involving a host on your monitored network and one or more external hosts if the flows use the same port, protocol, service, and if they were detected by the same detection engine (for flows detected by 3D Sensor) or were exported by the same NetFlow-enabled device and were processed by the same detection engine.

Keysight, EX-1004
Page 1760

This can reduce the space required to store flow data and can also speed up the rendering of flow data graphs. However, if you enable this option and you attempt to drill down to the table view of flow data (that is, access data on individual flows) for a flow summary that involves an external responder, the table view contains no information.

Note that you can also use the RNA detection policy to force your 3D Sensors to combine flow summaries involving external hosts **before** they transmit the data to the Defense Center, which can reduce the number of events sent to the Defense Center. However, keep in mind that setting this option in the RNA detection policy requires that your flow data mode is **Summary**, which prevents your 3D Sensors from transmitting individual flows to the Defense Center and therefore prevents you from taking advantage of any feature that requires data from individual flows. For more information, see Combining Flow Summaries from External Responders on page 1077 as well as Configuring RNA Detection Policy Settings on page 902.

Drop Duplicate RNA Flow Events

Select this check box if you want the Defense Center to drop duplicate flow events generated by 3D Sensors with RNA.

Duplicate flow events can be created if you use two RNA detection policies, each of which is monitoring a separate network segment using separate detection engines. In that scenario, each detection engine generates a flow event when RNA detects that a connection is terminated between a monitored host on one of the networks and a monitored host on the other network. On the other hand, if you use one policy to monitor both networks, only the reporting detection engine for the flow initiator generates a flow event.

Duplicate flow events can also be created if you overlap network segment coverage with your RNA detection engines in your RNA detection policy.

Note that best practices are to use only one detection policy and to not overlap network segment coverage; not following best practices can degrade performance as the Defense Center attempts to resolve the conflicts, and can also use excessive bandwidth.

Drop Duplicate NetFlow Events

Select this check box if you want the Defense Center to drop duplicate flow events that are based on NetFlow data. Duplicate NetFlow events can be created, for example, if two NetFlow-enabled devices export information about the same session.

Just as with RNA flow events, best practices are to avoid creating duplicate NetFlow events, For more information, see Drop Duplicate RNA Flow Events on page 1761.

Keysight, EX-1004
Page 1761

RNA Event Logging and Host Input Event Logging

Expand the **RNA Event Logging** and **Host Input Event Logging** sections to specify the RNA network discovery and RNA host input event types you want to log in the database.

See Understanding RNA Network Discovery Event Types on page 1005 and Understanding RNA Host Input Event Types on page 1010 for information about each event type

Understanding Vulnerability Impact Assessment Settings

If your deployment includes IPS, you can configure how the Sourcefire 3D System performs impact flag correlation with intrusion events.



Your options are as follows:

- Select the **Use RNA Vulnerability Mappings** check box if you want to use RNA vulnerability information to perform impact flag correlation.
- Select the **Third Party Vulnerability Mappings** check box if you want to use third-party vulnerability references to perform impact flag correlation. For more information, see Mapping Third-Party Vulnerabilities on page 1412 or the *Sourcefire 3D System Host Input API Guide*.

You can select either of the check boxes; if IPS generates an intrusion event and the Sourcefire 3D System can use either of the methods you specified to determine that the host involved in the event is vulnerable to the attack or exploit, the intrusion event will be marked with the red (Vulnerable) impact flag.

Note that if you clear both check boxes, intrusion events will never be marked with the red impact flag. For more information, see Using Impact Flags to Evaluate Events on page 307.

Understanding Multiple Fingerprint Settings

RNA matches fingerprints for operating systems and services against patterns in traffic to determine what operating system and which services are running on a particular host. To provide the most reliable operating system and service identity information, RNA collates fingerprint information from several sources.

RNA uses all passive data to derive operating system identities and assign a confidence value. For more information on current identities and how RNA selects the current identity, see Enhancing Your Network Map on page 1358.

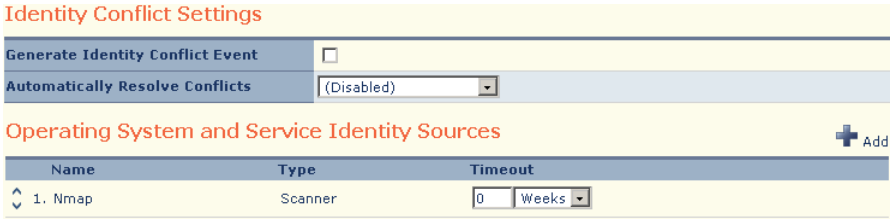
By default, unless there is an identity conflict, identity data added by a scanner or application overrides identity data detected by RNA. You can use the Multiple Fingerprinting page to rank scanner and application fingerprint sources by priority. RNA retains one identity for each source, but only data from the highest priority

Keysight, EX-1004
Page 1762

application or scanner source is used as the current identity. Note, however, that user input data overrides scanner and application data regardless of priority.

An identity conflict occurs when RNA detects an identity that conflicts with an existing identity that came from the active scanner or application sources listed on the Multiple Fingerprinting page or from a 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 RNA 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.

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 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 RNA vulnerabilities in the network map. For more information, see Mapping Third-Party Vulnerabilities on page 1412.



The following options comprise the RNA multiple fingerprint settings.

Generate Identity Conflict Event

Enable this option to generate an event when an identity conflict occurs on a host in the network map.

Keysight, EX-1004
Page 1763

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 RNA fingerprint when an identity conflict occurs, select **Passive** 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 **Active** from the **Automatically Resolve Conflicts** drop-down list.

Scanner/ Application List

You have several options:

- To add a new source, click **Add**. Type a name for the source.
- To delete a source that you added, click **Delete** next to the source.
- To change the type of source, select **Scanner** or **Application** from the **Type** drop-down list.
- 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.
- To promote a source and cause the operating system and service identities to be used in favor of sources below it in the list, click the up arrow next to the source name.
- To demote a source and cause the operating system and service identities to be used only if there are no identities provided by sources above it in the list, click the down arrow next to the source name.

Configuring RUA Settings

Platform: DC

Component: RUA

You can use the RUA settings in the system policy to filter which types of network activity cause RUA to add users to the database.

Sourcefire RUA (see Using Sourcefire RUA on page 1452) is an optional component of the Sourcefire 3D System that allows you to correlate network activity with user identity information. When RUA detects a user login for a user who is not already in the database, an RUA user is added to the Defense Center user database. RUA can add users to the database using the following types of detected protocols:

- LDAP
- AIM

Keysight, EX-1004
Page 1764

- POP3
- IMAP
- Oracle
- SIP (VoIP)

Note that although RUA detects SMTP logins, the Defense Center does not record them unless there is already a user with a matching email address in the database; RUA users are **not** added to the database based on SMTP logins.

The RUA feature license on the Defense Center (see Licensing RUA on page 1475) specifies the number of users you can monitor with RUA. After you reach your licensed limit, RUA stops adding new users to the Defense Center database.

Restricting RUA helps minimize username clutter and preserve RUA licenses. For example, obtaining usernames through protocols such as AIM, POP3, and IMAP can introduce usernames not relevant to your organization due to network access from contractors, visitors, and other guests. In addition, AIM, Oracle, and SIP logins always create duplicate user records. This is because these logins are not associated with any of the user metadata that RUA obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your 3D Sensors detect.

IMPORTANT! Sourcefire RUA Agents installed on Microsoft Active Directory LDAP servers collect **only** LDAP user login information. Therefore, unless your RUA implementation includes 3D Sensors with RUA, filtering non-LDAP logins has no effect. For more information on RUA Agents and 3D Sensors with RUA, see How Do I Choose an RUA Implementation? on page 1463.

To filter RUA users based on network activity type:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the RUA settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the RUA settings as part of a new system policy, click **Create Policy**.

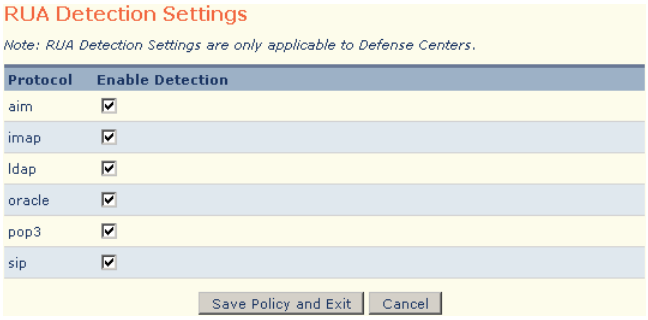
     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

Keysight, EX-1004
Page 1765

3. Click **RUA Settings**.

   The RUA Detection Settings page appears.



4. Select the check boxes that correspond to the types of logins that will create RUA users.

   By default, all login types cause RUA to add users to the database.

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 1736 for more information.

Configuring SNMP Polling

Platform: DC/MDC/Standalone Sensor

Component: None

You can enable Simple Network Management Protocol (SNMP) polling of an appliance through the system policy. If you enable SNMP polling on the appliance, your network management system can poll the appliance for basic appliance information. The SNMP feature supports use of versions 1, 2, or 3 of the SNMP protocol.

The system policy SNMP feature allows access to the standard management information base (MIB) for the device, which includes system details such as contact, administrative, location, and service information, IP addressing and routing information, and transmission protocol usage statistics. Note that enabling the system policy SNMP feature does not cause the appliance to send SNMP traps; it only makes the information of the standard MIB 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 1738.

Keysight, EX-1004
Page 1766

To configure SNMP polling:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the SNMP polling settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the SNMP polling settings as part of a new system policy, click **Create Policy**.

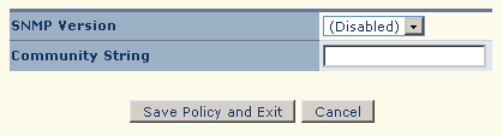
     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   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 1738.

4. Click **SNMP**.

   The SNMP page appears.

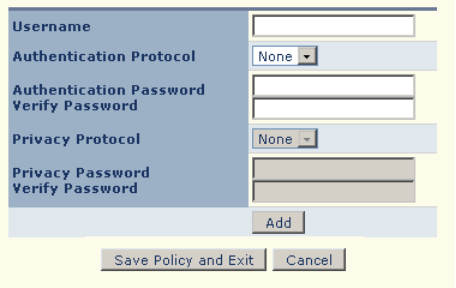


5. To enable SNMP polling, select the SNMP version you plan to use to poll the appliance: **Version 1**, **Version 2**, or **Version 3**.

   When the version is set to Disabled, SNMP polling is disabled.

6. Depending on the SNMP version you selected, you have different options:

   - If you selected **Version 1** or **Version 2**, enter the **Community String** that you want to be supplied by SNMP clients polling the appliance in order to connect. Skip to step 10.
   - If you selected **Version 3**, click **Add** to add an SNMP user. The user definition page appears. Continue with step 7.



Keysight, EX-1004
Page 1767

7. Enter a user name.
8. Optionally, select an authentication protocol: **SHA** for the Secure Hash Algorithm or **MD5** for the Message-Digest algorithm 5. If you select an authentication protocol, enter and confirm a password of at least eight characters.
9. If you selected an authentication protocol, you can also select a privacy protocol: **DES** for the Data Encryption Standard or **AES** for the Advanced Encryption Standard. If you select a privacy protocol, enter and confirm a password of at least eight characters.
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 1736 for more information.

Synchronizing Time

Platform: DC/MDC/Standalone Sensor

Component: None

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 must use native applications, such as command line interfaces or the operating system interface, to manage time settings for software sensors:

- For more information on configuring settings for Crossbeam Systems Switches, see the *Sourcefire 3D Sensor Software for X-Series Installation Guide*.
- You manage time settings on an Intrusion Agent through the operating system.

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. Connections to NTP servers do not use configured proxy settings. To use the

Keysight, EX-1004
Page 1768

Defense Center as an NTP server, see Serving Time from the Defense Center on page 1771.

Sourcefire recommends that you synchronize your virtual appliances to a physical NTP server. Do not synchronize your 3D Sensors (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 3D Sensor. Each procedure is explained separately below.

To synchronize time on the Defense Center:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

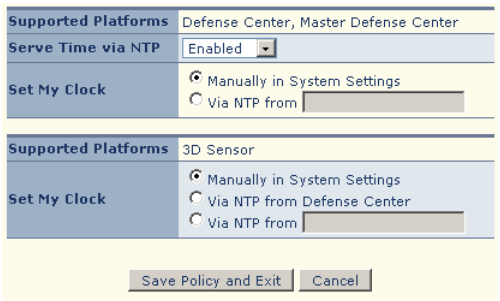
   - To modify the time settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the time settings as part of a new system policy, click **Create Policy**.

     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

3. Click **Time Synchronization**.

   The Time Synchronization page appears.



4. If you want to serve time from the Defense Center to your managed sensors, in the **Serve time via NTP** drop-down list, select **Enabled**.

   Note that if you set this option to **Enabled** and then apply the system policy to a sensor rather than a Defense Center, this value is ignored. Only Defense Centers can act as NTP servers.

Keysight, EX-1004
Page 1769

5. You have two options for specifying how the time is synchronized on the appliance:
   - To set the time manually, select **Manually in System Settings**. See Setting the Time Manually on page 1813 for information about setting the time after you apply the system policy.
   - To receive time through NTP from a different server, select **Via NTP Server 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, you should configure your DHCP server to set the same NTP server.

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 1736 for more information.

   IMPORTANT! It may take a few minutes for the appliance to synchronize with the configured NTP servers.

To synchronize time on a 3D Sensor:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:
   - To modify the time settings in an existing system policy, click **Edit** next to the system policy.
   - To configure the time settings as part of a new system policy, click **Create Policy**.

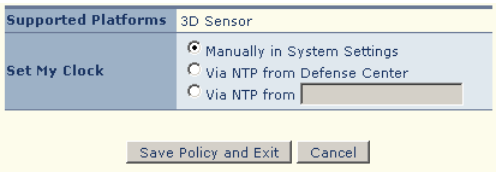
     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

Keysight, EX-1004
Page 1770

3. Click **Time Synchronization**.

   The Time Synchronization page appears.



4. You have two options for specifying how time is synchronized on the 3D Sensor:

   - To set the time manually, select **Manually in System Settings**. See Setting the Time Manually on page 1813 for information about setting the time after you apply the system policy.
   - To receive time through NTP from different servers, select **Via NTP Server from** and, 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.

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 1736 for more information.

IMPORTANT! It may take a few minutes for the 3D Sensor to synchronize with the configured NTP servers. In addition, if you are synchronizing the 3D Sensor 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 3D Sensor.

Serving Time from the Defense Center

Platform: DC/MDC

Component: None

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 3D Sensors.

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,

Keysight, EX-1004
Page 1771

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 sensors will still attempt to synchronize time with the Defense Center. You must disable NTP from the managed sensors' web interfaces to stop the synchronization attempts.

To configure the Defense Center as an NTP server:

Access: Admin

1. On the Defense Center, select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify the NTP server settings in an existing system policy, click **Edit** 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 1733, 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 sensors, 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 sensors. See Applying a System Policy on page 1736 for more information.

   IMPORTANT! It may take a few minutes for the Defense Center to synchronize with its managed sensors.

Keysight, EX-1004
Page 1772

Mapping Vulnerabilities for Services

Platform: DC/MDC

Component: RNA

RNA automatically maps vulnerabilities to a host for any service traffic received or sent by the host, when the service has a service ID in the RNA database and the packet header for the traffic includes a vendor and version.

However, many services do not include vendor and version information. For the services listed in the system policy, you can configure whether RNA associates vulnerabilities with service traffic for vendor and versionless services.

For example, a host receives SMTP traffic that does not have a vendor or version in the header. If you enable the SMTP service on the Vulnerability Mapping page of a system policy, then apply that policy to the Defense Center managing the sensor that detects the traffic, all vulnerabilities associated with SMTP applications are added to the host profile for the host.

Note that although RNA detectors collect service information and add it to host profiles, the service information will not be used for vulnerability mapping because you cannot specify a vendor or version for a custom service and cannot select the service for vulnerability mapping in the system policy.

To configure vulnerability mapping for services:

Access: Admin

1. Select **Operations > System Policy**.

   The System Policy page appears.

2. You have two options:

   - To modify active fingerprint source settings in an existing system policy, click **Edit** next to the system policy.
   - To configure active fingerprint source settings as part of a new system policy, click **Create Policy**.

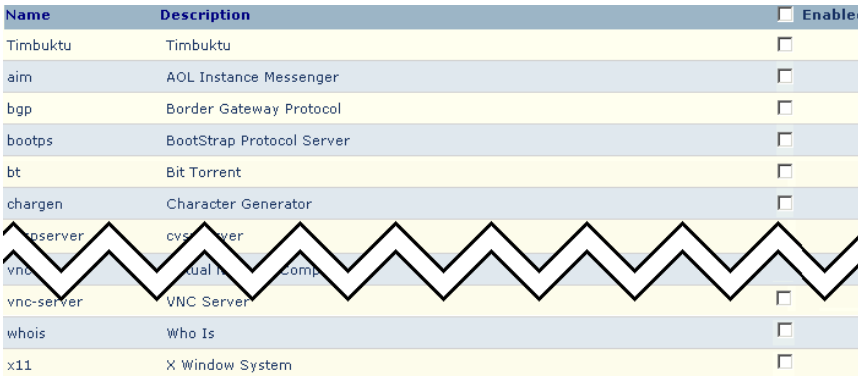
     Provide a name and description for the system policy as described in Creating a System Policy on page 1733, and click **Save**.

   In either case, the Access List page appears.

Keysight, EX-1004
Page 1773

3. Click **Vulnerability Mapping**.

   The Vulnerability Mapping page appears.



Name	Description	Enable
Timbuktu	Timbuktu	
aim	AOL Instance Messenger	
bgp	Border Gateway Protocol	
bootps	BootStrap Protocol Server	
bt	Bit Torrent	
chargen	Character Generator	
[illegible]	[illegible]	
[illegible]	[illegible]	
vnc-server	VNC Server	
whois	Who Is	
x11	X Window System	

4. You have two options:
   - To prevent vulnerabilities for a service from being mapped to hosts that receive service traffic without vendor or version information, clear the check box for that service.
   - To cause vulnerabilities for a service to be mapped to hosts that receive service traffic without vendor or version information, select the check box for that service.

TIP! You can select or clear all check boxes at once using the check box next to **Enable**.

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 sensors. See Applying a System Policy on page 1736 for more information.

Keysight, EX-1004
Page 1774

Chapter 41
Configuring System Settings

The system settings include a series of linked pages that you can use to view and modify settings on your appliance. Contrast the system settings, which are likely to be specific to a single appliance, with a system policy, which controls aspects of an appliance that are likely to be similar across a deployment. See Managing System Policies on page 1732 for more information.

The System Settings Options table describes the options you can configure in the system settings.

System Settings Options

Option	Description
Information	Allows you to view current information about the appliance. You can also change the appliance name. See Viewing and Modifying the Appliance Information on page 1777 for more information.
License	Provides you with options for managing your current licenses and for adding additional feature licenses on the platforms that support them. See Understanding Licenses on page 1780 for more information.
HTTPS Certificate	Allows you to request an HTTPS server certificate, if needed, from a trusted authority and upload certificates to your appliance. See Using Custom HTTPS Server Certificates on page 1794 for more information.

Keysight, EX-1004
Page 1775

System Settings Options (Continued)

Option	Description
Database	Lets you enable external read-only access to the appliance database, and provides a client driver for you to download. For more information, see Enabling Access to the Database on page 1799.
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. See Configuring Network Settings on page 1800 for more information.
Network Interface	Allows you to view and modify the settings for the network interfaces on your appliance. See Editing Network Interface Configurations on page 1803 for more information.
Process	Provides options that you can use to: • shut down the appliance • reboot the appliance • restart the Sourcefire 3D System-related processes See Shutting Down and Restarting the System on page 1806 for more information.
Remote Management	On the Defense Center, enables you to specify values for the internal network and management port that the Defense Center uses to communicate with its managed sensors and high availability peer. See Configuring the Communication Channel on page 1807 for more information. On the 3D Sensor, enables you to establish communications with a Defense Center from the sensor. See Configuring Remote Access to the Defense Center on page 1810 for more information.
Time	Displays the current time. If the time synchronization settings in the current system policy for the appliance is set to **Manual**, then you can use this page to change the time. See Setting the Time Manually on page 1813 for more information.
Health Blacklist	On the Defense Center, allows you to temporarily disable health monitoring for a 3D Sensor to prevent the Defense Center from generating unnecessary health events. See Blacklisting Health Modules on page 1815 for more information.

Keysight, EX-1004
Page 1776

System Settings Options (Continued)

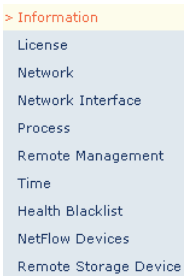
Option	Description
NetFlow Devices	On the Defense Center, allows you to specify the NetFlow-enabled devices you want to use to collect flow data. See Specifying NetFlow-Enabled Devices on page 1816 for more information.
Remote Storage Device	On Series 2 DC1000 and DC3000 Defense Centers, allows you to configure remote storage for backups and reports. See Managing Remote Storage on page 1817 for more information.
Lights-Out Management	On some Defense Centers and 3D Sensors, allows you to power cycle an appliance without being physically near the appliance. See Lights-Out Management on page 1823 for more information.

To configure the system settings:

Access: Admin

- Select **Operations > System Settings**.

  The Information page appears, with a list on the left side of the page that you can use to access other system settings. The Series 2 DC1000 or DC3000 Defense Center version of this the page is shown below.



Viewing and Modifying the Appliance Information

Platform: DC/3D Sensor

Component: None

The Information page provides you with information about the Defense Center or 3D Sensor. The information includes view-only information such as the product name and model number, the operating system and version, and the current

Keysight, EX-1004
Page 1777

appliance-level policies. The page also provides you with an option to change the name of the appliance.

IMPORTANT! You **cannot** view sensor information for Intrusion Agents.

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	Enable this check box to store event data on the Defense Center, but not the managed sensor. Clear this check box to store event data on both appliances.
Prohibit Packet Transfer to the Defense Center	Select this check box to prevent the managed sensor 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.
IP Address	The IP address of the appliance.
Current Policies	The appliance-level policies currently applied to the appliance. 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 can be important for troubleshooting.

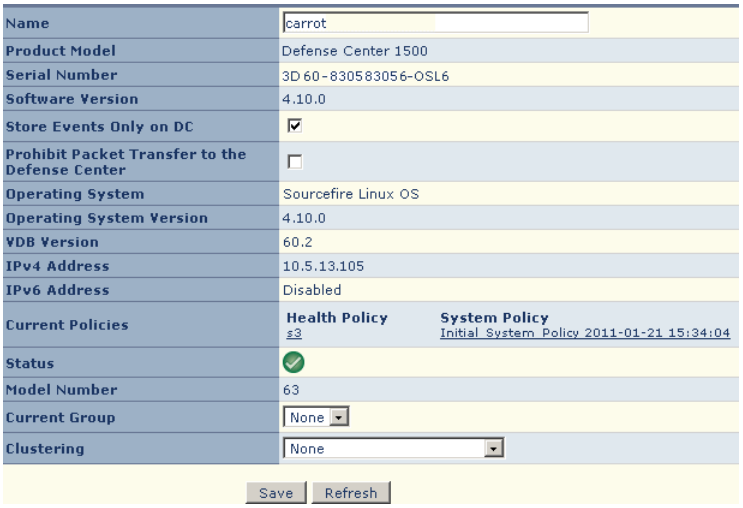
Keysight, EX-1004
Page 1778

To modify the appliance information:

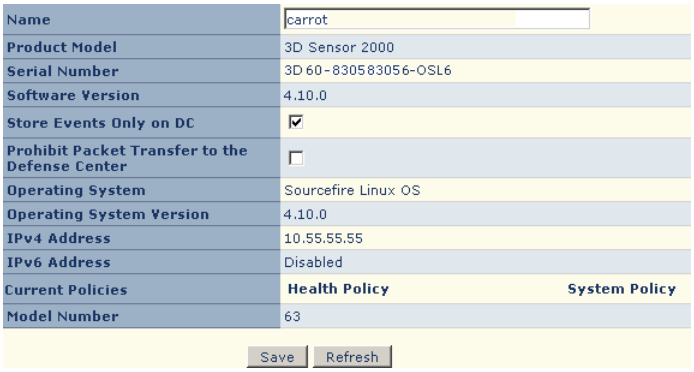
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears. The Defense Center version of the page is shown below.



   For comparison, the 3D Sensor version of the page is shown below.



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.

Keysight, EX-1004
Page 1779

Understanding Licenses

Platform: DC/MDC/3D Sensor

Component: None

You can license a variety of products and features to create your optimal deployment. 3D Sensors are licensed on the sensor at installation, with the exception of the 8000 Series sensors, which are licensed on the Defense Center. You can also add feature licenses such as RNA host licenses and Intrusion Agent licenses.

You can use a variety of appliances and optional features in your deployment. To understand why and when to use these licenses, see the following table.

Sourcefire Licenses

You apply a...	to...	so that you can...
Product License	a 3D Sensor other than 8000 Series sensors during installation	use IPS on that appliance. For information on adding a product license, see *Sourcefire 3D Sensor Installation Guide*, and *Sourcefire Defense Center Installation Guide*. Introduction to Sourcefire IPS on page 253.
Sensor License	a Defense Center at any time	manage one or more 8000 Series 3D Sensors. For information on managing a sensors, see Using the Defense Center on page 103.
Feature License	a Defense Center at any time	use additional features such as RNA, RUA, and so on. For information on how the various features function, see Understanding Feature Licenses on page 1784. For information on how to add a feature license, see Adding Feature Licenses on page 1788.
Virtual License	a Defense Center at any time	use virtual machines. For information on how to use virtual appliances, see *Sourcefire Virtual Defense Center and 3D Sensor Installation Guide*.

Keysight, EX-1004
Page 1780

You can view your licenses by using the Product Licensing widget in the dashboard. See Understanding the Product Licensing Widget on page 87 for more information. You can also use the License page to view the limits specified by the licenses for each feature or appliance you license. You can view your licenses from your Defense Center.

To view your current licenses:

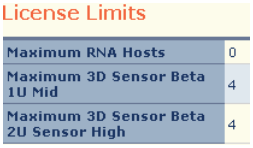
Access: Admin

1. On the Defense Center, select **Operations > System Settings**.

   The Information page appears.

2. Click **License**.

   The License page appears.



   See the following sections for more information:

   - Licensing an 8000 Series Sensor on page 1781
   - Understanding Feature Licenses on page 1784
   - Verifying Your Product License on page 1786
   - Managing Feature Licenses on page 1787

Licensing an 8000 Series Sensor

Platform: DC/8000 Series 3D Sensor

Component: None

A Defense Center manages the 8000 Series 3D Sensors, and sensors are licensed on the Defense Center that manages them. Without a valid license, the Defense Center does not process intrusion events.

The 8000 Series Sensor License Columns table describes each column that appears in a sensor license.

8000 Series Sensor License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.

Keysight, EX-1004
Page 1781

8000 Series Sensor License Columns (Continued)

Column	Description
Status	Indicates if the license is valid.
Type	Indicates the type of appliance associated with the license.
Max Connections	Lists the maximum number of connections allowed by the license.
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the feature license by clicking **Delete**.

After installing your 8000 Series sensor, use a Defense Center to license and register the sensor. If you need to obtain a sensor license, you can request it from the web interface. Before beginning, you should have the 12-digit sensor license serial number provided by Sourcefire when you purchased the sensor. If you do not have the serial number, you can find it by logging into the Sourcefire Support Site (https://support.sourcefire.com/), clicking **Account**, then clicking **Products & Contracts**. The serial number appears in the Sourcefire Software & Licenses section.

To register and license the 8000 Series sensor on the Defense Center:

Access: Admin

1. On the Defense Center, select **Operations > System Settings**.

   The Information page appears.

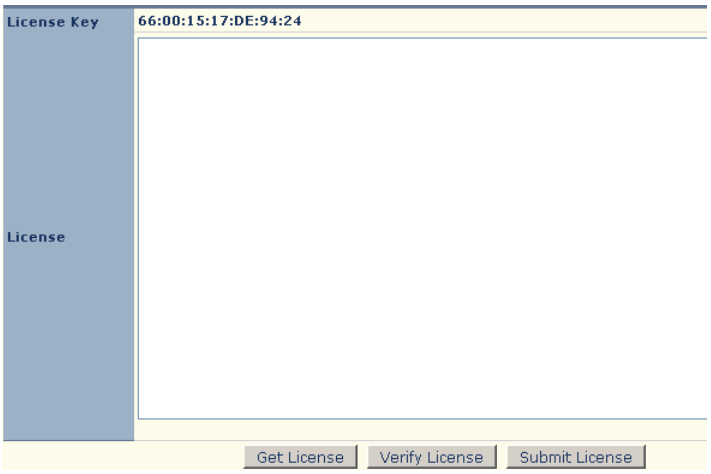
2. Click **License**.

   The License page appears.

Keysight, EX-1004
Page 1782

3. Click **Add New License**.

   The Add License page appears.



4. If you have received an email with your license file, copy the license file from the email, paste it into the **License** field, and click **Submit License**.
5. Optionally, if you do not have a license file, click **Get License**.

   The Licensing Center web site appears.

   IMPORTANT! If your web browser cannot access the Internet, you must switch to a host that can access it. Copy the license key at the bottom of the page and browse to https://keyserver.sourcefire.com/.

6. Follow the on-screen instructions for an appliance license to obtain your license file, which will be sent to you in an email.
7. Copy the license file from the email, paste it into the **License** field (see step 3), and click **Submit License**.

   If the license file is correct, the license is added to the Defense Center, and the sensor is available in the web interface.

Keysight, EX-1004
Page 1783

Understanding Feature Licenses

The Feature Licenses table describes how to determine which features to license for your deployment.

Feature Licenses

If you want to...	you need a license for...
capture and export data about the traffic that passes through NetFlow-enabled devices	NetFlow
monitor hosts on your network (including hosts discovered by NetFlow-enabled devices) to observe your network traffic to analyze a complete, up-to-the-minute profile of your network	RNA Hosts
correlate threat, endpoint, and network intelligence with user identity information	RUA Users
identify the source of policy breaches, attacks, or network vulnerabilities	RUA Users and either RNA Hosts or the product license (or both)
transmit events generated by open source Snort installations to the Defense Center	Intrusion Agents
IPS for use with Crossbeam Systems X-Series	IPS Software Sensors

For information on adding, viewing, and deleting feature licenses, see Managing Feature Licenses on page 1787. For detailed information on the feature licenses you can add, see the sections below.

NetFlow

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. The NetFlow cache stores a record of every flow (a sequence of packets that represents a connection between a source and destination host) that passes through the devices. You can deploy NetFlow-enabled devices on networks that your sensors cannot monitor, and use NetFlow data to monitor those networks.

You must use a Defense Center to configure NetFlow data collection and to view the collected data, and your deployment must include at least one 3D Sensor with RNA that can communicate with your NetFlow-enabled devices. Although

Keysight, EX-1004
Page 1784

you can use NetFlow-enabled devices exclusively to monitor your network, the Sourcefire 3D System uses RNA detection engines on 3D Sensors to analyze NetFlow data. For more information, see Introduction to NetFlow on page 1140.

RNA Host

Sourcefire RNA allows your organization to confidently monitor and protect your network using a combination of forensic analysis, behavioral profiling, and built-in alerting and remediation. 3D Sensors with RNA passively observe your organization's network traffic and analyze it to provide you with a complete, up-to-the-minute profile of your network.

By default, RNA is installed on most 3D Sensors. (The 3D9800 does not support RNA.) Sourcefire also makes key components of RNA available in installation packages for Crossbeam Systems security switches. However, to control how network intelligence is gathered and to view the resulting information, you must manage 3D Sensors with RNA with a Defense Center. In addition, to enable RNA functionality, that Defense Center must have an RNA host license installed and the 3D Sensor must have a product license installed. For more information, see Introduction to Sourcefire RNA on page 878.

RUA Host

Sourcefire Real-time User Awareness, also called RUA, allows your organization to correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, RUA can help you to identify the source of policy breaches, attacks, or network vulnerabilities, as well as 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.

All RUA deployments require a Defense Center that has an RUA feature license installed. If your organization uses LDAP, you can use the user information on your LDAP server to augment the Defense Center's database of user identity information with available metadata. For more information, see Using Sourcefire RUA on page 1452.

Intrusion Agent

If you have an existing installation of Snort®, you can install an Intrusion Agent to forward intrusion events to a Defense Center. You can then analyze the events detected by Snort alongside your other data.

Although you cannot manage policies or rules for an Intrusion Agent from the Defense Center, you can do analysis and reporting on those events. If the network map on the Defense Center has entries for the target host in a given event, the Defense Center assigns impact flags to the events. You can continue to manually tune Snort rules and preprocessors with the Intrusion Agent in place. For more information, see *Sourcefire 3D System Intrusion Agent Configuration Guide*.

Keysight, EX-1004
Page 1785

IPS Software Sensor

An IPS Software Sensor allows you to use 3D Sensor Software for X-Series on a Crossbeam® Next Generation Security Platform to gather network intelligence and intrusion information. For more information, see the *Sourcefire XOS Installation and Configuration Guide*.

Verifying Your Product License

Platform: DC/MDC/3D Sensor other than 8000 Series sensors

Component: None

During installation, the user who sets up the appliance adds the software license as part of the process. In most cases, you do not need to re-install the license.

To verify the product license file:

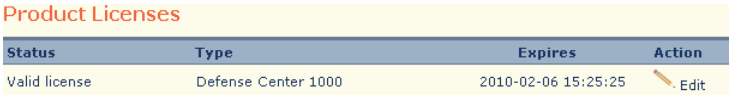
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

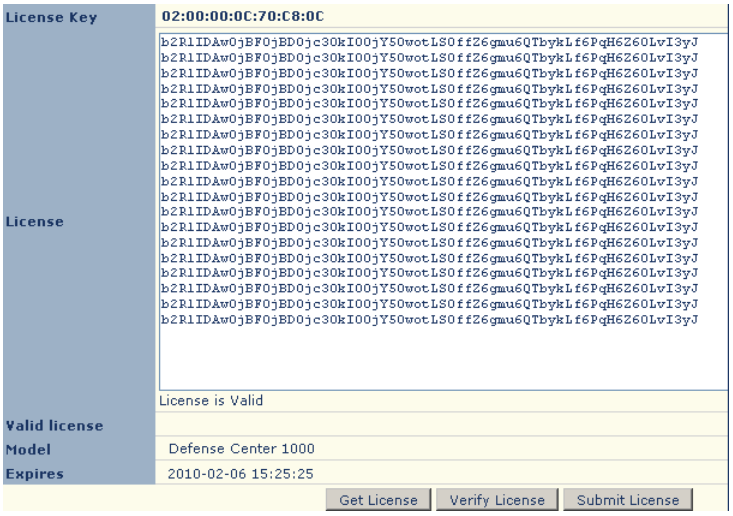
2. Click **License**.

   The License page appears.



3. Under Product Licenses, click **Edit**.

   The Manage License page appears.



Keysight, EX-1004
Page 1786

4. Click **Verify License**.
   - If the license file is **valid**, a message appears under the **License** field. Do **not** proceed to step 5.
   - If the license file is **invalid**, you will receive an error message. Continue with step 5 to obtain a license and install it.
5. Click **Get License**.

   The Licensing Center web site appears.

   IMPORTANT! If your web browser cannot access the Internet, you must switch to a host that can access it. Copy the license key at the bottom of the page and browse to https://keyserver.sourcefire.com/.
6. Follow the on-screen instructions for an appliance license to obtain your license file, which will be sent to you in an email.
7. Copy the license file from the email, paste it into the **License** field (as shown in Step 3), and click **Submit License**.

   If the license file is correct, the license is added to the appliance, and the features for the appliance are available in the web interface.

IMPORTANT! If you purchased a feature license, click **Add New License** and add it using the Add Feature License page. For more information about feature licenses, see Managing Feature Licenses on page 1787.

Managing Feature Licenses

Platform: DC/MDC/3D Sensor

Component: None

The Defense Center uses feature licenses to allow for additional features. Feature licenses include:

- NetFlow licenses, which specify the number of NetFlow-enabled devices you can use to gather flow data
- RNA host licenses, which specify the number of hosts that you can monitor with RNA
- RUA licenses, which allow you to use the RUA feature
- Intrusion Agent licenses, which allow you to use intrusion agents
- 3D Virtual Sensors, which allow you use virtual sensors in your deployment
- IPS licenses for Crossbeam, which allow you to use 3D Sensor Software with IPS on Crossbeam Systems security switches

Keysight, EX-1004
Page 1787

When you purchase license packs for any licensable feature, you must add them to the Defense Center from the web interface.

See the following sections for more information:

- Adding Feature Licenses on page 1788
- Viewing Feature Licenses on page 1790
- Configuring Network Settings on page 1800

Adding Feature Licenses

Platform: DC

Component: None

If you need to obtain a feature license for a feature you purchased, you can request it from the web interface. Before beginning, you should have the 12-digit feature license serial number provided by Sourcefire when you purchased the licensable feature. If you do not have the serial number, you can find it by logging into the Sourcefire Support Site (https://support.sourcefire.com/), clicking **Account**, then clicking **Products & Contracts**. The serial number appears in the Sourcefire Software & Licenses section.

IMPORTANT! Both Defense Centers in a high-availability pair **must** have NetFlow licenses for at least the number of NetFlow-enabled devices you are using. If one Defense Center does not have a NetFlow license, it will not receive data from your NetFlow-enabled devices.

To add a license:

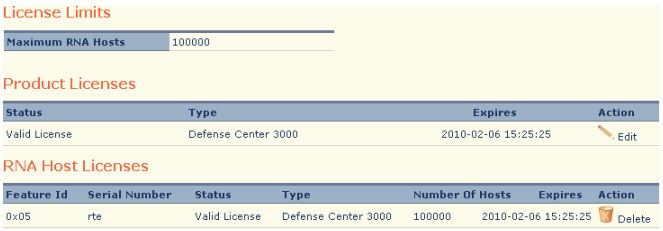
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **License**.

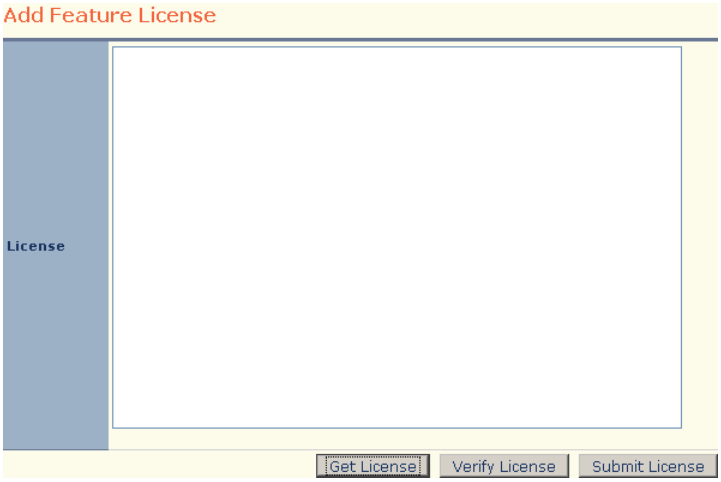
   The License page appears.



Keysight, EX-1004
Page 1788

3. Click **Add New License**.

   The Add Feature License page appears.



4. Click **Get License**.

   The Licensing Center web site appears.

IMPORTANT! If your web browser cannot access the Internet, you must switch to a host that can access it. Copy the license key at the bottom of the page and browse to https://keyserver.sourcefire.com/.

5. Follow the on-screen instructions for a feature license to obtain your license file, which will be sent to you in an email.

6. After you receive an email with the feature license file, copy the license file from the email, paste it into the **License** field, and click **Submit License**.

   If the license file is correct, the license is added to the appliance, and the licensed feature is available. You can repeat this process for each feature license you need to add.

TIP! Your Defense Center can have multiple feature licenses (for example, one or more licenses for RNA Hosts in addition to one or more licenses for Intrusion Agents, RUA, and so on). Note that there is only one product license.

Keysight, EX-1004
Page 1789

Viewing Feature Licenses

Platform: DC

Component: None

The licenses page displays the product and feature licenses that you have added to the Defense Center.

The first license that appears shows the Defense Center's product license which shows the license status, model code, node (MAC address), and expiration date, and provides a link that allows you to view or edit the license. For more information about viewing and modifying product licenses, see Verifying Your Product License on page 1786.

If you have feature or host licenses installed, they appear itemized below the product license. A summary of your licenses appears below the itemized list, and shows the total number of hosts, connections, exporters, virtual appliances, or users allowed by the sum of your feature or host licenses.

TIP! You can also view licenses by using the Product Licensing widget on the dashboard. See Understanding the Product Licensing Widget on page 87 for more information.

The NetFlow License Columns table describes each column that appears in a NetFlow license.

NetFlow License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.
Status	Indicates if the license is valid, invalid, or if a temporary license has expired.
Type	Indicates the type of appliance associated with the license
Allowed NetFlow Exporters	Lists the number of NetFlow-enabled devices that the license allows you to use.
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the feature license by clicking **Delete**.

Keysight, EX-1004
Page 1790

The RNA Host License Columns table describes each column that appears in an RNA host license.

RNA Host License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.
Status	Indicates if the license is valid, invalid, or if a temporary license has expired.
Type	Indicates the type of appliance associated with the license
Number of Hosts	Lists the number of monitored hosts added by the license.
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the host license by clicking **Delete**.

The RUA License Columns table describes each column that appears in an RUA host license.

RUA License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.
Status	Indicates if the license is valid, invalid, or if a temporary license has expired.
Type	Indicates the type of appliance associated with the license
Number of Users	Lists the number of monitored users added by the license.

Keysight, EX-1004
Page 1791

RUA License Columns (Continued)

Column	Description
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the feature license by clicking **Delete**.

The Intrusion Agent License Columns table describes each column that appears in an intrusion agent license.

Intrusion Agent License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.
Status	Indicates if the license is valid, invalid, or if a temporary license has expired.
Type	Indicates the type of appliance associated with the license
Swagent Max Connections	Lists the maximum number of software agent connections allowed by the license.
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the feature license by clicking **Delete**.

The Virtual 3D Sensor License Columns table describes each column that appears in an intrusion agent license.

Virtual 3D Sensor License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.

Keysight, EX-1004
Page 1792

Virtual 3D Sensor License Columns (Continued)

Column	Description
Status	Indicates if the license is valid, invalid, or if a temporary license has expired.
Type	Indicates the type of appliance associated with the license
Allowed Virtual Sensors	Lists the maximum number of Virtual 3D Sensors allowed by the license.
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the feature license by clicking **Delete**.

The IPS Software License Columns table describes each column that appears in an IPS Software license.

IPS Software License Columns

Column	Description
Feature ID	Displays the ID number that corresponds with the feature being licensed.
Serial Number	Displays the feature serial number.
Status	Indicates if the license is valid, invalid, or if a temporary license has expired.
Type	Indicates the type of appliance associated with the license
Expires	Displays the date and time that the feature license expires.
Action	Allows you to delete the feature license by clicking **Delete**.

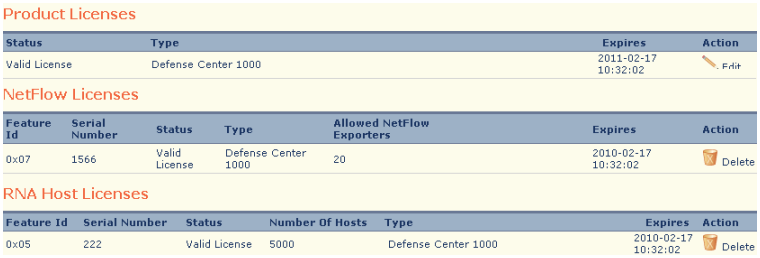
To view or delete your feature licenses:

Access: Admin

1. Select **Operations > System Settings**.
   The Information page appears.

Keysight, EX-1004
Page 1793

2. Click **License**.



The License page appears, showing the product license and any feature licenses you have added.

3. For the feature that you want to delete, click **Delete** in the Action column, and confirm that you want to delete the license.

Using Custom HTTPS Server Certificates

Platform: DC/MDC/3D Sensor

Component: None

Sourcefire Defense Centers and 3D Sensors 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 system settings for your appliance. For more information, see the following:

- Viewing the Current HTTPS Server Certificate on page 1795
- Generating a Certificate Request on page 1796
- Uploading Certificates on page 1797

Keysight, EX-1004
Page 1794

Viewing the Current HTTPS Server Certificate

Platform: DC/MDC/3D Sensor

Component: None

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

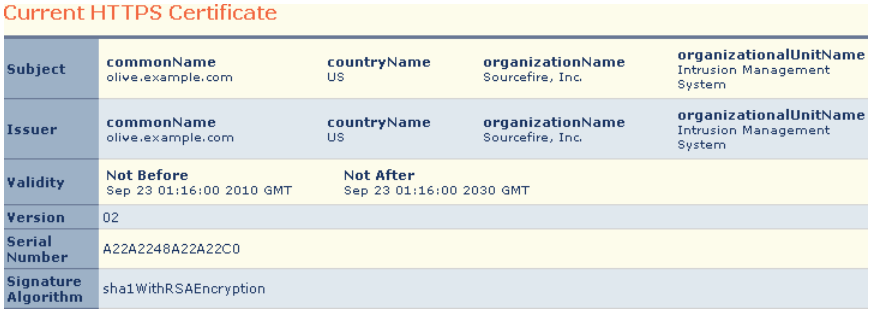
1. Select **Operations > System Settings**.

   The Information page appears.

Keysight, EX-1004
Page 1795

2. Click **HTTPS Certificate**.

   The Current HTTPS Certificate page appears, with the details of the current certificate for the appliance.



Current HTTPS Certificate

Subject	**commonName** olive.example.com	**countryName** US	**organizationName** Sourcefire, Inc.	**organizationalUnitName** Intrusion Management System
Issuer	**commonName** olive.example.com	**countryName** US	**organizationName** Sourcefire, Inc.	**organizationalUnitName** Intrusion Management System
Validity	**Not Before** Sep 23 01:16:00 2010 GMT	**Not After** Sep 23 01:16:00 2030 GMT		
Version	02			
Serial Number	A22A2248A22A22C0			
Signature Algorithm	sha1WithRSAEncryption			

Generating a Certificate Request

Platform: DC/MDC/3D Sensor

Component: None

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 system settings 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 **Operations > System Settings**.

   The Information page appears.

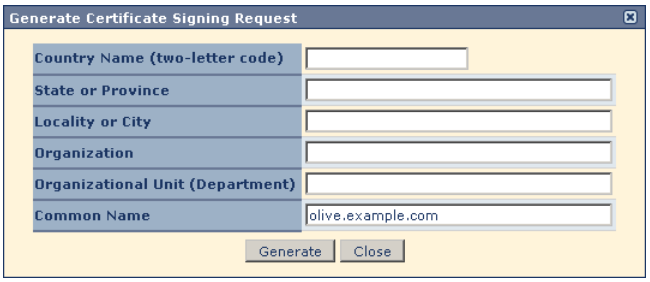
2. Click **HTTPS Certificate**.

   The Current HTTPS Certificate page appears.

Keysight, EX-1004
Page 1796

3. Click **Generate New CSR**.

   The Generate Certificate Signing Request pop-up window appears.



4. Type the two-letter country code for your country into the **Country Name** 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** 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 Certificates

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.

Keysight, EX-1004
Page 1797

To import a certificate:

Access: Admin

1. Select **Operations > System Settings**.

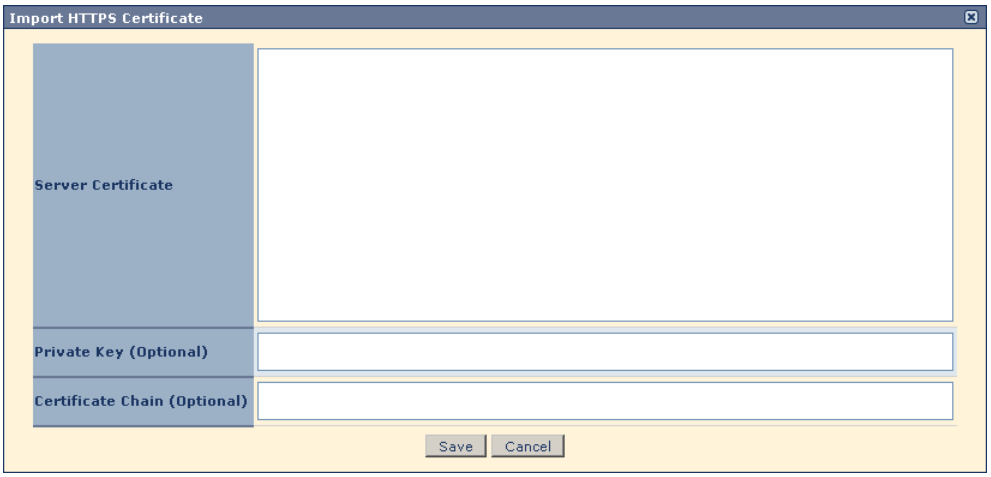
   The Information page appears.

2. Click **HTTPS Certificate**.

   The Current HTTPS Certificate page appears.

3. Click **Import HTTPS Certificate**.

   The Import HTTPS Certificate pop-up window appears.



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 past it into the **Certificate Chain** field.

7. Click **Import** to import the certificate.

   The certificate imports and the Current Certificate View updates to reflect the new certificate.

Keysight, EX-1004
Page 1798

Enabling Access to the Database

Platform: DC/MDC

Component: None

You can configure the Master Defense Center or 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 system settings page, you can enable database access as well as 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 1676.

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:

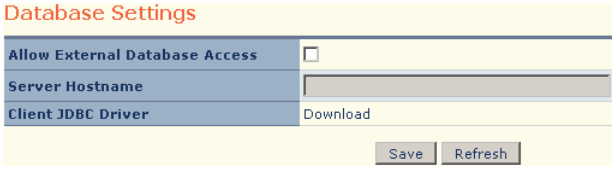
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **Database**.

   The Database Settings page appears.



3. Select the **Allow External Database Access** check box.

Keysight, EX-1004
Page 1799

4. If your Defense Center is configured for both IPv4 and IPv6, or if your database client and the Defense Center are separated by a network address translation (NAT) device, type the fully qualified domain name (FQDN) of the Defense Center in the **Server Hostname** field.

   You must make sure that the client can resolve the FQDN of the Defense Center, or you may not be able to connect to the database.

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.

6. Click **Save**.

   Your database access settings are saved.

TIP! Click **Refresh** to revert to the last saved database settings.

Configuring Network Settings

Platform: DC/MDC/3D Sensor

Component: None

With some exceptions, your Sourcefire 3D System provides a dual stack implementation so that you can choose IPv4, IPv6, or both IPv4 and IPv6 network settings in System Settings. The exceptions include software sensors or 3Dx800 sensors. You must use native applications, such as command line interfaces, third-party user interfaces, or the operating system interface, to manage network settings for software sensors or 3Dx800 sensors:

- For more information on configuring settings for Crossbeam-based software sensors, see the *Sourcefire 3D Sensor Software for X-Series Installation Guide*.
- For more information on configuring settings for Virtual 3D Sensors, see the *Virtual Defense Center and 3D Sensor Installation Guide*.
- For more information on configuring settings for 3Dx800 appliances, see the *3D Sensor Installation Guide*.
- For more information on configuring settings for Intrusion Agents, see the *Intrusion Agent Configuration Guide*.

You have the following configuration options:

- Disabled (IPv4 or IPv6)
- Manual (IPv4 and IPv6)
- DHCP (IPv4 and IPv6)
- Router assigned (IPv6 only)

Keysight, EX-1004
Page 1800

If you specify manual, you must manually configure all network properties. If you specify DHCP, the appliance automatically retrieves its network settings from a local DHCP server. If, in the case of IPv6, you specify Router assigned, the appliance retrieves its network settings from a local router.

Manual Network Configuration Settings

Setting	Description
Management Interface Address and either IPv4 Netmask or IPv6 Prefix Length	The IP address for the management interface. • 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). In most installations, the management interface is connected to an internal, protected network. This is the network through which Defense Centers and sensors communicate.
Default Network Gateway	The IP address of the gateway device for your network
Hostname	The DNS-resolvable name for the appliance **IMPORTANT!** If you change the hostname, the new name is not reflected in the syslog until after you reboot the appliance.
Domain	The fully qualified domain name where the appliance resides
Primary DNS Server	The IP address of the DNS server for the network where the appliance resides
Secondary DNS Server	A secondary DNS server's IP address
Tertiary DNS Server	A tertiary DNS server's IP address

If the appliance is not directly connected to the Internet, you can configure a proxy server to be used when downloading updates and SEUs. If needed, you can supply authentication credentials to connect to the proxy server via HTTP Digest authentication. By default, the appliance is configured to directly connect to the Internet.

Keysight, EX-1004
Page 1801

To configure network settings:

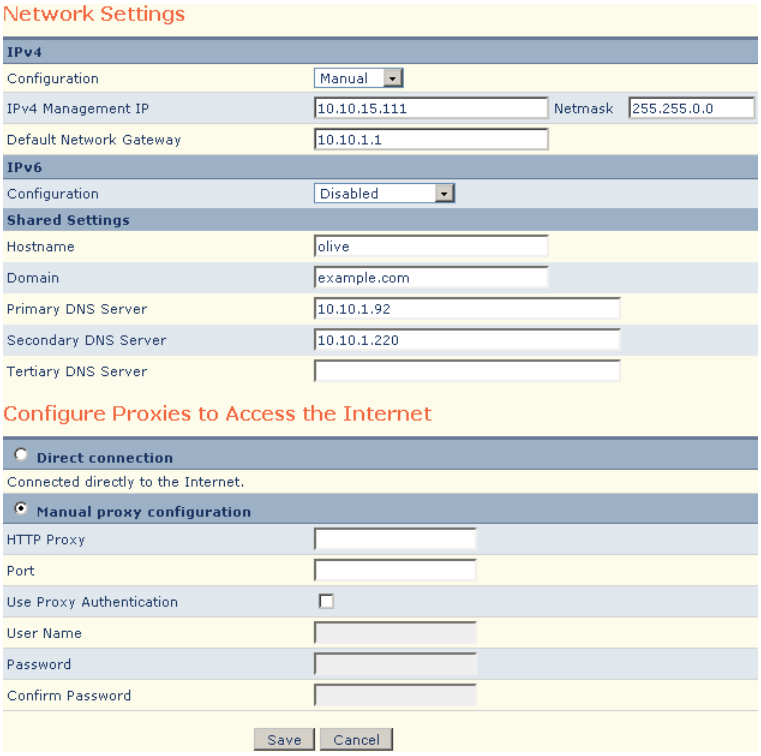
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **Network**.

   The Network page appears.



3. Specify which IP version (v4, v6, or both) you want to use by selecting the Configuration from the IPv4 and IPv6 settings:
   - 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.

Keysight, EX-1004
Page 1802

4. If you selected **Manual**, specify the network settings.

   See the Manual Network Configuration Settings table on page 1801 for a full description of each field you can configure. You can change the Shared Settings (hostname, domain, and domain servers) if you use manual or router assigned configurations.

5. If your appliance is not directly connected to the Internet, you can identify a proxy server to be used when downloading updates and rules. By default, appliances are configured to connect directly to the Internet. To configure a proxy server, you have two options:

   - If you have a direct connection from the appliance to the Internet, select **Direct connection**.
   - If your network uses a proxy, select **Manual proxy configuration** and 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** and supply a user name and password.

6. Click **Save**.

   The network settings are changed.

Editing Network Interface Configurations

Platform: DC/3D Sensor

Component: None

You can use the Network Interface page to modify the default settings for each network interface on your appliance. Any changes you make to the Auto Negotiate value are ignored for Gigabit interfaces. You must configure 3Dx800 interfaces on the 3Dx800 CLI.

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.

IMPORTANT! If you change the link mode for a sensing interface, the sensor drops traffic while the network interface card renegotiates its network connection.

Keysight, EX-1004
Page 1803

To edit a network interface:

Access: Admin

1. You have two choices:
   - To configure network interfaces from a 3D Sensor, select **Operations > System Settings**.
   - To configure network interfaces from a Defense Center, select **Operations > Sensor**, then click **Edit** next to the 3D Sensor.

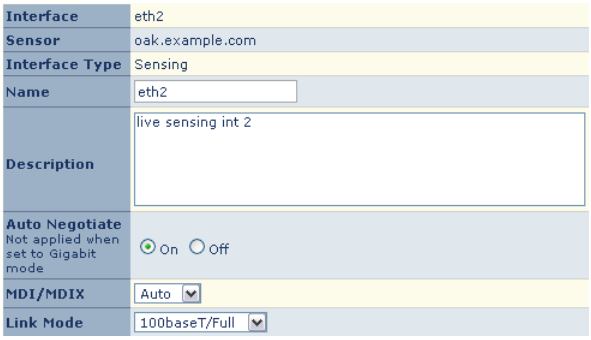
   The System Settings page appears.
2. Click **Network Interface**.

   The Network Interface page appears, listing the current settings for each interface on your appliance.

Name	Description	Interface	Type	MDI/MDIX	LinkMode	Sensor	
eth0	live management interface	eth0	Management	Auto	100baseT/Full	oak.example.com	Edit
eth1	live sensing int 1	eth1	Sensing	Auto	100baseT/Full	oak.example.com	Edit
eth2	live sensing int 2	eth2	Sensing	Auto	100baseT/Full	oak.example.com	Edit
eth3	unused int 3	eth3	N/A	Auto	N/A	oak.example.com	Edit
eth4	reserve int 4	eth4	N/A	MDIX	N/A	oak.example.com	Edit

3. Click **Edit** next to the interface that you want to modify.

   The current settings for the interface appear:



   These settings include:
   - interface name
   - sensor name
   - interface type, either Sensing or Management
   - interface description
   - whether the interface is configured to auto-negotiate speed and duplex settings

Keysight, EX-1004
Page 1804

- whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto mode (Series 2 3D Sensors only); N/A in this column indicates that the interface does not support MDI/MDIX
- 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 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. You must configure 3Dx800 interfaces on the 3Dx800 CLI.
- For Series 2 3D Sensors only, if you disable auto negotiation and specify a link mode, you must also set the MDI/MDIX field to the required MDI or MDIX mode.

  Normally, **MDI/MDIX** is set to **Auto**, which automatically handles switching between MDI and MDIX to attain link. However, when you set a specific link mode, automatic MDI/MDIX handling is disabled, making it impossible for the endpoints to attain link unless you manually set the required MDI/MDIX mode.

4. Click **Save**.

   The Network Interface page appears again.

Keysight, EX-1004
Page 1805

Shutting Down and Restarting the System

Platform: DC/MDC/3D Sensor

Component: None

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 RNA and Snort processes (Snort runs on the 3D Sensor only if you are licensed to use IPS)

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:

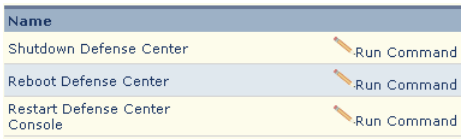
Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **Process**.

   The Appliance Process page appears. The Defense Center version of the page is shown below.



Keysight, EX-1004
Page 1806

3. Specify the command you want to perform:

   For the DC/MDC:

   - To shut down the Defense Center, click **Run Command** next to **Shutdown Defense Center**.
   - To reboot the system, click **Run Command** next to **Reboot Defense Center**. Note that this logs you out of the Defense Center.
   - To restart the Defense Center, 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.

   For the 3D Sensor:

   - To shut down the 3D Sensor, click **Run Command** next to **Shutdown Appliance**.
   - To reboot the system, click **Run Command** next to **Reboot Appliance**. Note that this logs you out of the 3D Sensor.
   - To restart the 3D Sensor, click **Run Command** next to **Restart Appliance Console**.
   - To restart the Snort and RNA processes, click **Run Command** next to **Restart Detection Engines**.

Configuring the Communication Channel

Platform: DC/3D Sensor

Component: None

Version 4.8 and earlier Defense Centers and sensors use a range of internal network IP addresses called the *management virtual network* to transmit third-party communications such as NTP to managed sensors and, in high availability deployments, to its Defense Center peer. The default address range is 172.16.0.0/16. The default port for communications between the Defense Center, its managed sensors, and if high availability is enabled, its high availability peer is 8305/tcp. The communication on port 8305 is bi-directional.

Enhancements in the current software eliminate the need for the management virtual network provided both the Defense Center and the sensors it manages are both using the current software. However, if your Defense Center is using the current version of the software and the sensors it manages are using an older version of the software, you will need to use a management virtual network and ensure that it does not conflict with other communications on your network.

IMPORTANT! The management virtual network is required only when the Defense Center must communicate with sensors running an older version. If both the Defense Center and all sensors have been upgraded to the current version, the management virtual network is unnecessary.

Keysight, EX-1004
Page 1807

For more information, refer to:

- Setting Up the Management Virtual Network on page 1808
- Editing the Management Virtual Network on page 1809

Setting Up the Management Virtual Network

Platform: DC/3D Sensor

Component: None

If the IP address range or the port conflicts with other communications on your network, you can specify different values. This is usually configured as part of the installation process, but you can change it later.

WARNING! The IP address range you specify for the Management Virtual Network must not conflict with any other local network, including your management network. The user interface prevents you from entering the address range for the management network, but make sure you **do not** to enter a range that overlaps other local networks. Doing so may break communications between hosts on the local network.

You must use native applications, such as command line interfaces, third-party user interfaces, or the operating system interface, to manage the communication channel sensor settings for Crossbeam-based software sensors, 3Dx800 sensors, and Intrusion Agents. For more information on configuring settings for Crossbeam-based software sensor, see the *Sourcefire 3D Sensor Software for X-Series Installation Guide*. For more information on configuring settings for Intrusion Agents, see the *Intrusion Agent Configuration Guide*.

IMPORTANT! Master Defense Centers do not currently use a Management Virtual Network. You can not edit the **Management Virtual Network** field of a Master Defense Center. The field is filled with 0.0.0.0/24 to indicate that the Management Virtual Network is disabled on a Master Defense Center.

To configure the communications channel:

Access: Admin

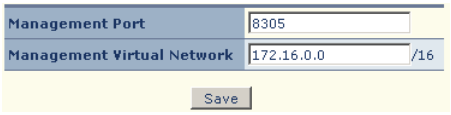
1. Select **Operations > System Settings**.

   The Information page appears.

Keysight, EX-1004
Page 1808

2. Click **Remote Management**.

   The Remote Management page appears.



3. In the **Management Port** field, enter the port number that you want to use.

   WARNING! Changing the management port on the Defense Center requires that you also manually change the management port on every managed sensor.

4. In the **Management Virtual Network** field, enter the IP address range that you want to use.

   TIP! The subnet mask is fixed at /16 (sixteen bits).

5. Click **Save** to save your changes for both the IP address range and the port number.

   The new values are saved.

Editing the Management Virtual Network

Platform: DC/3D Sensor

Component: None

You can change the host IP or host name of the connected appliance. You can also regenerate the Virtual IP address, a feature that is especially useful after network reconfigurations or appliance updates.

WARNING! If the Management Virtual Network is functioning properly, it should not be edited. Typically, this function is used only under the direction of Sourcefire Support.

Previous versions of Sourcefire 3D Systems used a default /24 CIDR address space, which provided enough addresses for 127 appliances. The current version uses a default /16 CIDR address space, which provides for a much greater number of appliances.

Master Defense Centers do not currently use a Management Virtual Network. You can not edit the **Management Virtual Network** field of a Master Defense Center. The field is filled with 0.0.0.0/24 to indicate that the Management Virtual Network is disabled on a Master Defense Center.

Keysight, EX-1004
Page 1809

To edit the remote management virtual network:

Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

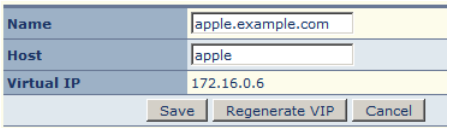
2. Click **Remote Management**.

   The Remote Management page appears.



3. Click **Edit** next to the host whose Management Virtual Network you want to change.

   The Edit Remote Management page appears.



4. Edit the name or host ID in the **Name** or **Host** fields as required.
5. Optionally, click **Regenerate VIP** to regenerate the IP address used by the virtual network.

   TIP! The regenerate VIP option is useful after you reconfigure your network or change the Sourcefire 3D System to take advantage of a larger address space.

6. After appropriate management virtual network edits are made, click **Save**.

Configuring Remote Access to the Defense Center

Platform: DC/3D Sensor

Component: None

You must begin the procedure for setting up the management relationship between a Defense Center and a sensor on the sensor.

Three fields are provided for setting up communications between appliances:

- **Management Host** - the hostname of IP address.
- **Registration Key** - registration key
- **Unique NAT ID** - a unique alphanumeric ID for use when registering sensors in NAT environments. See Working in NAT Environments on page 117 for more information.

Keysight, EX-1004
Page 1810

Valid combinations include:

- **Management Host** and **Registration Key** used on both appliances
- **Registration Key** and **Unique NAT ID** used on the 3D Sensor with **Host**, **Registration Key**, and **Unique NAT ID** used on the Defense Center.
- **Management Host**, **Registration Key**, and **Unique NAT ID** used on the 3D Sensor with **Registration Key** and **Unique NAT ID** used on the Defense Center.

Note that you must use the **Management Host** or **Host** field (hostname or IP address) on at least one of the appliances.

TIP! If you register a sensor to a Defense Center using a **Registration Key** and **Unique NAT ID**, but without a hostname or IP address, the Remote Management page displays the **Unique NAT ID** in the Host field.

Sourcefire **strongly** recommends that you read Using the Defense Center on page 103 before you add sensors to the Defense Center.

To set up sensor management from the sensor:

Access: Admin

1. On the sensor's web interface, select **Operations > System Settings**.

   The Information page appears.

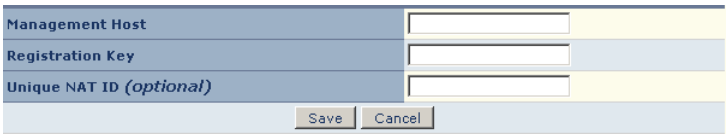
2. Click **Remote Management**.

   The Remote Management page appears.

   WARNING! Leave the **Management Port** field at the top of the Remote Management page in the default setting in nearly all cases. If you must change the **Management Port**, see Setting Up the Management Virtual Network on page 1808

3. Click **Add Manager**.

   The Add Remote Management page appears.



Keysight, EX-1004
Page 1811

4. In the **Management Host** field, type the IP address or the hostname of the Defense Center that you want to use to manage the sensor.

WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

TIP! You can leave the **Management 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.

5. In the **Registration Key** field, type the one-time use registration key that you want to use to set up a communications channel between the sensor and the Defense Center.
6. Optionally, in the **Unique NAT ID** field, type a unique alphanumeric NAT ID that you want to use to identify the sensor.
7. Click **Save**.

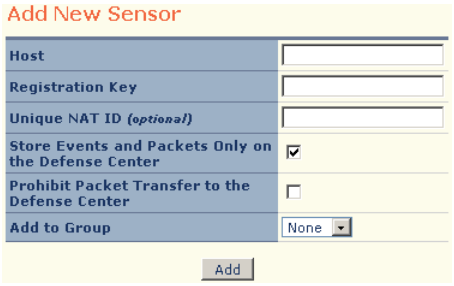
   After the sensor confirms communication with the Defense Center, the Pending Registration status appears.



8. Access the Defense Center web interface and select **Operations > Sensors**.

   The Sensors page appears.
9. Click **New Sensor**.

   The Add New Sensor page appears.



10. Type the IP address or the hostname of the sensor you want to add in the **Host** field.

WARNING! Make sure you use hostnames rather than IP addresses if your network uses DHCP to assign IP addresses.

Keysight, EX-1004
Page 1812

11. In the **Registration Key** field, type the same one-time use registration key that you used in step 5.
12. If you used a unique ID in step 6, type the same value in the **Unique NAT ID** field.
13. You can store IPS data on both the Defense Center and the sensor by clearing the **Store Events and Packets Only on the Defense Center** check box.

    By default, IPS data is stored only on the Defense Center and not on the sensor. Note that RNA data is never stored on the sensor.

    IMPORTANT! 3Dx800 sensors and Crossbeam-based software sensors cannot store IPS data locally. You **must** store events on the Defense Center.

14. You can prevent packet data from leaving a sensor by checking the **Prohibit Packet Transfer to the Defense Center** check box.

    IMPORTANT! If you elect to prohibit sending packets and you do not store events on the 3D Sensor, packet data is not retained. Packet data is often important for forensic analysis.

15. To add the sensor to a group, select the group from the **Add to Group** list.

    For more information about groups, see Managing Sensor Groups on page 137.
16. Click **Add**.

    The sensor is added to the Defense Center. It can take up to two minutes for the Defense Center to verify the sensor's heartbeat and establish communication.

    IMPORTANT! In some high availability deployments where network address translation is used, you may need to use the Add Manager feature to add the secondary Defense Center. Contact Sourcefire Support for more information.

Setting the Time Manually

Platform: DCMDC/3D Sensor

Component: None

If the Time Synchronization setting in the currently applied system policy is set to **Manual**, then you can manually set the time for the appliance using the Time page in the system settings.

Keysight, EX-1004
Page 1813

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: • **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 1768 for more information about the time settings in the system policy.

To manually configure the time:

Access: Admin

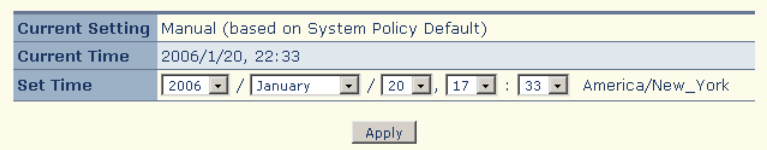
1. Select **Operations > System Settings**.

   The Information page appears.

Keysight, EX-1004
Page 1814

2. Click **Time**.

   The Time page appears.



3. From list boxes that appear, select the following:
   - 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. Select your time zone and click **Save** and, after the time zone setting is saved, click **Close** to close the pop-up window.

   For more information about using the time zone page, see Setting Your Default Time Zone on page 50.

Blacklisting Health Modules

Platform: DC/MDC

Component: None

If you want to disable health events 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. Once the blacklist settings take effect, the appliances report a disabled status in the Health Monitor Summary. For information on blacklisting individual or groups of appliances see Blacklisting Health Policies or Appliances on page 1936.

You can also 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. For example, if an appliance is temporarily disconnected from the management network, you can blacklist the Appliance Heartbeat module during that maintenance window. For information on blacklisting an individual policy modules, see Blacklisting a Health Policy Module on page 1938

Keysight, EX-1004
Page 1815

Specifying NetFlow-Enabled Devices

Platform: DC

Component: RNA + NetFlow

If you have enabled the NetFlow feature on your NetFlow-enabled devices), you can use the flow data that these devices collect to supplement the flow data collected by 3D Sensors with RNA by specifying the devices and the networks they monitor in your RNA detection policy.

One of the prerequisites for using NetFlow data is to use the system settings to specify the NetFlow-enabled devices you are going to use to collect the data. You must configure these NetFlow-enabled devices to export NetFlow version 5 data.

For more information on using NetFlow data with the Sourcefire 3D System, including information on additional prerequisites, see Introduction to NetFlow on page 1140.

To add NetFlow-enabled devices for flow data collection:

Access: Admin

1. Select **Operations > System Settings**.

   The Information page appears.

2. Click **NetFlow Devices**.

   The NetFlow Devices page appears.



3. Click **Add Device** to add a NetFlow-enabled device.



4. In the **IP Address** field, enter the IP address of the NetFlow-enabled device you want to use to collect flow data.

5. To add additional NetFlow-enabled devices, repeat steps 3 and 4.

   To remove a NetFlow-enabled device, click **Delete** next to the device you want to remove. Keep in mind that if you remove a NetFlow-enabled device from the system policy, you should also remove it from your RNA detection policy. For more information, see Editing an RNA Detection Policy on page 921.

6. Click **Save**.

   The list of NetFlow-enabled devices is saved.

Keysight, EX-1004
Page 1816

Managing Remote Storage

Platform: Series 2 DC

Component: None

On Defense Centers and Master 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 a 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 1985.

TIP! After configuring and selecting remote storage, you can switch back to local storage **only** if you **have not** increased the RNA flow database limit.

You must ensure that your external remote storage system is functional and accessible from the Defense Center or Master Defense Center.

Select one of the backup and report storage options:

- To disable external remote storage and use the local Defense Center or Master Defense Center for backup and report storage, see Using Local Storage on page 1817.
- To use NFS for backup and report storage, see Using NFS for Remote Storage on page 1818.
- To use SSH for backup and report storage, see Using SSH for Remote Storage on page 1820.
- To use SMB for backup and report storage, see Using SMB for Remote Storage on page 1821.

IMPORTANT! You cannot use remote backup and restore to manage data on Crossbeam-based software sensors.

Using Local Storage

Platform: Series 2 DC

Component: None

You can store backups and reports on the local Defense Center and Master Defense Center.

Keysight, EX-1004
Page 1817

To store backups and reports locally:

Access: Admin

1. Select **Operations** > **System Settings**.

   The Information page appears.

2. Click **Remote Storage Device**.

   The Remote Storage Device page appears.



3. At **Storage Type**, select **Local (No Remote Storage)**.
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

Platform: Series 2 DC

Component: None

You can select Network File System (NFS) protocol to store your reports and backups.

To store backups and reports using NFS:

Access: Admin

1. Select **Operations** > **System Settings**.

   The Information page appears.

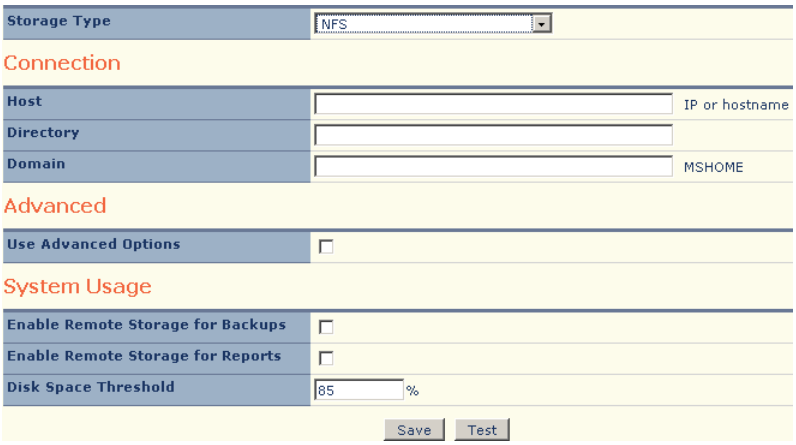
2. Click **Remote Storage Device**.

   The Remote Storage Device page appears.

Keysight, EX-1004
Page 1818

3. At **Storage Type**, select **NFS**.

   The page refreshes to display the NFS storage configuration options.



4. Add the connection information:
   - Enter the IP or hostname of the storage system in the **Host** field.
   - Enter the path to your storage area in the **Directory** 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 commands.
6. Under System Usage, select either or both of the following:
   - Select **Enable Remote Storage for Backups** to store backups on the designated host.
   - Select **Enable Remote Storage 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.

Keysight, EX-1004
Page 1819

Using SSH for Remote Storage

Platform: Series 2 DC

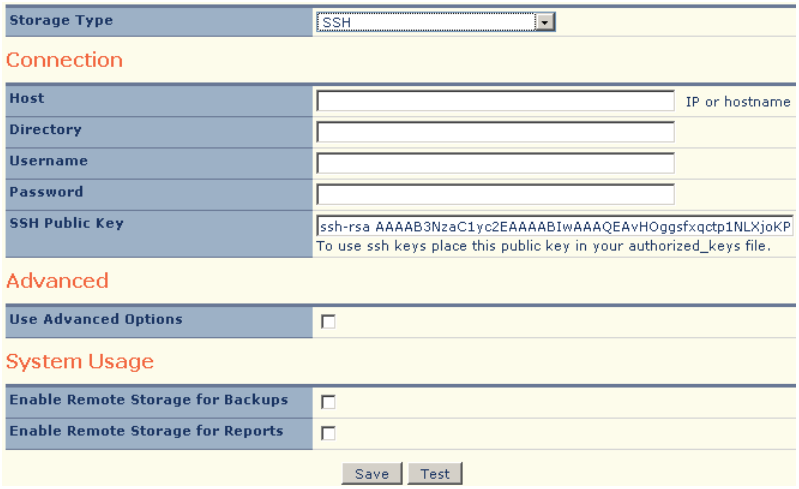
Component: None

You can select Secure Shell (SSH) protocol to store your reports and backups.

To store backups and reports using SSH:

Access: Admin

1. Select **Operations** > **System Settings**.
   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.



4. Add the connection information:
   - Enter the IP or hostname of the storage system in the **Host** field.
   - Enter the path to your storage area in the **Directory** 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 the commands.

Keysight, EX-1004
Page 1820

6. Under System Usage, select either or both of the following:
   - Select **Enable Remote Storage for Backups** to store backups on the designated host.
   - Select **Enable Remote Storage 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

Platform: Series 2 DC

Component: None

You can select Server Message Block (SMB) protocol to store your reports and backups.

To store backups and reports using SMB:

Access: Admin

1. Select **Operations** > **System Settings**.

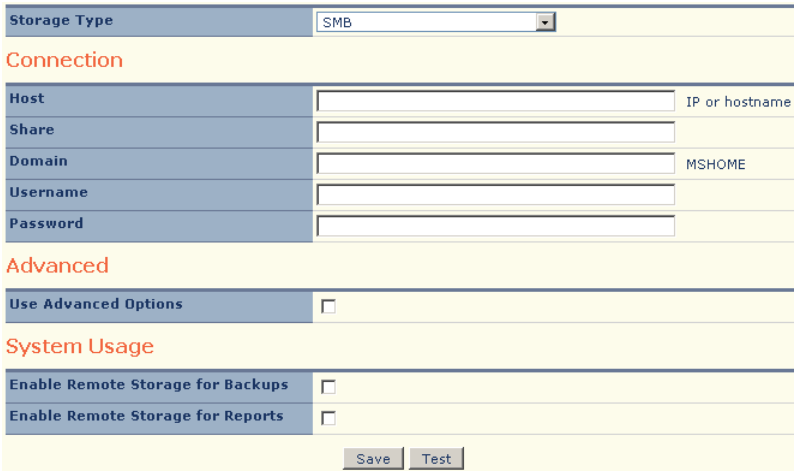
   The Information page appears.
2. Click **Remote Storage Device**.

   The Remote Storage Device page appears.

Keysight, EX-1004
Page 1821

3. At **Storage Type**, select **SMB**.

   The page refreshes to display the SMB storage configuration options.



4. Add the connection information:
   - Enter the IP or hostname of the storage system in the **Host** field.
   - Enter the share of your storage area in the **Share** field.
   - 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 commands.
6. Under System Usage, select either or both of the following:
   - Select **Enable Remote Storage for Backups** to store backups on the designated host.
   - Select **Enable Remote Storage 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.

Keysight, EX-1004
Page 1822

Lights-Out Management

Platform: 8000 Series Sensor

Component: None

You can use Lights-Out Management on a Serial Over LAN (SOL) connection to remotely monitor or manage your 8000 Series sensor without logging into the sensor. You can perform limited tasks, such as viewing the chassis serial number or monitoring such conditions as fan speed and temperature. Optionally, you can use the SOL connection to remotely log into the sensor at the command line interface.

If all attempts to access to your sensor have failed, you can use Lights-Out Management to restart your sensor remotely. Note that if a system is restarted while the SOL connection is active, the Lights-Out Management session may disconnect or time out.

WARNING! Do **not** restart your sensor unless it does not respond to any other attempts to restart. Remotely restarting the sensor does not gracefully reboot the system and you may lose data.

You must enable Lights-Out Management for both the sensor and the user you want to manage the sensor. After you enable the sensor and the user, you use a third-party Intelligent Platform Management Interface (IPMI) utility to access and manage your sensor. Examples of IPMI utilities are IPMItool for Linux-like or Macintosh environments or IPMIutil for Windows environments. For more information on using the utility for Lights-Out Management, see:

- Understanding a Serial Over LAN Connection on page 1823
- Understanding Lights-Out Management on page 1824
- Using Serial Over LAN or Lights-Out Management on page 1830

For more information on enabling Lights-Out Management, see:

- Enabling Lights-Out Management on the 3D Sensor on page 1827
- Enabling Lights-Out Management User Access on page 1828

Understanding a Serial Over LAN Connection

You use a third-party IPMI utility on your computer to create a Serial Over LAN connection to the sensor. If your computer uses a Linux-like or Macintosh environment, use IPMItool; for Windows environments, use IPMIutil.

Linux

IPMItool is standard with many distributions and is ready to use.

Keysight, EX-1004
Page 1823

Macintosh

IPMItool is included in the operating system, but is not configured by default for connection security. Recompile IPMItool with the `-I` flag turned on to enable encrypted sessions.

Windows

You must compile IPMIutil on Windows. If you do not have access to a compiler, you can use IPMIutil itself to compile. The IPMIutil binary file is available in the ipmirastools package located here:

`http://ipmiutil.sourceforge.net/FILES/ipmirastools-40.zip`

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 -P password
command
```

where

`ipmitool` invokes the utility.

`-I lanplus` enables encryption for the session.

`-H IP_address` indicates the IP address of the sensor you want to access.

`-U user_name` is the name of an authorized user.

`-P password` is the password of an authorised user.

`-command` is the name of the command you want to give.

The same command for Windows looks like this:

```
ipmiutil -J 3 -N IP_address -U user_name -P password command
```

This command connects you to the command line on the sensor and you can log in as if you were physically present at the sensor.

Understanding Lights-Out Management

Lights-Out Management provides the ability to perform a limited set of actions over a SOL connection without the need to log into the sensor. You use the 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.

For example, to display a l ist of sensor information, the IPMItool command is:

```
ipmitool -I lanplus -H IP_address -U user_name -P password
sdr
```


Keysight, EX-1004
Page 1824

The same command with the IPMIutil utility is :

```
ipmiutil -J 3 -N IP_address -U user_name -P password sensor
```

See the Sample Sensor Information Display on page 1826 for an example of the output from this command.

The Lights-Out Management Commands table describes the supported commands for different utilities.

Lights-Out Management Commands

IPMItool	IPMIutil	Description
`-I lanplus`	`-J 3`	Enables encryption for the IPMI session
`-H`	`-N`	Indicates the IP address of the remote sensor
`-U`	`-U`	Indicates the username of an authorized LOM account
`-P`	`-P`	Indicates the password for the username
`sol activate`	`sol -a`	Starts the SOL session
`sol deactivate`	`sol -d`	Ends the SOL session
`chassis power cycle`	`-c`	Restarts the sensor
`chassis power on`	`-u`	Powers up the sensor
`chassis power off`	`-d`	Powers down the sensor
`sdr`	`sensor`	Displays sensor information, such as fan speeds and temperatures

Keysight, EX-1004
Page 1825

Sample Sensor Information Display

This is an example of the output from running the `sdr` or `sensor` command:

```
BB +1.1V IOH     | 1.09 Volts        | ok
BB +1.1V P1 Vccp | 1.07 Volts        | ok
BB +1.1V P2 Vccp | 1.08 Volts        | ok
BB +1.5V P1 DDR3 | 1.53 Volts        | ok
BB +1.5V P2 DDR3 | 1.53 Volts        | ok
BB +1.8V AUX     | 1.77 Volts        | ok
BB +3.3V         | 3.28 Volts        | ok
BB +3.3V STBY    | 3.35 Volts        | ok
BB +3.3V Vbat    | 3.00 Volts        | ok
BB +5.0V         | 4.99 Volts        | ok
BB +5.0V STBY    | 4.95 Volts        | ok
BB +12.0V        | 11.88 Volts       | ok
BB -12.0V        | -12.24 Volts      | ok
BB +1.35v P1 MEM | disabled          | ns
BB +1.35v P2 MEM | disabled          | ns
Baseboard Temp   | 25 degrees C      | ok
Front Panel Temp | 17 degrees C      | ok
IOH Therm Margin | -52 degrees C     | ok
Mem P1 Thrm Mrgn | -31 degrees C     | ok
Mem P2 Thrm Mrgn | -38 degrees C     | ok
Processor 1 Fan  | 5764 RPM          | ok
Processor 2 Fan  | 5456 RPM          | ok
Memory Fan 2     | 4830 RPM          | ok
Memory Fan 1     | 4690 RPM          | ok
System Fan 1     | 8925 RPM          | ok
System Fan 2     | 6688 RPM          | ok
P1 Therm Margin  | -58 degrees C     | ok
P2 Therm Margin  | -60 degrees C     | ok
P1 Therm Ctrl %  | 0 unspecified     | ok
P2 Therm Ctrl %  | 0 unspecified     | ok
Pwr Unit Status  | 0x00              | ok
IPMI Watchdog    | 0x00              | ok
Physical Scrty   | 0x00              | ok
FP NMI Diag Int  | 0x00              | ok
SMI Timeout      | 0x00              | ok
System Event Log | 0x00              | ok
System Event     | 0x00              | ok
Button           | 0x00              | ok
```


Keysight, EX-1004
Page 1826

```
P1 Status          | 0x80              | ok
P2 Status          | 0x80              | ok
P1 VRD Hot         | 0x00              | ok
P2 VRD Hot         | 0x00              | ok
CATERR             | 0x00              | ok
CPU Missing        | 0x00              | ok
IOH Therm Trip     | 0x00              | ok
NM Capabilities    | Not Readable      | ns
```

Enabling Lights-Out Management on the 3D Sensor

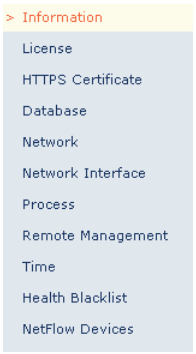
Use the web interface of the sensor you want to remotely manage to enable Lights-Out Management. The settings are prepopulated with the configuration (DHCP or static), IP address, netmask, and default gateway of the sensor.

To enable Lights-Out Management on an 8000 Series Sensor:

Access: Admin

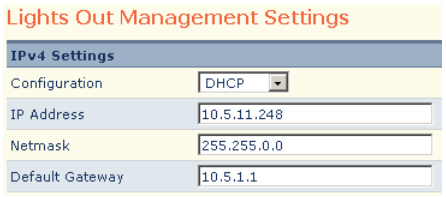
1. Select **Operations > System Settings** .

   The Information page appears, with a list on the left side of the page that you can use to access other system settings.



2. Select **Lights Out Management.**

   The Lights Out Management Settings page appears.



3. Confirm or edit the sensor's settings: (DHCP or static), sensor IP address, netmask, and default gateway, and click **Save**.

   The sensor is now enabled for Lights-Out Management.

Keysight, EX-1004
Page 1827

Enabling Lights-Out Management User Access

Use the web interface of the sensor you want to manage to enable an admin user with Lights-Out Management access. Note that Lights-Out Management access must be selected in addition to selecting the admin role. You must enable a Lights-Out Management user on each sensor you want to manage.

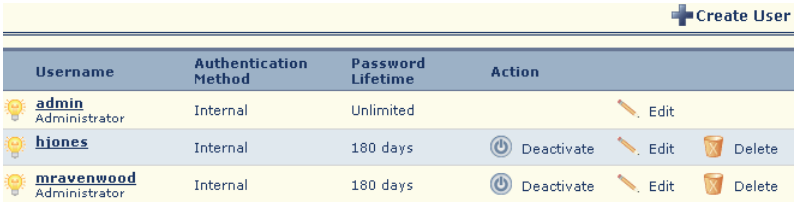
Each sensor may have up to 13 users with Lights-Out Management access. Each user name may contain up to 16 alphanumeric characters, and each password may contain up to 20 alphanumeric characters. Your Sourcefire 3D System user account is synchronized with your Lights-Out Management user account. Changing your system password changes your Lights-Out Management password.

To enable or view Lights-Out Management user access:

Access: Admin

1. Select **Operations > User Management.**

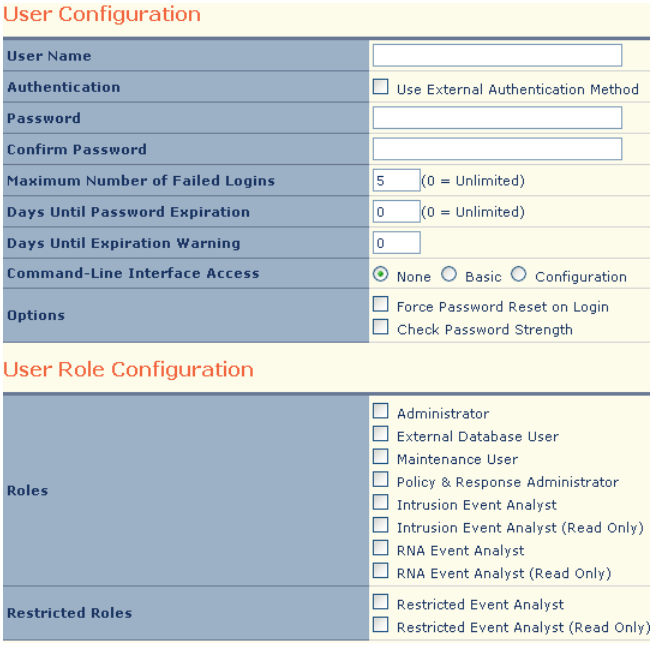
   The User Management page appears.



2. You have the following options:
   - To grant Lights-Out Management user access to an existing user, click **Edit** next to a user name in the list.

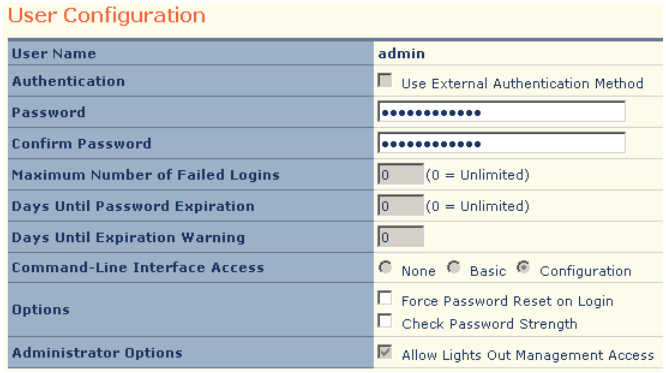
Keysight, EX-1004
Page 1828

- To grant Lights-Out Management user access to a new user, click **Create User**.



3. Enable the Administrator role.

   The User Configuration changes to display Lights-Out Management access in Administrator Options.



4. Check **Allow Lights Out Management Access** to allow the user access to Lights-Out Management.
5. Click **Save**.

   The user has Lights-Out Management access for this sensor.

Keysight, EX-1004
Page 1829

Using Serial Over LAN or Lights-Out Management

Select one of the commands from the Lights-Out Management Commands on page 1825 to use.

To create a Serial Over LAN connection:

Access: Admin with Lights-Out Management access

- Enter the following command:

  For IPMItool:

  ```
  ipmitool -I lanplus -H IP_address -U user_name -P password
  sol activate
  ```

  For IPMIutil:

  ```
  ipmiutil -J 3 -H IP_address -U username -P password sol -a
  ```

  The command line login for the sensor appears.

To use Lights-Out Management:

Access: Admin with Lights-Out Management access

- Enter the following command:

  For IPMItool:

  ```
  ipmitool -I lanplus -H IP_address -U user_name -P password
  command
  ```

  For IPMIutil:

  ```
  ipmiutil -J 3 -H IP_address -U username -P password command
  ```

  where *command* is one of the commands from the Lights-Out Management Commands table.

  The corresponding action as noted in the table is performed.

Keysight, EX-1004
Page 1830

Chapter 42
Updating System Software

Sourcefire electronically distributes several different types of updates, including major and minor updates to the system software itself, as well as Security Enhancement Updates (SEU) and 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 or the SEU, Sourcefire, you **must** read the release notes or advisory text that accompanies the update. The release notes provide specific information that you must know, including: supported platforms, compatibility, prerequisites, warnings, and specific installation and uninstallation instructions.

See the following sections for more information:

- Understanding Update Types on page 1832
- Installing Software Updates on page 1833
- Uninstalling Software Updates on page 1845
- Updating the Vulnerability Database on page 1846
- Importing SEUs and Rule Files on page 1848

Keysight, EX-1004
Page 1831

Understanding Update Types

Platform: DC/MDC/3D Sensor

Component: None

Sourcefire electronically distributes several different types of updates, including major and minor updates to the system software itself, as well as to the SEU and VDB.

The following table describes the types of updates provided by Sourcefire. Note that for most update types, you can schedule their download and installation; see Scheduling Tasks on page 1699. Also 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 or SEU. 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.

Sourcefire 3D System Update Types

Update Type	Description	Scheduling Supported?	Uninstall Supported?
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, 4.10.0.1).	yes	yes
feature updates to the Sourcefire 3D System	Feature updates more comprehensive than patches and generally include new features (and usually change the third digit in the version number; for example, 4.10.1)	yes	yes
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 version number; for example, 4.10 or 5.0)	no	no
VDB	VDB updates affect the vulnerabilities reported by RNA as well as the operating systems, client applications, and services that RNA detects.	yes	no
SEU	SEU updates provide new and updated standard text rules and shared object rules that you can use to detect potential attacks against your network and its assets. In addition, an SEU can also provide an updated version of Snort®, as well as features such as new preprocessors and decoders.	yes	no

Keysight, EX-1004
Page 1832

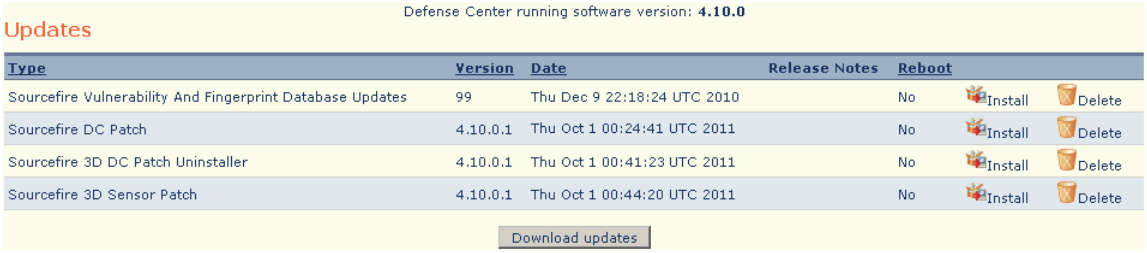
Installing Software Updates

Platform: DC/MDC/3D Sensor

Component: None

Sourcefire periodically issues software updates to the Sourcefire 3D System. Updating an appliance does not modify its configuration; the policies and network settings on the appliance remain intact.

To install Sourcefire 3D System updates, use the Patch Update Management page.The following graphic shows the Defense Center version of the page.



When you upload updates obtained from Sourcefire Support to your appliance, they appear on the page. Note that uploaded VDB updates (see Updating the Vulnerability Database on page 1846) also appear on the page, as do uninstaller updates (see Uninstalling Software Updates on page 1845), which are created when you install a patch to a Sourcefire appliance.

The list of updates shows the type of each update, the version number, and the date and time it was generated. It also indicates whether a reboot is required as part of the update.

If your deployment includes a Defense Center, you can use it to install Sourcefire 3D System updates on its managed 3D Sensors, including software sensors. However, for major updates to software sensors, you may need to uninstall the previous version and install the new version.

Note that when you install or uninstall updates on 3D Sensors, traffic inspection, link state, and traffic flow can be affected. The manner and duration of network traffic interruption depends on the following factors:

- whether the sensor's interface sets are configured as inline, inline with fail open, or passive
- the components of the Sourcefire 3D System that the update affects
- whether the update reboots the 3D Sensor

For specific information on how and when network traffic is affected for a particular update, see the release notes for that update.

This section explains how to plan for and perform manual software updates on your Sourcefire appliances. For patches and feature updates, you can take

Keysight, EX-1004
Page 1833

advantage of the automated update feature; see Automating Software Updates on page 1704.

To update your Sourcefire 3D System appliances:

Access: Admin

1. **Read the release notes for the update.**

   Available on the Sourcefire Support Site, the release notes describe supported platforms, new features and functionality, known and resolved issues, and product compatibility; they also contain information on any prerequisites, warnings, and specific installation and uninstallation instructions.

2. **Make sure your appliances (including software sensors) are running the correct version of the Sourcefire 3D System.**

   The release notes for the update indicate the required version. If you are running an earlier version, you can obtain updates from the Sourcefire Support Site.

3. **Install the latest SEU on your appliances.**

   You must install the latest SEU (see Importing SEUs and Rule Files on page 1848) on your appliances before you begin the update. You can obtain the SEU from the Sourcefire Support Site.

4. **Make sure the computers or appliances where you installed software sensors are running the correct versions of their operating systems.**

   Make sure that any Crossbeam Systems platforms you are using to host Sourcefire software sensors are running the correct version of the operating system, as described in the release notes.

5. **Delete any backups that reside on the appliance, then back up current event and configuration data to an external location.**

   Sourcefire **strongly** recommends that you delete or move any backup files that reside on your appliance, then back up current event and configuration data to an external location. Event data is **not** backed up as part of the update process.

   For more information on the backup and restore feature, including the types of backups that are supported for your appliance, see Using Backup and Restore on page 1985.

6. **Make sure you have enough free disk space and allow enough time for the update.**

   When you update a managed sensor, the update requires additional disk space on the Defense Center. The release notes for the update indicate space and time requirements.

7. **Update your Master Defense Centers.**

   Always update Master Defense Centers first; see Updating a Defense Center or Master Defense Center on page 1835.

Keysight, EX-1004
Page 1834

8. **Update your Defense Centers.**

   After you update any Master Defense Centers in your deployment, you can update your Defense Centers they manage; see Updating a Defense Center or Master Defense Center on page 1835.

   Note that 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 Center at the same time. First, complete the update procedure for one of the Defense Centers, then update the second Defense Center.

9. **Update your managed 3D Sensors.**

   After you update the Master Defense Centers and Defense Centers in your deployment, you can update your managed sensors (including software sensors). Sourcefire **strongly** recommends that you use your Defense Centers to update the sensors they manage; see Updating Managed Sensors on page 1838.

   Note that you **must** use the Defense Center to update sensors that do not have a web interface, including Crossbeam-based software sensors and Virtual 3D Sensors. However, for major updates to software sensors, you may need to uninstall the previous version and install the new version.

10. **Update your unmanaged 3D Sensors.**

    See Updating Unmanaged 3D Sensors on page 1840.

Updating a Defense Center or Master Defense Center

Platform: DC/MDC

Component: None

Use the procedure in this section to update your Defense Centers and Master Defense Centers. If your deployment includes Master Defense Centers, you **must** update them before you update the Defense Centers that they manage.

You 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 Support Site. Choose this option if your Defense Center has access to the internet and you are not performing a major update. 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.

Keysight, EX-1004
Page 1835

Note that 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 Center at the same time. First, complete the update procedure for one of the Defense Centers, then update the second Defense Center.

IMPORTANT! For major updates, updating the Defense Center removes existing uninstallers.

To update the Defense Center or Master Defense Center:

Access: Admin

1. Read the release notes for the update and complete any required pre-update tasks.

   Pre-update tasks can include making sure that the Defense Center is running the correct version of the Sourcefire software, making sure you have enough free disk space to perform the update, making sure you have set aside adequate time to perform the update, backing 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 **Operations > Update** to display the Patch Update Management page, then click **Download Updates** to check for the latest updates on the Support Site.
   - For major updates, or if your Defense Center does not have access to the Internet, first manually download the update from the Sourcefire Support Site. Select **Operations > Update** to display the Patch Update Management page, then click **Upload Update**. Browse to the update and click **Upload**.

   IMPORTANT! Download the update directly from the Support Site, either manually or by clicking **Update** on the Patch Update Management page. If you transfer an update file by email, it may become corrupted.

   The update is uploaded to the Defense Center. The Patch Update Management page shows the type of update you just uploaded, its version number, and the date and time it was generated. The page also indicates whether a reboot is required as part of the update.

Keysight, EX-1004
Page 1836

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 **Operations > 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 **Operations > Update**.

   The Patch Update Management page appears.
6. Click **Install** next to the update you uploaded.

   The Install Update page appears.
7. Under Selected Update, 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 1832 and the release notes to determine your update type.

   - For minor updates, you can monitor the update's progress in the task queue (**Operations > 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 1843 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. Even for minor updates, the web interface may become unavailable during the update process, or the Defense Center 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, continue to refrain from using the web interface until the update has completed. 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.
8. After the update finishes, if necessary, log into the Defense Center.

Keysight, EX-1004
Page 1837

9. Clear your browser cache and force a reload of the browser. Otherwise, the user interface may exhibit unexpected behavior.
10. Select **Operations > Help > About** and confirm that the software version is listed correctly.
11. Verify that all managed sensors are successfully communicating with the Defense Center.
12. Re-apply intrusion policies to the IPS detection engines on your managed 3D Sensors.

    Unless you enabled the **Inspect Traffic During Policy Apply** option when you created your IPS detection engines (this option is supported on many sensor models; see Creating a Detection Engine on page 206), applying an intrusion policy causes IPS detection engines to restart. This can cause a short pause in processing and, for most detection engines with inline interface sets, may cause a few packets to pass through the sensor uninspected.
13. Update the VDB on your Defense Centers and the 3D Sensors with RNA that they manage; see Updating the Vulnerability Database on page 1846.
14. Continue with the next section, Updating Managed Sensors, to update the Sourcefire software on the sensors that the Defense Center manages.

Updating Managed Sensors

Platform: DC

Component: None

After you update your Defense Centers, Sourcefire **strongly** recommends that you use them to update the sensors they manage. Updating managed sensors is a two-step process. First, download the update from the Support Site and upload it to the managing Defense Center. Next, install the software. Note that you can update multiple 3D Sensors at once, but only if they use the same update.

IMPORTANT! You **must** use the Defense Center to update sensors that do not have a web interface, including Crossbeam-based software sensors and Virtual 3D Sensors. However, for major updates to software sensors, you may need to uninstall the previous version and install the new version.

Keysight, EX-1004
Page 1838

To update managed 3D Sensors:

Access: Admin

1. Read the release notes for the update and complete any required pre-update tasks.

   Pre-update tasks can include updating your managing Defense Center, making sure that the 3D Sensors are running the correct version of the Sourcefire software, making sure software sensors are running the correct version of their operating systems, making sure you have enough free disk space to perform the update, you have set aside adequate time to perform the update, backing up event and configuration data, and so on.

2. Update the Sourcefire software on the sensors' managing Defense Center; see Updating a Defense Center or Master Defense Center on page 1835.

3. Download the update from the Sourcefire Support Site.

   Different 3D Sensor models use different updates. For information on the updates you can download, see the release notes.

   IMPORTANT! Download the update directly from the 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 **Operations > Update**.

   The Patch Update Management 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 Patch Update Management page shows the type of update you just uploaded, its version number, and date and time it was generated. The page also indicates whether a reboot is required as part of the update.

7. Click **Install** next to the update you are installing.

   The Install Update page appears.

Keysight, EX-1004
Page 1839

8. Select the sensors where you want install the update, then click **Install**. If prompted, confirm that you want to install the update and reboot the 3D Sensors.

   The update process begins.

   IMPORTANT! Traffic inspection, link state, and traffic flow can be affected during the update, depending on the types of interface sets configured on the sensors, the components that the update affects, and whether the update reboots the sensors. For specific information on how and when network traffic is affected for a particular update, see the release notes for that update.

   Depending on the size of the file, it may take some time to install the update on all sensors. You can monitor the update's progress in the Defense Center's task queue (**Operations > Monitoring > Task Status**).

   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. Select **Operations > Sensors** and confirm that the sensors you updated have the correct version listed.
10. Verify that the sensors you updated are successfully communicating with the Defense Center.
11. Reapply intrusion policies to the IPS detection engines on your managed 3D Sensors.

    Unless you enabled the **Inspect Traffic During Policy Apply** option when you created your IPS detection engines (this option is supported on many sensor models; see Creating a Detection Engine on page 206), applying an intrusion policy causes IPS detection engines to restart. This can cause a short pause in processing and, for most detection engines with inline interface sets, may cause a few packets to pass through the sensor uninspected.

Updating Unmanaged 3D Sensors

Platform: Standalone Sensor

Component: None

Use the procedure in this section to update unmanaged 3D Sensors only; Sourcefire **strongly** recommends that you update managed 3D Sensors using their managing Defense Centers. For more information, see Updating Managed Sensors on page 1838.

Keysight, EX-1004
Page 1840

You update the 3D Sensor in one of two ways, depending on the type of update and whether your 3D Sensor has access to the internet:

- You can use the 3D Sensor to obtain the update directly from the Support Site. Choose this option if your 3D Sensor has access to the internet and you are not performing a major update. 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 3D Sensor. Choose this option if your 3D Sensor does not have access to the internet or if you are performing a major update.

IMPORTANT! For major updates, updating the 3D Sensor removes existing uninstallers.

To update an unmanaged 3D Sensor:

Access: Admin

1. Read the release notes for the update and complete any required pre-update tasks.

   Pre-update tasks can include making sure that the 3D Sensor is running the correct version of the Sourcefire software, making sure you have enough free disk space to perform the update, making sure you have set aside adequate time to perform the update, backing up event and configuration data, and so on.

2. Upload the update to the 3D Sensor. You have two options, depending on the type of update and whether your 3D Sensor has access to the internet.
   - For all except major releases, and if your 3D Sensor has access to the Internet, select **Operations > Update** to display the Patch Update Management page, then click **Download Updates** to check for the latest updates on the Support Site.
   - For major releases, or if your 3D Sensor does not have access to the Internet, first manually download the update from the Sourcefire Support Site. Select **Operations > Update** to display the Patch Update Management page, then click **Upload Update**. Browse to the update and click **Upload**.

   IMPORTANT! Download the update directly from the Support Site, either manually or by clicking **Update** on the Patch Update Management page. If you transfer an update file by email, it may become corrupted.

   The update is uploaded to the 3D Sensor. The Patch Update Management page shows the type of update you just uploaded, its version number, and the date and time it was generated. The page also indicates whether a reboot is required as part of the update.

Keysight, EX-1004
Page 1841

3. Select **Operations > 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.

4. Select **Operations > Update**.

   The Patch Update Management page appears.

5. Click **Install** next to the update you just uploaded. If prompted, confirm that you want to install the update and reboot the 3D Sensor.

   The update process begins.

   IMPORTANT! Traffic inspection, link state, and traffic flow can be affected during the update, depending on the types of interface set configured on the sensors, the components that the update affects, and whether the update reboots the sensors. For specific information on how and when network traffic is affected for a particular update, see the release notes for that update.

   How you monitor the update depends on whether the update is a major update or a more minor update; see the Sourcefire 3D System Update Types table on page 1832 and the release notes to determine your update type.

   - For minor updates, you can monitor the update's progress in the task queue (**Operations > Monitoring > Task Status**).
   - For major updates, you can begin monitoring the update's progress in the task queue. However, after the 3D Sensor 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 1843 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 3D Sensor reboots. Even for minor updates, the web interface may become unavailable during the update process, or the 3D Sensor 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, continue to refrain from using the web interface until the update has completed. 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.

6. After the update finishes, if necessary, log into the 3D Sensor.

Keysight, EX-1004
Page 1842

7. Clear your browser cache and force a reload of the browser. Otherwise, the user interface may exhibit unexpected behavior.
8. Select **Operations > Help > About** and confirm that the software version is listed correctly.
9. Re-apply intrusion policies to your IPS detection engines.

   Unless you enabled the **Inspect Traffic During Policy Apply** option when you created your IPS detection engines (this option is supported on many sensor models; see Creating a Detection Engine on page 206), applying an intrusion policy causes IPS detection engines to restart. This can cause a short pause in processing and, for most detection engines with inline interface sets, may cause a few packets to pass through the sensor uninspected.

Monitoring the Status of Major Updates

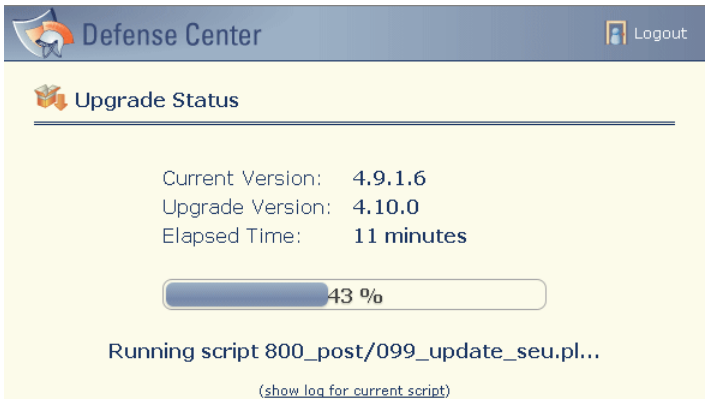
Platform: DC/MDC/3D Sensor

Component: None

For major updates, the Sourcefire 3D System provides you with a streamlined web interface so that you can easily monitor the update process, but also so that you are prohibited from using the web interface to perform tasks other than monitoring the update.

If you are upgrading a Master Defense Center, Defense Center, or unmanaged 3D Sensor, you can begin monitoring the update's progress in the task queue (**Operations > 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 and maintenance users, when you log back in, the Upgrade Status page appears. The following graphic shows the Defense Center version of the page.



Keysight, EX-1004
Page 1843

If you are using a Defense Center to update a managed 3D Sensor, Sourcefire recommends that you monitor the update's progress from the Defense Center's task queue. Note, however, that if you attempt to log in to the sensor's local web interface after the appliance finishes its pre-update checks, the Upgrade Status page appears and you can use it to monitor the update's progress just as you would for an unmanaged sensor.

The Upgrade Status 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 as well as gives details about the script currently being run.

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 Update Status 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, which can include: clearing your browser's cache and reloading the browser, verifying that the update completed successfully, verifying that the appliances in your deployment are communicating, reapplying intrusion policies, updating the VDB, and so on.

The procedures in Updating a Defense Center or Master Defense Center on page 1835 and Updating Unmanaged 3D Sensors on page 1840 outline the typically required post-update steps. However, for specific information on the post-update steps you must complete, see the release notes for that update.

Keysight, EX-1004
Page 1844

Uninstalling Software Updates

Platform: DC/MDC/3D Sensor

Component: None

When you install a patch to a Sourcefire appliance, the update process creates an uninstaller update that allows you to uninstall the patch from that appliances's web interface.

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.

You must use the local web interface to uninstall patches, as described by the procedure in this section; you cannot use the Defense Center to uninstall patches from managed sensors. For information on uninstalling patches from appliances that do not have local web interfaces, including Crossbeam-based software sensors and Virtual 3D Sensors, see the release notes.

In addition, you must uninstall a patch from the appliances in your deployment in the reverse order of how you installed it. That is, first uninstall the patch from your managed 3D Sensors, then your Defense Centers, and finally your Master Defense Centers.

When you uninstall a patch, 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 4.10.0 to Version 4.10.0.2. Uninstalling the Version 4.10.0.2 patch might result in an appliance running Version 4.10.0.1, even though you never installed the Version 4.10.0.1 update. For information on the resulting Sourcefire software version when you uninstall an update, see the release notes.

To uninstall a patch using the local web interface:

Access: Admin

1. Select **Operations > Update**.

   The Patch Update Management page appears.

Keysight, EX-1004
Page 1845

2. Click **Install** next to the uninstaller for the update you want to remove.
   - On the Defense Center, the Install Update page appears. Under Selected Update, select the Defense Center and click **Install**.
   - On the 3D Sensor, 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 (**Operations > Monitoring > Task Status**).

   If the uninstall for a 3D Sensor requires a reboot, the sensor uses IPS detection engines with inline interface sets, and the sensor does not have a fail-open network card, traffic is interrupted while the sensor reboots. If the sensor has a fail-open network card, some traffic may pass through the sensor uninspected while it reboots.

   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. Before the uninstall completes, the web interface may become unavailable, or the appliance may log you out. This is expected behavior. If this occurs, log in again and view the task queue. If the uninstall is still running, continue to refrain from using the web interface until the uninstall has completed. If you encounter issues with the uninstall, for example, if the task queue indicates that the uninstall has failed or if a manual refresh of the task queue shows no progress, do **not** restart the uninstall. Instead, contact Support.

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 **Operations > 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 sensors (for the Defense Center) or its managing Defense Center (for 3D Sensors).

Updating the Vulnerability Database

Platform: DC

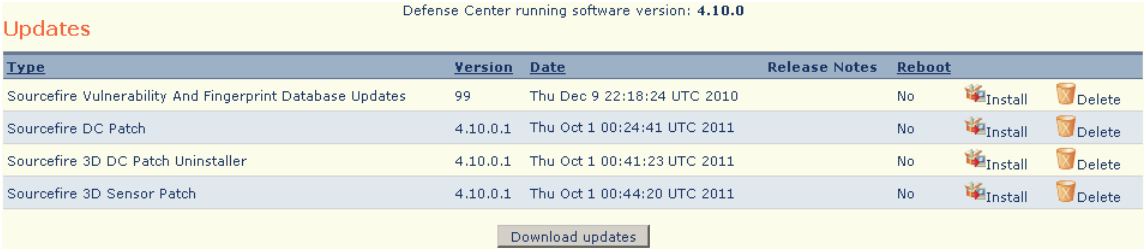
Component: RNA

The Sourcefire Vulnerability Database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for RNA-detection operating systems, client applications, and services. RNA correlates the operating system and services detected on each host with the vulnerability database to

Keysight, EX-1004
Page 1846

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 install VDB updates, use the Patch Update Management page.



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.

Because you should install the same version of the VDB on all the appliances in your deployment, you should use your Defense Centers to install the VDB on all managed 3D Sensors with RNA, including software sensors. Because you cannot view RNA data on Master Defense Centers or on unmanaged 3D Sensors, you do not need to update the VDB on these appliances.

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.

TIP! This section explains how to plan for and perform manual VDB updates on your Sourcefire 3D System appliances. You can take advantage of the automated update feature to schedule VDB updates; see Automating Vulnerability Database Updates on page 1710.

To update the vulnerability database:

Access: Admin

1. Read the VDB Update Advisory Text for the update.

   The VDB Update Advisory Text includes information about the changes to the VDB made in the update, as well as product compatibility information.

2. Select **Operations > Update**.

   The Patch Update Management page appears.

Keysight, EX-1004
Page 1847

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 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 Support Site, either manually or by clicking **Update**. If you transfer an update file by email, it may become corrupted.

   The VDB update is saved on the Defense Center and appears in the Updates section.
4. Click **Install** next to the VDB update.

   The Install Update page appears.
5. Select the Defense Center, as well as the sensors where you want to install the VDB update, 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 (**Operations > 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, confirm that the VDB build number matches the update you installed.
   - To check the VDB build number on the Defense Center, select **Operations > Help > About**.
   - To check the VDB build number on your managed sensors, select **Operations > Sensors** on the Defense Center, then click **Edit** next to each sensor you updated.

Importing SEUs and Rule Files

Platform: DC/MDC/Standalone Sensor

Component: IPS

As new vulnerabilities become known, the Sourcefire Vulnerability Research Team (VRT) releases Security Enhancement Updates (SEUs).

Keysight, EX-1004
Page 1848

An SEU contains new and updated standard text rules and shared object rules that you can use to detect potential attacks against your network and its assets. An SEU can also provide IPS with an updated version of Snort®, as well as features such as new preprocessors and decoders. Importing an SEU installs the changes in the SEU on the appliance where you import it. When you apply a policy that implements new or modified SEU components supported by the version of system software running on a managed 3D Sensor, and the SEU has not been installed on the sensor, the system pushes the necessary SEU components to the sensor.

IMPORTANT! SEUs may contain new binaries. Make sure your process for downloading and installing SEUs complies with your security policies. In addition, SEUs can be quite large, so make sure you import SEUs during periods of low network use.

The following are important points you should keep in mind when you import an SEU:

- VRT often includes new rules in SEUs, with the rule state set for 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 362 for more information.

  Sometimes VRT also uses an SEU to change the default state of existing rules. For information on choosing whether to allow SEUs to change the default states of existing rules in intrusion policies you create, see Allowing SEUs to Modify the Base Policy on page 363.
- SEUs are cumulative, so the newest SEU contains the intrusion rules and new features of all previous SEUs. You cannot import an SEU with a version number lower than the version of the currently installed SEU.
- Optionally, an SEU can add new and updated features to the base policy in intrusion policies you create, including new default feature settings, new rules set to their default rule states, and modified default rule states of existing rules. An SEU adds new rules and other features, and updates default feature settings and rule states, according to whether you choose for your base policy to be updated based on updates in the SEU. See Understanding the Base Policy on page 361 for more information.
- Importing an SEU discards any unsaved edits to an intrusion policy. See Committing Intrusion Policy Changes on page 345 for more information.
- If your Sourcefire 3D System deployment includes two Defense Centers configured as a high availability pair, you only need to import the SEU on one of the Defense Centers. The second Defense Center receives the SEU as part of the regular synchronization process.

Keysight, EX-1004
Page 1849

- Optionally, when the import completes you can automatically reapply intrusion policies owned by the appliance where you import the SEU. Note that an appliance owns policies applied from that appliance. For example, a Defense Center owns a policy that it applies to a detection engine on a managed 3D Sensor, and a standalone 3D Sensor owns a policy that it applies to its own detection engine.
- Applying an intrusion policy from a Defense Center to a managed sensor after you import an SEU does not apply the SEU to the sensor. However, any new rules or features provided by the SEU that are enabled in the policy you apply to the sensor are also enabled on the sensor by that policy.
- You can use the **Reapply All** or **Reapply Intrusion Policies** option on the Detection Engines page (**Operations > Configuration > Detection Engines**) to reapply your intrusion policies. These options can be convenient, for example, if you import an SEU and do not automatically reapply your intrusion policies after the import completes. These options conveniently reapply filtered and non-filtered intrusion policies simultaneously for each detection engine. See Understanding Detection Engines on page 199 and Defining VLANs and Subnetworks on page 393 for more information.

In addition to configuring SEU imports on the Import SEU page, you can also schedule SEU imports on the Scheduling page. The Comparing SEU Import Features table compares the options available with the two SEU import features.

Comparing SEU Import Features

Option	Import SEU Page	Scheduling Page
Download and install a single SEU	Yes	Yes
Schedule a single SEU import	No	Yes
Schedule recurring daily, week, or monthly imports	Yes	Yes
Automatically re-apply your intrusion policies when the import completes	Yes	Yes
Configure separate SEU download and install tasks	No	Yes
View scheduled imports on the scheduling calendar	No	Yes
View import details in the Task Details table	No	Yes

Keysight, EX-1004
Page 1850

There are two methods that you can use to import SEUs on the Import SEU page.

- Using One-Time SEU Imports on page 1851 explains how to import a single SEU from the Sourcefire Support web site.
- Using Recurring SEU Imports on page 1855 explains how to use an automated feature on the web interface to download and install SEUs from the Sourcefire Support site and, optionally, set the rules state for new rules and re-apply your intrusion policy.

You can also import a copy of a standard text rules file that you have created on a local machine. See Importing Local Rule Files on page 1857 for more information.

Using One-Time SEU Imports

Platform: DC/MDC/Standalone Sensor

Component: IPS

There are two methods that you can use for one-time SEU imports.

- Using Manual One-Time SEU Imports on page 1851 explains how to manually download an SEU from the Sourcefire Support web site to your local machine and then manually install the SEU.
- Using Automatic One-Time SEU Imports on page 1853 explains how to use an automated feature on the web interface to search the Sourcefire Support site for new SEUs and upload them.

Using Manual One-Time SEU Imports

Platform: DC/MDC/Standalone Sensor

Component: IPS

The following procedure explains how to import a new SEU manually. This procedure is especially useful if your Defense Center or 3D Sensor does not have Internet access.

To manually import an SEU:

Access: P&R Admin/Admin

1. From a computer that can access the Internet, access and log into Sourcefire Support (https://support.sourcefire.com/).
2. Click **Downloads**.
3. Navigate to the latest SEU.

TIP! SEUs are cumulative, so the newest SEU contains the intrusion rules and new features of all previous SEUs. You cannot import an SEU with a version number lower than the version of the currently installed SEU.

Keysight, EX-1004
Page 1851

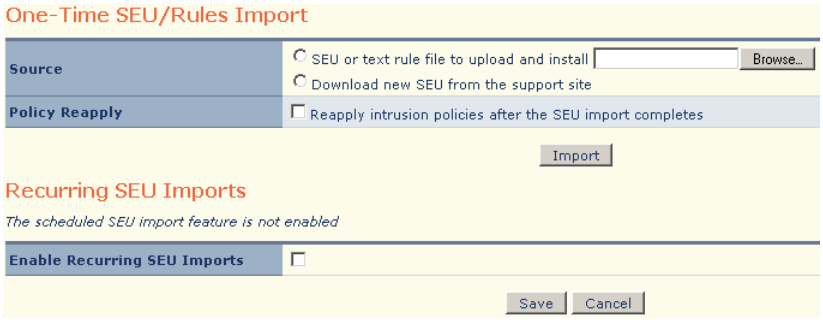
4. Click the SEU file that you want to download and save it to your computer.
5. Log into your appliance's web interface.
6. Select **Policy & Response > IPS > SEU**.

   The SEU page appears.



7. Click **Import SEU**.

   The Import SEU page appears.



TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policy & Response > IPS > Rule Editor**.

8. 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 868 for more information.
9. Select **SEU or text rule file to upload and install** and click **Browse** to navigate to and select the SEU file.
10. Optionally, select **Reapply intrusion policies after the SEU import completes** to automatically reapply intrusion policies currently applied from this appliance when the SEU import completes.

    Note that you cannot apply intrusion policies to a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). See Using Clustered 3D Sensors on page 250 for more information.

Keysight, EX-1004
Page 1852

11. Click **Import**.

    The SEU is installed and the rules are updated. The system displays the SEU Detail View workflow. See Understanding the SEU Detail View Table on page 1863 for more information.

    Unless you selected **Reapply intrusion policies after the SEU import completes** in step 10, any rule changes are not implemented until the next time you apply the affected intrusion policies. See Applying an Intrusion Policy on page 347 for procedures.

IMPORTANT! Contact Support if you receive an error message while installing the SEU.

Using Automatic One-Time SEU Imports

Platform: DC/MDC/Standalone Sensor

Component: IPS

The following procedure explains how to import a new SEU by automatically connecting to the Sourcefire Support site. You can use this procedure only if the appliance has Internet access.

To automatically import an SEU:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > SEU**.

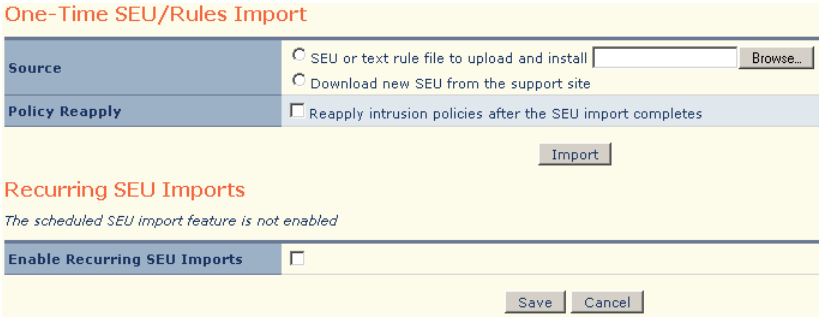
   The SEU page appears.



Keysight, EX-1004
Page 1853

2. Click **Import SEU**.

   The Import SEU page appears.



TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policy & Response > IPS > Rule Editor**.

3. 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 868 for more information.
4. Select **Download new SEU from the support site**.
5. Optionally, select **Reapply intrusion policies after the SEU import completes** to automatically reapply intrusion policies currently applied from this appliance when the SEU import completes.

   Note that you cannot apply intrusion policies to a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). See Using Clustered 3D Sensors on page 250 for more information.
6. Click **Import**

   The SEU is installed and the rules are updated. The system displays the SEU Detail View workflow. See SEU Detail View Fields on page 1863 for more information.

   Unless you selected **Reapply intrusion policies after the SEU import completes** in step 5, any rule changes are not implemented until the next time you apply the affected intrusion policies. See Applying an Intrusion Policy on page 347 for procedures.

IMPORTANT! Contact technical support if you receive an error message while installing the SEU.

Keysight, EX-1004
Page 1854

Using Recurring SEU Imports

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can configure daily, weekly, or monthly SEU imports on the Import SEU page. This feature is conveniently located on the Import SEU page to be near SEU import log information. You can also schedule SEU imports from the Scheduling page. For more information on scheduling SEU imports on the Scheduling page, see Automating SEU Imports on page 1716.

If your Sourcefire 3D System deployment includes two Defense Centers configured as a high availability pair, you only need to import the SEU on one of the Defense Centers. The second Defense Center receives the SEU as part of the regular synchronization process.

To schedule recurring SEU imports:

Access: P&R Admin/Admin

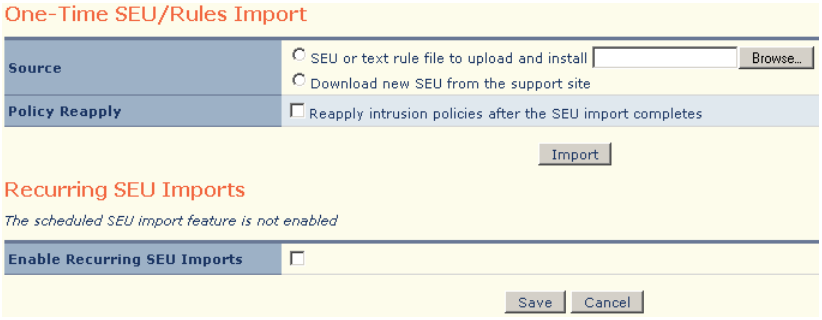
1. Select **Policy & Response > IPS > SEU**.

   The SEU page appears.



2. Click **Import SEU**.

   The Import SEU page appears.



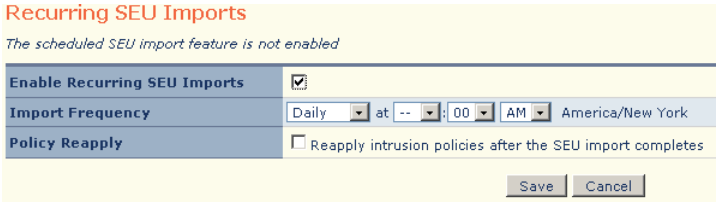
TIP! You can also click **Import Rules** on the Rule Editor page, which you access by selecting **Policy & Response > IPS > Rule Editor**.

3. 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 868 for more information.

Keysight, EX-1004
Page 1855

4. Select **Enable Recurring SEU Imports**.

   The page expands to display options for configuring recurring imports.



   Import status messages appear beneath the **Recurring SEU Imports** section heading. Recurring imports are enabled when you save your settings.

   IMPORTANT! **To disable** recurring imports, clear the **Enable Recurring SEU Imports** check box and click **Save**.

5. 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.

6. 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 SEUs.
7. 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 SEUs.
8. In the **Import Frequency** field, specify the time when you want to start your recurring SEU import.
9. Optionally, select **Reapply intrusion policies after the SEU import completes** to automatically reapply intrusion policies currently applied from this appliance when the SEU import completes.

   Note that you cannot apply intrusion policies to a clustered pair running different versions of the Sourcefire 3D System (for example, if an upgrade on one of the sensors fails). See Using Clustered 3D Sensors on page 250 for more information.

Keysight, EX-1004
Page 1856

10. Click **Save** to enable recurring SEU imports using your settings.

    The status message under the Recurring SEU Imports section heading changes to indicate that the SEU has not yet run.

    The SEU 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 SEU Import Log page displays a red status icon. See Viewing the SEU Import Log on page 1859 for more information. During an import you can also view messages as they occur in the SEU Detail View. See SEU Detail View Fields on page 1863 for more information.

    IMPORTANT! Depending on SEU size and content, several minutes can pass before status messages appear in the SEU Import Log or SEU Detail View.

    Unless you selected **Reapply intrusion policies after the SEU import completes** in step 9, any rule changes are not implemented until the next time you apply the affected intrusion policies. See Applying an Intrusion Policy on page 347 for procedures.

    Applicable subtasks in the SEU import occur in the following order: download, install, base policy update, and policy re-apply. Once 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 SEU.

Importing Local Rule Files

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 (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.

  IPS will automatically assign the rule the next available custom rule SID of 1000000 or greater, and a revision number of 1.

Keysight, EX-1004
Page 1857

- You **must** include the SID assigned by IPS 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 (**Policy & Response > IPS > 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 IPS and a revision number greater than the current revision number. Note that IPS 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 (**Policy & Response > IPS > Rule Editor**), click on the deleted rule category to expand the folder, then click **Edit** next to the rule.
- If you import an intrusion rule with a SID greater than 2147483647, IPS truncates the SID to 2147483647.
- IPS 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 435 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.
- IPS imports local rules preceded with a single pound character (#).
- IPS ignores local rules preceded with two pound characters (##) and does not import them.

This procedure explains how to import local standard text rules from a local machine.

To import local rule files:

Access: P&R Admin/Admin

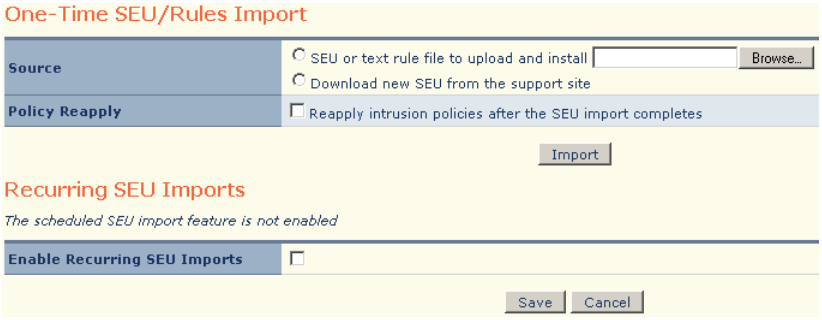
1. Select **Policy & Response > IPS > Rule Editor**.

   The Rule Editor page appears.

Keysight, EX-1004
Page 1858

2. Click **Import Rules**.

   The Import SEU page appears.



TIP! You can also click **Import SEU** on the SEU page, which you access by selecting **Policy & Response > IPS > SEU**.

3. Select **SEU 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! 3D Sensors do **not** use the new rule set for inspection until after you apply their intrusion policies. See Applying an Intrusion Policy on page 347 for procedures.

Viewing the SEU Import Log

Platform: DC/MDC/Standalone Sensor

Component: IPS

The Defense Center or 3D Sensor with IPS generates a record for each SEU 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 SEUs and local rule files that you import, delete any record from the list, and access detailed records for all imported rules and SEU

Keysight, EX-1004
Page 1859

components. The fields in the SEU Import Log table are described in the SEU Import Log Actions table.

SEU Import Log Actions

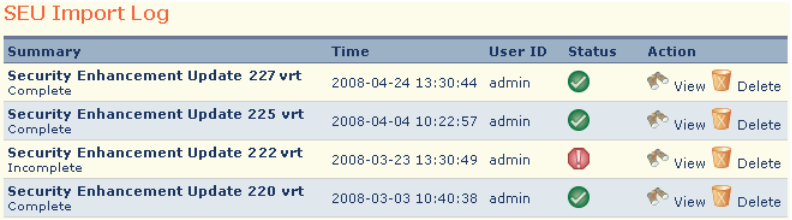
To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the SEU Import Log Table on page 1861.
delete an import file record from the import log, including detailed records for all objects included with the file	click **Delete** next to the file 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 an SEU or local rule file	click **View** next to the file name for the import file.

To view the SEU import log:

Access: P&R Admin/Admin

- Select **Policy & Response > IPS > SEU**.

The SEU page appears.



Keysight, EX-1004
Page 1860

Understanding the SEU Import Log Table

Platform: DC/MDC/Standalone Sensor

Component: IPS

The fields in the list of SEUs and local rule files that you import are described in the SEU Import Log Fields table.

SEU Import Log Fields

Field	Description
Summary	The name of the import file. If the import fails, a brief statement of the reason for the failure appears under the file name.
Time	The time and date that the import started.
User ID	The user name of the user that triggered the import.
Status	Whether the import: • succeeded () • failed or is currently in progress () **TIP!** The red status icon indicating an unsuccessful or incomplete import appears on the SEU page during the import and is replaced by the green icon only when the import has successfully completed.
Action	Allows you to view the SEU Details page for an SEU or local rule file by clicking **View**, or delete the file record and all detailed object records imported with the file by clicking **Delete**. **TIP!** You can view import details as they appear while an SEU import is in progress.

Viewing SEU Details

Platform: DC/MDC/Standalone Sensor

Component: IPS

The SEU Detail View lists a detailed record for each object imported in an SEU 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.

Keysight, EX-1004
Page 1861

The SEU Detail View Actions table below describes specific actions you can perform on an SEU Detail View workflow page.

SEU Detail View Actions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the SEU Detail View Table on page 1863.
sort and constrain records on the current workflow page	find more information in Sorting Drill-down Workflow Pages on page 1612.
drill down to the next page in the workflow	use one of the following methods: • On a drill-down page in a custom workflow, click a value within a row, which constrains the table view and does **not** drill down to the next page. • To drill down to the next workflow page constraining on some rules, select the check boxes next to the rules 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**. The SEU Detail page contains a single workflow page, but you can create a custom workflow with multiple pages. For information, see Creating Custom Workflows on page 1618.
temporarily use a different workflow	click **Workflows**. For information on selecting workflows, see Selecting Workflows on page 1590. For information on creating custom workflows, see Creating Custom Workflows on page 1618.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.

Keysight, EX-1004
Page 1862

SEU Detail View Actions (Continued)

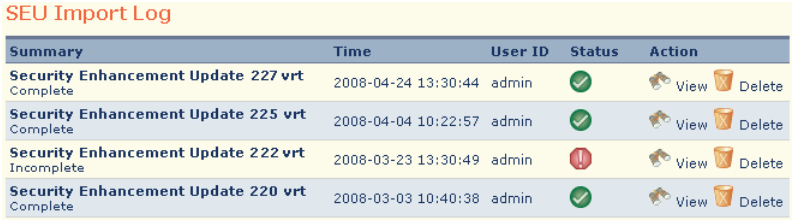
To...	You can...
search the entire SEU log database for SEU import records	click **Search**. From more information, see Searching the SEU Import Log on page 1865.
open a search page prepopulated with the current single constraint	select **Edit Search** or **Save Search** next to Search Constraints. From more information, see the Table View and Drill-down Page Features table on page 1594.

To view the SEU Details page:

Access: P&R Admin/Admin

1. Select **Policy & Response > IPS > SEU**.

   The SEU page appears.



2. Click **View** next to the file whose detailed records you want to view.

   The table view of detailed records appears.

Understanding the SEU Detail View Table

Platform: DC/MDC/Standalone Sensor

Component: IPS

You can view a detailed record for each object imported in an SEU or local rule file. The fields in the SEU Detail View are described in the SEU Detail View Fields table.

SEU Detail 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 SEU components is the component name, such as online help, or Snort.

Keysight, EX-1004
Page 1863

SEU Detail View Fields (Continued)

Field	Description
Type	The type of imported object, which can be `rule`, `SEU component`, `SEU edit` (that is, an updated feature), or `variable`.
Action	An indication that one of the following has occurred for the object type: • `new` (for a rule or variable, this is the first time the object has been stored on this appliance) • `changed` (for an SEU component or rule, the SEU component has been modified, or the rule has a higher revision number and the same GID and SID) • `collision` (for an SEU 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 SEU) • `enabled` (for an SEU edit, a preprocessor, rule, or other feature provided by the SEU 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 an SEU or local rule file, the import failed) • `apply` (the **Reapply intrusion policies after the SEU import completes** option was enabled for the import)
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 489 for more information.
SID	The SID for a rule.
Rev	The revision number for a rule.
Policy	`All`, indicating that a rule is included in all default policies.

Keysight, EX-1004
Page 1864

SEU Detail View Fields (Continued)

Field	Description
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 SEU Detail View is constrained by default to SEU records.

Searching the SEU Import Log

Platform: DC/MDC/Standalone Sensor

Component: IPS

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 re-use later.

TIP! You search the entire SEU import log database even when you initiate a search by clicking **Search** on the toolbar from the SEU Detail 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 1556 for more information.

Keysight, EX-1004
Page 1865

The search criteria you can use are described in the SEU Import Log Search Criteria table. Note that record searches are case-insensitive. For example, searching for `RULE` or `rule` yields the same results.

SEU 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 1556 for the syntax for entering time.	`> 2006-01-15 13:30:00` returns all rule records imported after January 15, 2006 at 1:30pm.
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.	`*dhcp*` returns all rule records with DHCP in the Message field.
Type	Specify the type of record, which can be `rule, SEU component, SEU edit` (that is, an updated feature), or `variable`.	`seu*` returns all imported SEU components, such as updated online help or updated versions of Snort.
Action	Specify an action for the object you want to view. See the SEU Detail View Actions table on page 1862 for a list of actions you can specify.	`new` returns all rules imported for the first time on the appliance.
GID	Specify the generator ID for the rule.	`3` returns all shared object rules.
SID	Specify a single signature ID for a rule. You cannot specify a list of SIDs.	`923` returns the record for the rule with the SID 923.
Rev	Specify the revision number for the rule.	`3` returns rules with the revision number 3.
Policy	Specify the default policy the rules is imported into.	`All` returns rules imported into all default policies.
SEU	Specify the SEU filename.	*`filename`* returns all records for the specified import file.
Details	Specify details on the imported object.	`previously*` 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 1550.

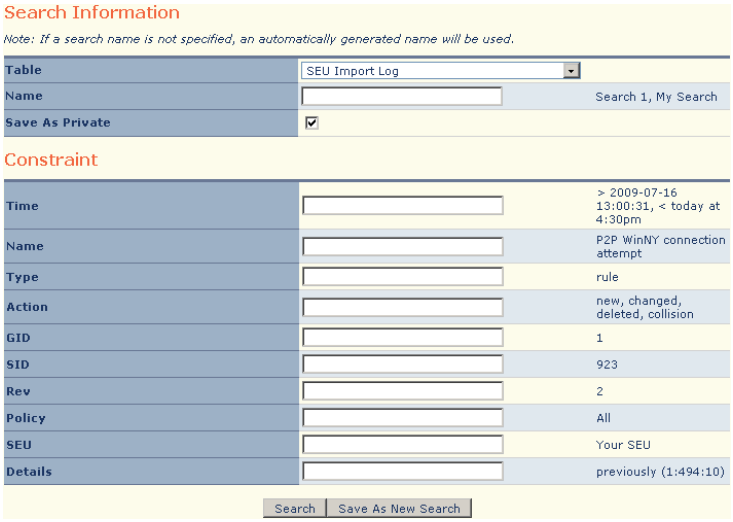
Keysight, EX-1004
Page 1866

To search the SEU Import Log:

Access: IPS/Admin

1. Select **Analysis & Reporting > Searches > SEU Import Log**.

   The SEU Import Log search page appears.



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 web interface automatically creates one when you save it.

3. Enter your search criteria in the appropriate fields, as described in the SEU Import Log Search Criteria table. If you enter multiple criteria, the search returns 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 SEU import log workflow. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

Keysight, EX-1004
Page 1867

- 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.

Keysight, EX-1004
Page 1868

Chapter 43
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, intrusion event information, and statistics for the Data Correlator and RNA 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 3D Sensor. The following sections provide more information about the monitoring features that the system provides:

- Viewing Host Statistics on page 1870 describes how to view host information such as:
  - system uptime
  - disk and memory usage
  - RNA process statistics
  - 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 1891.

Keysight, EX-1004
Page 1869

- Monitoring System Status and Disk Space Usage on page 1874 describes how to view basic event and disk partition information.
- Viewing System Process Status on page 1875 describes how to view basic process status.
- Understanding Running Processes on page 1877 describes the basic system processes that run on the appliance.
- Viewing IPS Performance Statistics on page 1883 describes how to view IPS performance statistics and how to generate graphs based on these statistics.
- Viewing RNA Performance Statistics on page 1886 describes how to view RNA performance statistics and how to generate graphs based on these statistics.

Viewing Host Statistics

Platform: DC/MDC/3D Sensor

Component: None

The Statistics page lists the current status of following:

- general host statistics; see the Host Statistics table on page 1870 for details
- Data Correlator statistics (Defense Center only - requires RNA); see the Data Correlator Process Statistics table on page 1871 for details
- RNA process statistics (Defense Center only - requires RNA); see the RNA Process Statistics table on page 1872 for details
- intrusion event information (requires IPS); see the Intrusion Event Information table on page 1873 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.

Keysight, EX-1004
Page 1870

Host Statistics (Continued)

Category	Description
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 1874 for more information.
Processes	A summary of the processes running on the system. See Viewing System Process Status on page 1875 for more information.

If your Sourcefire 3D System deployment includes a Defense Center managing 3D Sensors with RNA, you can also view statistics about the Data Correlator and RNA processes for the current day. As the 3D Sensors perform data acquisition, decoding, and analysis, the RNA process correlates the data with the fingerprint and vulnerability databases, and 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 RNA network map.

The statistics that appear for RNA and the Data Correlator are averages for the current day, using statistics gathered between 12:00AM and 11:59PM for each detection engine.

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 RNA events that the Data Correlator receives and processes per second
Flows/Sec	Number of flows 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

Keysight, EX-1004
Page 1871

The RNA Process Statistics table describes the statistics displayed for the RNA process.

RNA Process Statistics

Category	Description
Packets Dropped (%)	Average percentage of packets dropped by the RNA process for the current day
Mbits/Second	Average number of megabits per second processed by the RNA process for the current day
Packets/Second	Average number of packets per second processed by the RNA process for the current day
CPU Usage - User (%)	Average percentage of CPU time spent by user processes for the current day
CPU Usage - System (%)	Average percentage of CPU time spent by system processes for the current day
VmSize (KB)	Average size of memory allocated to the RNA process for the current day, in kilobytes
VmRSS (KB)	Average amount of memory used by the RNA process for the current day, in kilobytes

On 3D Sensors with IPS and on Defense Centers that manage sensors with IPS, you can also view the time and date 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 in the database.

IMPORTANT! The information in the Intrusion Event Information section of the Statistics page is based on intrusion events stored on the sensor rather than those sent to the Defense Center. If you manage your sensor so that intrusion events are not stored locally, no intrusion event information is listed on this page. This is also the case for 3D Sensors that cannot store events locally.

Keysight, EX-1004
Page 1872

The Intrusion Event Information table describes the statistics displayed in the Intrusion Event Information section of the Statistics page.

Intrusion Event Information

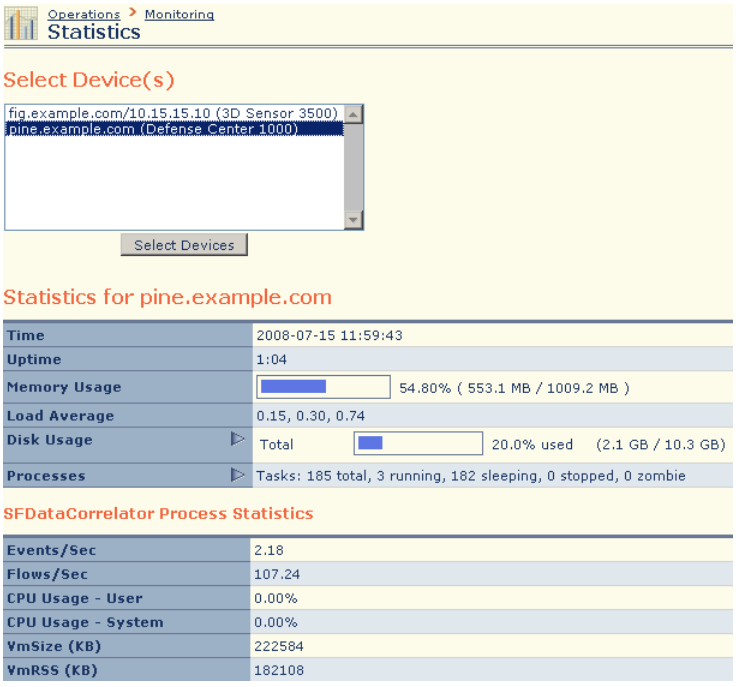
Statistic	Description
Last Alert Was	The date and time that the last event occurred
Total Events Last Hour	The total number of events that occurred in the past hour
Total Events Last Day	The total number of events that occurred in the past twenty-four hours
Total Events in Database	The total number of events in the events database

To view the Statistics page:

Access: Maint/Admin

1. Select **Operations > Monitoring > Statistics**.

   The Statistics page appears. The Defense Center version of the page is shown below.



Keysight, EX-1004
Page 1873

2. On the Defense Center, you can also list statistics for managed sensors. From the Select Device(s) box and 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 that you selected.

Monitoring System Status and Disk Space Usage

Platform: DC/MDC/3D Sensor

Component: None

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 1891.

To access disk usage information:

Access: Maint/Admin

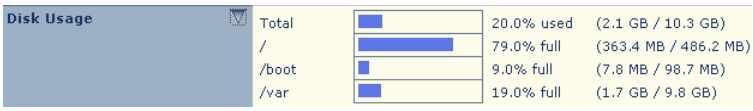
1. Select **Operations > Monitoring > Statistics**.

   The Statistics page appears.

2. Click the down arrow next to **Disk Usage** to expand it.



   The Disk Usage section expands.



On the Defense Center, to view disk usage information for a specific sensor:

Access: Maint/Admin

1. Select the sensor name from the **Select Device(s)** box, and click **Select Devices**.

   The page reloads, listing host statistics for each sensor you selected.

2. Click the down arrow next to Disk Usage to expand it.

   The Disk Usage section expands.

Keysight, EX-1004
Page 1874

Viewing System Process Status

Platform: DC/MDC/3D Sensor

Component: None

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 sensors with a Defense Center, you can use the Defense Center's web interface to view the process status for any managed sensor.

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
Nice	The *nice* value, which is a value that indicates the scheduling priority of a process. Values range between -20 (highest priority) and 19 (lowest priority)
Size	The memory size used by the process (in kilobytes, unless the value is followed by `m`, which indicates megabytes)
Res	The amount of resident paging files in memory (in kilobytes, unless the value is followed by `m`, which indicates megabytes)
State	The process state: • D - process is in uninterruptible sleep (usually Input/Output) • N - process has a positive nice value • R - process is runnable (on queue to run) • S - process is in sleep mode • T - process is being traced or stopped • W - process is paging • X - process is dead • Z - process is defunct • < - process has a negative nice value
Time	The amount of time (in hours:minutes:seconds) that the process has been running

Keysight, EX-1004
Page 1875

Process Status (Continued)

Column	Description
Cpu	The percentage of CPU that the process is using
Command	The executable name of the process

To expand the process list:

Access: Maint/Admin

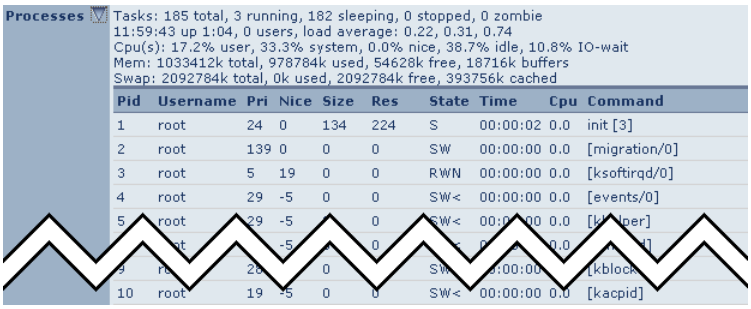
1. Select **Operations > Monitoring > Statistics**.

   The Statistics page appears.

2. On the Defense Center, select the device or devices you want to view process statistics for and click **Select Devices**.

3. Click the down arrow next to **Processes**.

   The process list expands, listing general process status that includes the number and types of running tasks, the current time, the current system uptime, the system load average, CPU, memory, and swap information, and specific information about each running process.



Cpu(s) lists the following CPU usage information:

- user process usage percentage
- system process usage percentage
- nice usage percentage (CPU usage of processes that have a negative nice value, indicating a higher priority)

  Nice values indicate the scheduled priority for system processes and can range between -20 (highest priority) and 19 (lowest priority).
- idle usage percentage

Mem lists the following memory usage information:

- total number of kilobytes in memory
- total number of used kilobytes in memory

Keysight, EX-1004
Page 1876

- 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 1877.

To collapse the process list:

Access: Maint/Admin

▶ Click the up arrow next to **Processes**.

The process list collapses.

Understanding Running Processes

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 1878
- Understanding Executables and System Utilities on page 1880

Keysight, EX-1004
Page 1877

Understanding System Daemons

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
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
rnareportd	Manages RNA reports

Keysight, EX-1004
Page 1878

System Daemons (Continued)

Daemon	Description
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
sfestreamer (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
sfreactd	Manages Check Point OPSEC integration; only seen if Checkpoint SAM support is enabled
SFRemediateD (Defense Center only - requires RNA)	Manages remediation responses
sftimeserviced (Defense Center only)	Forwards time synchronization messages to managed sensors
sfmbservice (requires IPS)	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 3D Sensor to a Defense Center or, in a high availability environment, between Defense Centers
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

Keysight, EX-1004
Page 1879

Understanding Executables and System Utilities

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
correlator (Defense Center only - requires RNA)	Analyzes binary files created by RNA to generate events, flow data, and the network map
cp	Utility that copies files
df	Utility that lists the amount of free space on the appliance
echo	Utility that writes content to standard output
egrep	Utility that searches files and folders for specified input; supports extended set of regular expressions not supported in standard grep
find	Utility that recursively searches directories for specified input
grep	Utility that searches files and directories for specified input
halt	Utility that stops the server
httpsdctl	Handles secure Apache Web processes
hwclock	Utility that allows access to the hardware clock

Keysight, EX-1004
Page 1880

System Executables and Utilities (Continued)

Executable	Description
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 1738 for more information about access configuration.
iptables-restore	Handles iptables file restoration
iptables-save	Handles saved changes to the iptables
kill	Utility that can be used to end a session and process
killall	Utility that can be used to end all sessions and processes
ksh	Public domain version of the Korn shell
logger	Utility that provides a way to access the syslog daemon from the command line
md5sum	Utility that prints checksums and block counts for specified files
mv	Utility that moves (renames) files
myisamchk	Indicates database table checking and repairing
mysql	Indicates a database process; multiple instances may appear
openssl	Indicates authentication certificate creation
perl	Indicates a perl process
ps	Utility that writes process information to standard output
RNA (requires RNA)	Captures packets, decodes and performs session reassembly, correlating acquired data with the RNA fingerprint database, then generates binary files that the Data Correlator processes to generate the network map and to populate the database with events and flow data
sed	Utility used to edit one or more text files

Keysight, EX-1004
Page 1881

System Executables and Utilities (Continued)

Executable	Description
sfheartbeat	Identifies a heartbeat broadcast, indicating that the appliance is active; heartbeat used to maintain contact between a sensor and Defense Center
sfmb	Indicates a message broker process; handles communication between Defense Centers and sensor.
sfsnort (requires IPS)	Indicates that Snort is running
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
ssh	Indicates a Secure Shell (SSH) connection to the appliance
sudo	Indicates a sudo process, which allows users other than admin to run executables
top	Utility that displays information about the top CPU processes
touch	Utility that can be used to change the access and modification times of specified files
vim	Utility used to edit text files
wc	Utility that performs line, word, and byte counts on specified files

Keysight, EX-1004
Page 1882

Viewing IPS Performance Statistics

Platform: DC/3D Sensor

Component: IPS

The IPS performance statistics page allows you to generate graphs that depict performance statistics for IPS 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, and the percent of packets uninspected by Snort. These graphs can show statistics for the last hour, last day, last week, or last month of operation.

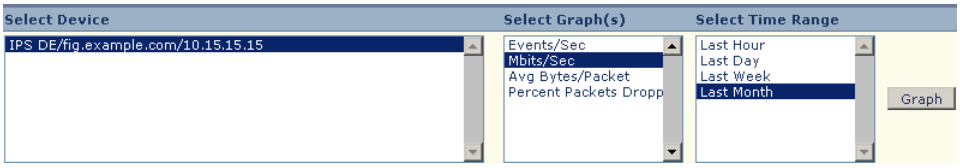
IMPORTANT! Because of the way traffic is processed on 3Dx800 sensors, performance statistics for those sensors are under reported. IPS performance statistics refer only to the data stored locally on the 3D Sensor.

To view the IPS performance statistics:

Access: Maint/Admin

▶ Select **Operations > Monitoring > Performance > IPS**.

The IPS page appears. The Defense Center version of the page is shown below.



See the following sections for more information:

- Generating IPS Performance Statistics Graphs on page 1884
- Saving IPS Performance Statistics Graphs on page 1885

Keysight, EX-1004
Page 1883

Generating IPS Performance Statistics Graphs

Platform: DC/3D Sensor

Component: IPS

You can generate graphs that depict performance statistics for a Defense Center or a 3D Sensor with IPS based on the number of events per second, megabits per second, or average bytes per packet.

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 IPS Performance Statistics Graph Types table lists the available graph types.

IPS Performance Statistics Graph Types

Graph Type	Output
Events/Sec	Displays a graph that represents the number of events that are generated on the sensor per second
Mbits/Sec	Displays a graph that represents the number of megabits of traffic that pass through the sensor per second
Avg Bytes/Packet	Displays a graph that represents the average number of bytes included in each packet
Percent Packets Dropped	This graph depicts the average percentage of uninspected packets across all detection resources (instances of Snort) assigned to the selected detection engine. If you assign two detection resources to a detection engine that has two interface sets and each interface set is connected to a different network segment, then an average of 50% may indicate that one segment has a 90% drop rate and the other has a 10% drop rate. It may also indicate that both segments have a drop rate of 50%. The graph only represents the total % drop when there is a single detection resource assigned to a selected detection engine.

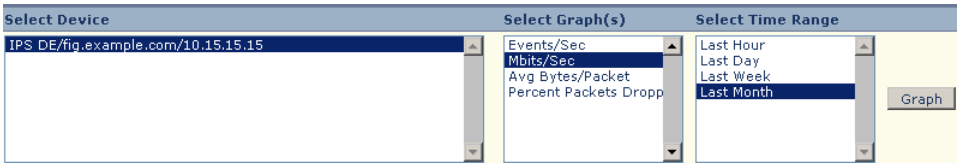
Keysight, EX-1004
Page 1884

To generate IPS performance statistics graphs:

Access: Maint/Admin

1. Select **Operations > Monitoring > Performance > IPS.**

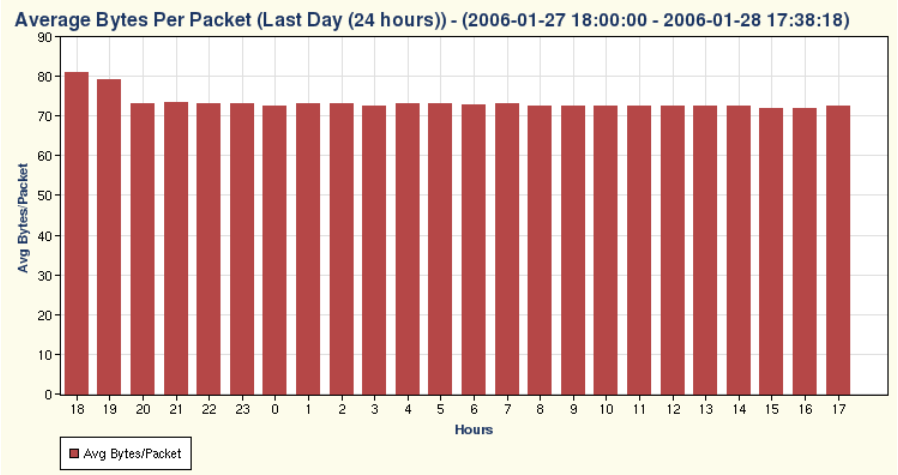
   The IPS page appears. The Defense Center version of the page is shown below.



2. From the **Select Device** list, select the detection engines 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.



Saving IPS Performance Statistics Graphs

Platform: DC/3D Sensor

Component: IPS

After you have generated an IPS performance statistics graph, you can save the graph as a graphic file for later use.

Keysight, EX-1004
Page 1885

To save the graph:

Access: Maint/Admin

▶ Right-click on the graph and follow the instructions for your browser to save the image.

Viewing RNA Performance Statistics

Platform: DC

Component: RNA

The RNA Performance page allows you to generate graphs that display RNA-related performance statistics over a specific period of time. Graphs can be generated to display:

- the number of events generated by the Data Correlator per second
- the number of megabits analyzed by the RNA process per second
- average number of bytes included in each packet analyzed by the RNA process
- the percentage of packets dropped by RNA
- the number of packets, in thousands, analyzed by the RNA process per second
- the number of established connections analyzed by the RNA process per second

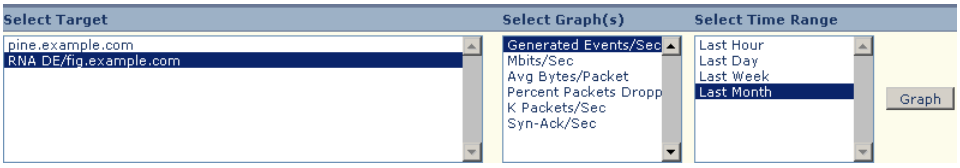
These graphs can show statistics for the last hour, last day, last week, or last month of operation.

To access the RNA Performance page:

Access: Maint/Admin

▶ Select **Operations > Monitoring > Performance > RNA.**

The RNA page appears.



See the following sections for more information:

- Generating RNA Performance Statistics Graphs on page 1887
- Saving RNA Performance Statistics Graphs on page 1889

Keysight, EX-1004
Page 1886

Generating RNA Performance Statistics Graphs

Platform: DC

Component: RNA

You can generate graphs that display performance statistics for managed 3D Sensors with RNA.

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 RNA Performance Statistics Graph Types table lists the available graph types.

RNA Performance Statistics Graph Types

Graph Type	Output
Processed Events/Sec	Displays a graph that represents the number of events that the Data Correlator processes per second
Processed Flows/Sec	Displays a graph that represents the number of flows that the Data Correlator processes per second
Generated Events/Sec	Displays a graph that represents the number of events that RNA generates per second
Mbits/Sec	Displays a graph that represents the number of megabits of traffic that are analyzed by the RNA process per second
Avg Bytes/Packet	Displays a graph that represents the average number of bytes included in each packet analyzed by the RNA process
Percent Packets Dropped	Displays a graph that represents the percentage of packets dropped by RNA
K Packets/Sec	Displays a graph that represents the number of packets analyzed by the RNA process per second, in thousands
Syn/Ack/Sec	Displays a graph that represents the number of established connections observed by the RNA process per second

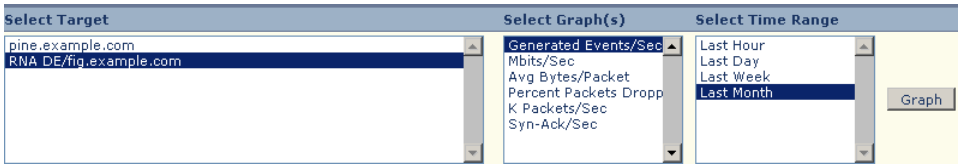
Keysight, EX-1004
Page 1887

To generate RNA performance statistics graphs:

Access: Maint/Admin

1. Select **Operations > Monitoring > Performance > RNA.**

   The RNA page appears.



2. From the **Select Target** list, select the Defense Center, the managed 3D Sensors, or the detection engines that you want to include.

   Depending on whether you select a detection engine or a sensor, 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.

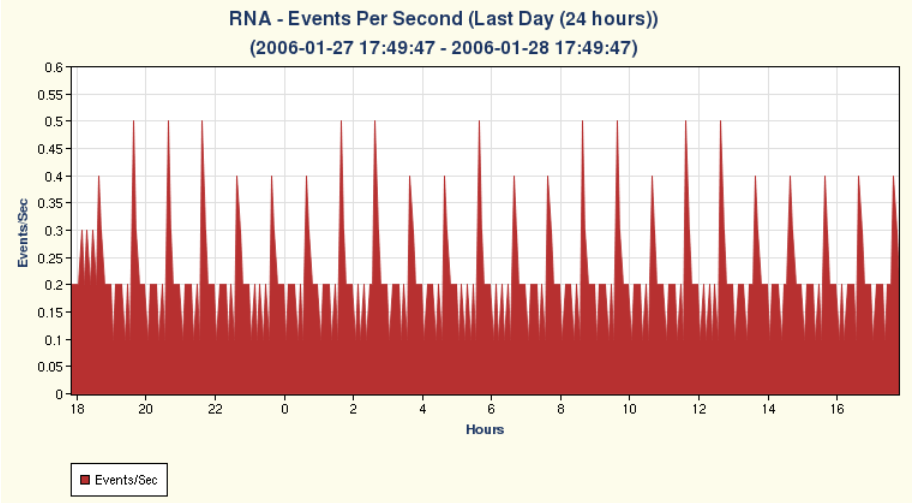
   TIP! You can select multiple graphs by holding down the Ctrl or Shift keys while clicking on the graph type.

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**.

   The graph appears, displaying the information you specified. If you selected multiple graphs, each graph appears on the page.



Keysight, EX-1004
Page 1888

Saving RNA Performance Statistics Graphs

Platform: DC

Component: RNA

After you have generated an RNA performance statistics graph, you can save the graph as a graphic file for later use.

To save the graph:

Access: Maint/Admin

1. Create an RNA performance statistic graph as described in Generating RNA Performance Statistics Graphs on page 1887.
2. Right-click on the graph and follow the instructions for your browser to save the image.

Keysight, EX-1004
Page 1889

Chapter 44
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 one of the default health policies. 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.

At 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:

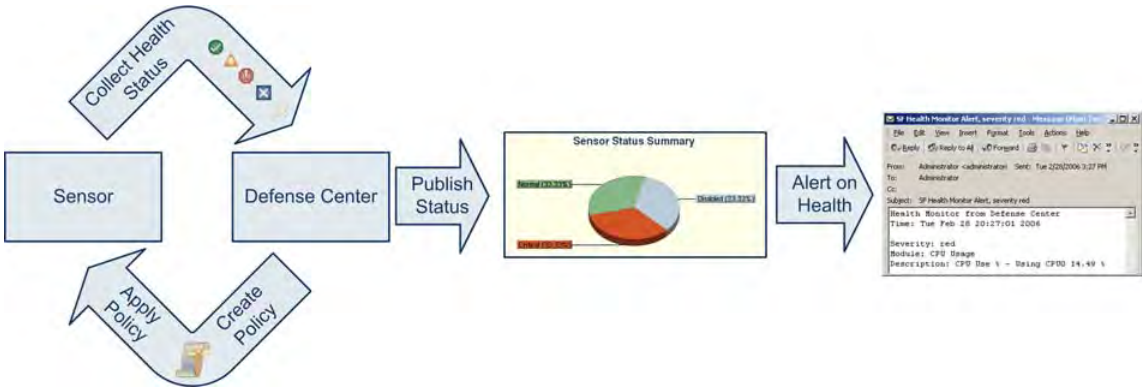
- Understanding Health Monitoring on page 1891
- Configuring Health Policies on page 1895

Keysight, EX-1004
Page 1890

- Using the Health Monitor Blacklist on page 1935
- Configuring Health Monitor Alerts on page 1940
- Using the Health Monitor on page 1945
- Using Appliance Health Monitors on page 1947
- Working with Health Events on page 1957

Understanding Health Monitoring

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 appliances 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

Keysight, EX-1004
Page 1891

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 Admin access privileges can access system health data. For more information on assigning user privileges, see Modifying User Privileges and Options on page 1684.

IMPORTANT! Except for the Defense Center, Sourcefire 3D System appliances do not have health monitoring policies applied to them by default. If you want to monitor the health of a managed appliance, you have to apply a health policy to that appliance. For more information on available default health policies you can apply to an appliance, see Predefined Health Policies on page 1896. For more information on creating customized health policies, see Creating Health Policies on page 1902. For details on applying policies, see Applying Health Policies on page 1928.

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 1892
- Understanding Health Modules on page 1893
- Understanding Health Monitoring Configuration on page 1895

Understanding Health Policies

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 one of the five default health policies to each appliance.

Keysight, EX-1004
Page 1892

Understanding Health Modules

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
Appliance Heartbeat	This module determines if an appliance heartbeat is being heard from the sensor and alerts based on the sensor heartbeat status.
Automatic Application Bypass Status	This module determines if a detection engine 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 Temperature	This module determines if the CPU on the sensor is overheated and alerts when the temperature exceeds temperatures configured for the module. This module only runs on 3Dx800 sensors.
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.
Card Reset	This module checks for network cards which have restarted due to hardware failure and alerts when a reset occurs.
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.
Event Stream Status	This module compares the number of events per second to the limits configured for this module and alerts if the limits are exceeded. If the Event Stream is zero, the eStreamer process may be down or the Defense Center may not be sending events. This module only runs on Master Defense Centers.
Fan Alarm	This module determines if fans need to be replaced on the sensor and alerts based on the fan status. This module only runs on 3Dx800 sensors.
Hardware Alarms	This module determines if hardware needs to be replaced on a 3D9900 sensor and alerts based on the hardware status. The module also reports on the status of hardware-related daemons. This module only runs on 3D9900 sensors. For more information on the details reported for 3D9900 sensors, see Interpreting Hardware Alert Details for 3D9900 Sensors on page 1962.

Keysight, EX-1004
Page 1893

Health Modules (Continued)

Module	Description
Health Monitor Process	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. This module only runs on Defense Centers.
IPS Event Rate	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 IPS Event Rate is zero, the IPS process may be down or the 3D Sensor may not be sending events. Select **Analysis & Reporting > Event Summary > Intrusion Event Statistics** to check if events are being received from the sensor.
Link State Propagation	This module determines when a link in a paired inline interface set fails and triggers the link state propagation mode.
Memory Usage	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.
PEP Status	This module monitors the application of PEP rules to interface sets on a 3D9900. If PEP rules cannot be applied to interfaces in an interface set, the module generates an alert.
Power Supply	This module determines if power supplies on the sensor require replacement and alerts based on the power supply status. This module only runs on Series 2 DC3000s, MDC3000s, 3Dx800, 3D9900, 3D3500, 3D4500, and 3D6500 appliances.
Process Status	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 crashed 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.
RNA Event Status	This module indicates whether a specified period of time has passed since any RNA events have been detected by a sensor.
RNA Host License Limit	This module determines if sufficient RNA host licenses remain and alerts based on the warning level configured for the module.
Time Synchronization Status	This module tracks the synchronization of a sensor 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.
Traffic Status	This module determines if the sensor currently collects traffic and alerts based on the traffic status.

Keysight, EX-1004
Page 1894

Understanding Health Monitoring Configuration

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 one of the default policies provided for that purpose.

   For more information on setting up health policies, see Configuring Health Policies on page 1895.

2. Apply a health policy to each appliance where you want to track health status. For information on the default health policies available for immediate application, see Predefined Health Policies on page 1896.

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 1940.

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 Table Events View. For more information about viewing system health data, see the following topics:

- Using the Health Monitor on page 1945
- Using Appliance Health Monitors on page 1947
- Working with Health Events on page 1957

Configuring Health Policies

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 1891.

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 policies provided for you. You can also import a health policy exported from another Defense Center.

Keysight, EX-1004
Page 1895

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 and Master Defense Center automatically, see Default Health Policy on page 1899.

For more information, see the following topics:

- Predefined Health Policies on page 1896
- Creating Health Policies on page 1902
- Applying Health Policies on page 1928
- Editing Health Policies on page 1931
- Deleting Health Policies on page 1934

Predefined Health Policies

The Defense Center health monitor includes several default health policies 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. To also monitor sensor health, you can push health policies to 3D Sensors.

IMPORTANT! You cannot apply a health policy to RNA Software for Red Hat Linux or Crossbeam-based software sensors.

For more information, see the following topics:

- Default 3D Sensor Health Policy on page 1897
- Default 3Dx800 Health Policy on page 1897
- Suggested 3D9900 Health Policy on page 1898
- Default Health Policy on page 1899
- Default Intrusion Sensor Health Policy on page 1900
- Default IPS (3Dx800 only) Health Policy on page 1901
- Default RNA Sensor Health Policy on page 1901

Keysight, EX-1004
Page 1896

Default 3D Sensor Health Policy

Use the Default 3D Sensor Health Policy to monitor health on any 3D Sensor. Enabled health modules for this policy are listed in the Enabled Health Modules: 3D Sensor Health Policy table.

Enabled Health Modules: 3D Sensor Health Policy

Module	For more information, see...
Automatic Application Bypass Status	Configuring Automatic Application Bypass Monitoring on page 1907
Disk Usage	Configuring Disk Usage Monitoring on page 1911
IPS Event Rate	Configuring IPS Event Rate Monitoring on page 1917
Link State Propagation	Configuring Link State Propagation Monitoring on page 1919
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Power Supply	Configuring Power Supply Monitoring on page 1922
Process Status	Configuring Process Status Monitoring on page 1923
Traffic Status	Configuring Traffic Status Monitoring on page 1927

Default 3Dx800 Health Policy

Use the Default 3Dx800 Health Policy to monitor health on 3Dx800 sensors. Enabled health modules for this policy are listed in the Enabled Health Modules: Default 3Dx800 Health Policy table. Note that the Hardware Alarm module should be used instead of the Power Supply module to monitor power supply health on the 3Dx800 sensor models.

Enabled Health Modules: Default 3Dx800 Health Policy

Module	For more information, see...
Automatic Application Bypass Status	Configuring Automatic Application Bypass Monitoring on page 1907
CPU Temperature	Configuring CPU Temperature Monitoring on page 1908

Keysight, EX-1004
Page 1897

Enabled Health Modules: Default 3Dx800 Health Policy (Continued)

Module	For more information, see...
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Fan Alarm	Configuring Fan Monitoring on page 1914
Hardware Alarms	Configuring Hardware Monitoring on page 1915
IPS Event Rate	Configuring IPS Event Rate Monitoring on page 1917
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Process Status	Configuring Process Status Monitoring on page 1923
Traffic Status	Configuring Traffic Status Monitoring on page 1927

Suggested 3D9900 Health Policy

The Defense Center interface does not include a default health policy specifically for 3D9900 sensors. Sourcefire recommends that you start with the default 3D Sensor policy and enable the Hardware Alarms module. If the sensor will be running RNA, enable the RNA Process module as well.

Health modules that should be enabled when creating a policy for this type of sensor are listed in the Suggested Health Modules: 3D9900 Health Policy table. Note that the CPU Usage module cannot be enabled when monitoring 3D9900 sensor models. CPU usage for a 3D9900 may reach 100% during normal sensor operation, so the data provided by the module would generate misleading events.

Suggested Health Modules: 3D9900 Health Policy

Module	For more information, see...
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Hardware Alarms	Configuring Hardware Monitoring on page 1915
IPS Event Rate	Configuring IPS Event Rate Monitoring on page 1917
Link State Propagation	Configuring Link State Propagation Monitoring on page 1919
Memory Usage	Configuring Memory Usage Monitoring on page 1920

Keysight, EX-1004
Page 1898

Suggested Health Modules: 3D9900 Health Policy (Continued) (Continued)

Module	For more information, see...
PEP Status	Configuring PEP Status Monitoring on page 1921
Power Supply	Configuring Power Supply Monitoring on page 1922
Process Status	Configuring Process Status Monitoring on page 1923
Traffic Status	Configuring Traffic Status Monitoring on page 1927

Default Health Policy

Use the Default Health Policy to monitor health on a Defense Center. Enabled health modules for this policy are listed in the Enabled Defense Center Health Modules - Default Health Policy table.

Enabled Defense Center Health Modules - Default Health Policy

Module	For more information, see...
Automatic Application Bypass Status	Configuring Automatic Application Bypass Monitoring on page 1907
Appliance Heartbeat	Configuring Appliance Heartbeat Monitoring on page 1906
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Link State Propagation	Configuring Link State Propagation Monitoring on page 1919
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Time Synchronization Status	Configuring Time Synchronization Monitoring on page 1927
Power Supply	Configuring Power Supply Monitoring on page 1922
Process Status	Configuring Process Status Monitoring on page 1923
RNA Host License Limit	Configuring RNA Host Usage Monitoring on page 1925

Keysight, EX-1004
Page 1899

Use the Default Health Policy to monitor health on a Master Defense Center. Enabled health modules for this policy are listed in the Enabled MDC Health Modules - Default Health Policy table.

Enabled MDC Health Modules - Default Health Policy

Module	For more information, see...
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Event Stream	Configuring Event Stream Monitoring on page 1913
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Process Status	Configuring Process Status Monitoring on page 1923
RNA Host License Limit	Configuring RNA Host Usage Monitoring on page 1925

Default Intrusion Sensor Health Policy

Use the Default IPS Health Policy to monitor health on legacy Intrusion Sensors that you have not upgraded to Version 4.10. Enabled health modules for this policy are listed in the Enabled Health Modules: Default Intrusion Sensor Health Policy table.

Enabled Health Modules: Default Intrusion Sensor Health Policy

Module	For more information, see...
Automatic Application Bypass Status	Configuring Automatic Application Bypass Monitoring on page 1907
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Health Monitor Process	Configuring Health Status Monitoring on page 1916
IPS Event Rate	Configuring IPS Event Rate Monitoring on page 1917
Link State Propagation	Configuring Link State Propagation Monitoring on page 1919
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Power Supply	Configuring Power Supply Monitoring on page 1922

Keysight, EX-1004
Page 1900

Enabled Health Modules: Default Intrusion Sensor Health Policy (Continued)

Module	For more information, see...
Process Status	Configuring Process Status Monitoring on page 1923
Traffic Status	Configuring Traffic Status Monitoring on page 1927

Default IPS (3Dx800 only) Health Policy

Use the Default IPS (3Dx800 only) Health Policy to monitor IPS health on 3Dx800 sensors. Enabled health modules for this policy are listed in the Enabled Health Modules: Default 3Dx800 Health Policy table. Note that the Hardware Alarm module should be used instead of the Power Supply module to monitor power supply health on the 3Dx800 sensor models.

Enabled Health Modules: Default IPS (3Dx800 only) Health Policy

Module	For more information, see...
Automatic Application Bypass Status	Configuring Automatic Application Bypass Monitoring on page 1907
CPU Temperature	Configuring CPU Temperature Monitoring on page 1908
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Fan Alarm	Configuring Fan Monitoring on page 1914
Hardware Alarms	Configuring Hardware Monitoring on page 1915
IPS Event Rate	Configuring IPS Event Rate Monitoring on page 1917
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Process Status	Configuring Process Status Monitoring on page 1923
Traffic Status	Configuring Traffic Status Monitoring on page 1927

Default RNA Sensor Health Policy

Use the Default RNA Sensor Health Policy to monitor health on legacy RNA Sensors that you have not upgraded to Version 4.10. Enabled health modules for

Keysight, EX-1004
Page 1901

this policy are listed in the Enabled Health Modules: Default RNA Sensor Health Policy table.

Enabled Health Modules: Default RNA Sensor Health Policy

Module	For more information, see...
Automatic Application Bypass Status	Configuring Automatic Application Bypass Monitoring on page 1907
Disk Usage	Configuring Disk Usage Monitoring on page 1911
Link State Propagation	Configuring Link State Propagation Monitoring on page 1919
Memory Usage	Configuring Memory Usage Monitoring on page 1920
Power Supply	Configuring Power Supply Monitoring on page 1922
Process Status	Configuring Process Status Monitoring on page 1923
RNA Host License Limit	Configuring RNA Host Usage Monitoring on page 1925
Traffic Status	Configuring Traffic Status Monitoring on page 1927

Creating Health Policies

Platform: DC/MDC

Component: None

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 1997.

Keysight, EX-1004
Page 1902

To create a health policy:

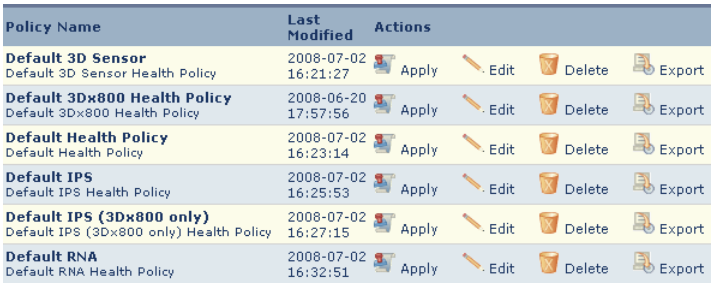
Access: Maint/Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

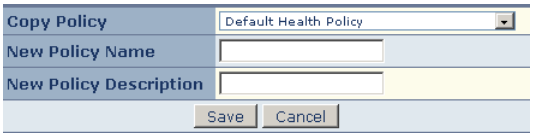
2. On the toolbar, click **Health Policy**.

   The Health Policy page appears.



3. Click **Create Policy** to create a new policy.

   The Create Health Policy page appears.

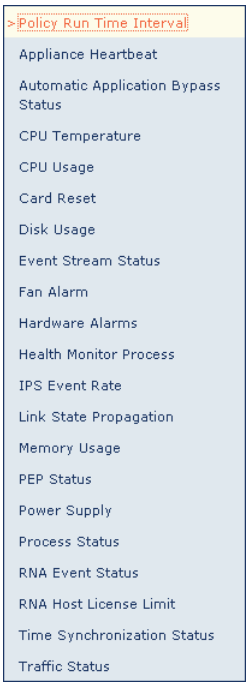


4. Select the existing policy that you want to use as the basis for the new policy from the **Copy Policy** drop-down list.
5. Enter a name for the policy.
6. Enter a description for the policy.

Keysight, EX-1004
Page 1903

7. Select **Save** to save the policy information.

   The Health Policy Configuration page appears, including a list of the modules.



8. 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 1905
   - Configuring Appliance Heartbeat Monitoring on page 1906
   - Configuring Automatic Application Bypass Monitoring on page 1907
   - Configuring CPU Temperature Monitoring on page 1908
   - Configuring CPU Usage Monitoring on page 1909
   - Configuring Card Reset Monitoring on page 1911
   - Configuring Disk Usage Monitoring on page 1911
   - Configuring Event Stream Monitoring on page 1913
   - Configuring Fan Monitoring on page 1914
   - Configuring Hardware Monitoring on page 1915
   - Configuring Health Status Monitoring on page 1916
   - Configuring IPS Event Rate Monitoring on page 1917
   - Configuring Link State Propagation Monitoring on page 1919
   - Configuring Memory Usage Monitoring on page 1920
   - Configuring PEP Status Monitoring on page 1921

Keysight, EX-1004
Page 1904

- Configuring Power Supply Monitoring on page 1922
- Configuring Process Status Monitoring on page 1923
- Configuring RNA Event Status Monitoring on page 1924
- Configuring RNA Host Usage Monitoring on page 1925
- Configuring Time Synchronization Monitoring on page 1927
- Configuring Traffic Status Monitoring on page 1927

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.

9. Click **Save** 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 1928.

Configuring Policy Run Time Intervals

Platform: DC/MDC

Component: None

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: Maint/Admin

1. On the Health Policy Configuration page, select **Policy Run Time Interval**.

   The Health Policy Configuration - Policy Run Time Interval page appears.



2. In the **Run Interval (mins)** field, enter the time in minutes that you want to elapse between automatic repetitions of the test.

Keysight, EX-1004
Page 1905

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 1928 for more information.

Configuring Appliance Heartbeat Monitoring

Platform: DC

Component: None

Supported Platform: DC

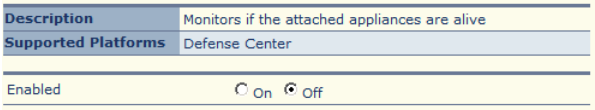
The Defense Center receives heartbeats from its managed appliances once every two minutes or every 200 events, whichever comes first, as an indicator that the appliance is running and communicating properly with the Defense Center. Use the Appliance Heartbeat health status module to track whether the Defense Center receives heartbeats from managed appliances. If the Defense Center does not detect a heartbeat from a appliance, 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: Maint/Admin

1. In the Health Policy Configuration page, select **Appliance Heartbeat**.

   The Health Policy Configuration - Appliance Heartbeat page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

Keysight, EX-1004
Page 1906

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 1928 for more information.

Configuring Automatic Application Bypass Monitoring

Platform: DC/MDC

Component: None

Supported Platform: All 3D Sensors

Use this module to detect when a detection engine 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.

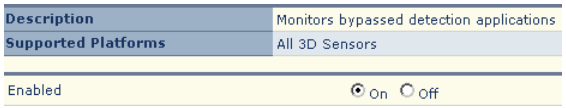
For more information on automatic application bypass, see Automatic Application Bypass on page 229.

To configure automatic application bypass monitoring status:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Automatic Application Bypass Status**.

   The Automatic Application Bypass Status page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

Keysight, EX-1004
Page 1907

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 3D Sensor if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring CPU Temperature Monitoring

Platform: DC/MDC

Component: None

Supported Platform: 3Dx800

The temperature of the central processing unit (CPU) on your 3Dx800 sensor provides an important barometer for the health of your sensor. Overheating a CPU can damage the processing unit. Use the CPU Temperature health status module to set CPU temperature limits.

If the CPU temperature on the monitored sensor exceeds the Warning limit, the status classification for that module changes to Warning. If the CPU temperature on the monitored sensor exceeds the Critical limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

By default, the Critical limit is set to 52 degrees Celsius and the Warning limit is set to 50 degrees Celsius. The maximum temperature you can set for either limit is 100 degrees Celsius, and the Critical limit must be greater than the Warning limit.

WARNING! Sourcefire recommends that you do not set the Critical limit higher than 65 degrees Celsius and that you do not set the Warning limit higher than 55 degrees Celsius.

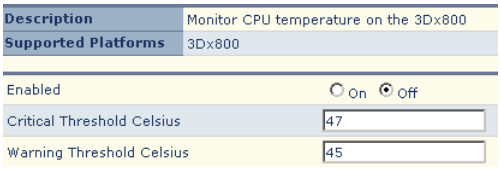
Keysight, EX-1004
Page 1908

To configure CPU temperature health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **CPU Temperature**.

   The Health Policy Configuration - CPU Temperature page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.
3. In the **Critical Threshold Celsius** field, enter the number of degrees, in Celsius, that should trigger a critical health status.
4. In the **Warning Threshold Celsius** field, enter the number of degrees, in Celsius, 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 sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring CPU Usage Monitoring

Platform: DC/MDC

Component: None

Supported Platform: All except 3D9900

Excessive CPU usage can 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.

Keysight, EX-1004
Page 1909

The maximum percentage you can set for either limit is 100 percent, and the Critical limit must be higher than the Warning limit.

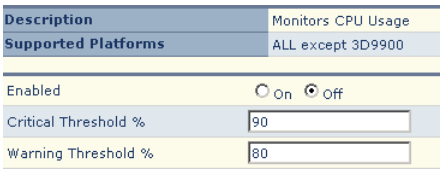
Note that this module is not available for health policies applied to 3D9900 sensors.

To configure CPU Usage health module settings:

Access: Maint/Admin

1. On the Health Policy Configuration page, select **CPU Usage**.

   The Health Policy Configuration - CPU Usage page appears.



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 1928 for more information.

Keysight, EX-1004
Page 1910

Configuring Card Reset Monitoring

Platform: DC/MDC

Component: None

Supported Platform: All 3D Sensors except 3Dx800

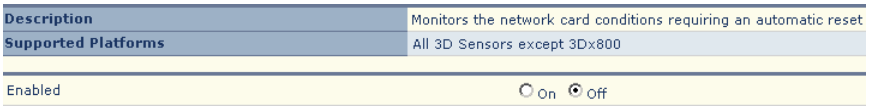
Use the card reset monitoring health status module to track when the network card restarts because of hardware failure. If a reset occurs, this module generates an alert. That status data feeds into the health monitor.

To configure card reset monitoring:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Card Reset**.

   The 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 1928 for more information.

Configuring Disk Usage Monitoring

Platform: DC/MDC

Component: None

Supported Platform: All

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.

Keysight, EX-1004
Page 1911

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.

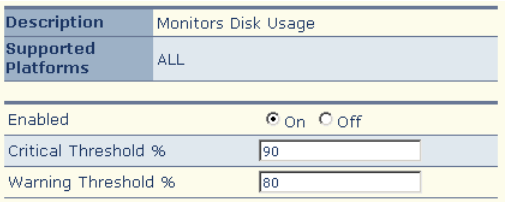
The maximum percentage you can set for either limit is 100 percent, and the Critical limit must be higher than the Warning limit.

To configure Disk Usage health module settings:

Access: Maint/Admin

1. On the Health Policy Configuration page, select **Disk Usage**.

   The Health Policy Configuration - Disk Usage page appears.



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.

Keysight, EX-1004
Page 1912

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 1928 for more information.

Configuring Event Stream Monitoring

Platform: DC/MDC

Component: None

Supported Platform: MDC

Use the Event Stream Status module to monitor the health of the event stream process on a Defense Center by generating alerts when too many seconds elapse between events received by the Master 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 600, and the Critical limit must be higher than the Warning limit. The minimum number of seconds is 300.

Note that if you do not configure any events under eStreamer on your Master Defense Center and you enable this module, the status for the module turns yellow and the event contains a message indicating that events are not configured.

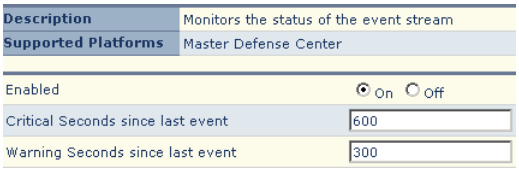
Keysight, EX-1004
Page 1913

To configure Event Stream Status health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Event Stream Status**.

   The Health Policy Configuration - Event Stream Status page appears.



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 Master Defense Center for your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Fan Monitoring

Platform: DC/MDC

Component: None

Supported Platform: 3Dx800

Use the Fan Alarm health status module to warn of fan failure on a 3Dx800 sensor. If the Fan Alarm module finds a fan that has failed, the status classification for that module changes to Critical. That status data feeds into the health monitor.

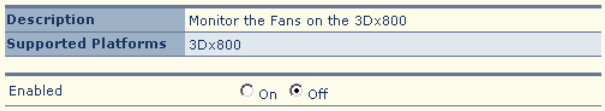
Keysight, EX-1004
Page 1914

To configure Fan Alarm health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select Fan Alarm.

   The Health Policy Configuration - Fan Alarm 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 appropriate sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Hardware Monitoring

Platform: DC/MDC

Component: None

Supported Platform: 3Dx800, 3D8140, 3D8250, 3D9900

Use the Hardware Alarm health status module to detect hardware failure on a 3Dx800, 3D8140, 3D8250, or 3D9900 sensor. If the Hardware Alarm module finds a hardware component that has failed, the status classification for that module changes to Critical. That status data feeds into the health monitor.

Note that the Hardware Alarm module can be used in addition to the Power Supply module to monitor power supply health on the 3Dx800 sensor models.

For more information on the hardware status conditions that can cause hardware alerts on 3D9900 sensors, see Interpreting Hardware Alert Details for 3D9900 Sensors on page 1962.

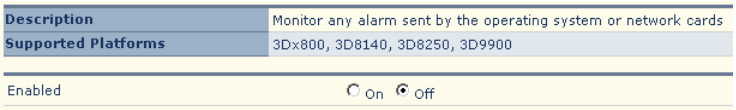
Keysight, EX-1004
Page 1915

To configure Hardware Alarm health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select Hardware Alarms.

   The Health Policy Configuration - Hardware Alarm 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 appropriate sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Health Status Monitoring

Platform: DC/MDC

Component: None

Supported Platform: DC

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 sensor (`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

Keysight, EX-1004
Page 1916

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.

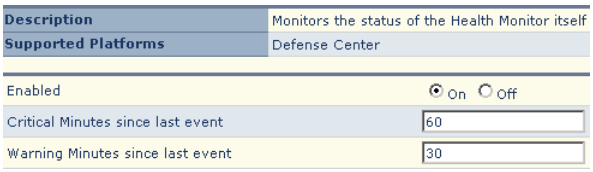
To configure Health Monitor Process module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Health Monitor Process**.

   The Health Policy Configuration - Health Monitor Process page appears.

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 1928 for more information.

Configuring IPS Event Rate Monitoring

Platform: DC/MDC

Component: IPS

Supported Platform: IPS

Use the IPS 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 for

Keysight, EX-1004
Page 1917

the IPS process on the monitored sensor 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 sensor (**Operations > 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

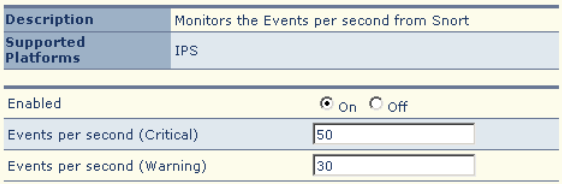
The maximum number of events you can set for either limit is 999, and the Critical limit must be higher than the Warning limit.

To configure IPS Event Rate Monitor health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **IPS Event Rate**.

   The Health Policy Configuration - IPS Event Rate page appears.



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.

Keysight, EX-1004
Page 1918

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 sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Link State Propagation Monitoring

Platform: DC/MDC

Component: None

Supported Platform: IPS

Use the Link State Propagation health status module to detect the interface link state propagation status on an inline interface pair. If a link state propagates to the paired interface, the status classification for that module changes to Critical and the state reads:

`Module Link State Propagation: ethx_ethy is Triggered`

where *x* and *y* are the paired interface numbers.

To configure Link State Propagation health module settings:

Access: Maint/Admin

1. On the Health Policy Configuration page, select **Link State Propagation**.

   The Health Policy Configuration - Link State Propagation monitor page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

Keysight, EX-1004
Page 1919

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 sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Memory Usage Monitoring

Platform: DC/MDC

Component: None

Supported Platform: All

Use the Memory Usage health status module to set memory usage limits. The module calculates free memory by adding free memory and cached 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.

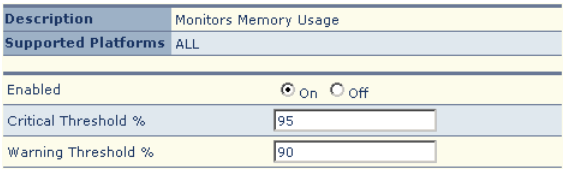
The maximum percentage you can set for either limit is 100 percent, and the Critical limit must be higher than the Warning limit.

To configure Memory Usage health module settings:

Access: Maint/Admin

1. On the Health Policy Configuration page, select **Memory Usage**.

   The Health Policy Configuration - Memory Usage page appears.



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.

Keysight, EX-1004
Page 1920

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 1928 for more information.

Configuring PEP Status Monitoring

Platform: DC/MDC

Component: None

Supported Platform: 3D8140, 3D8250, 3D9900

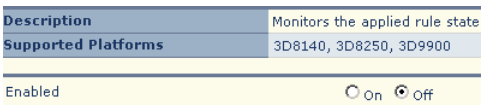
Use the PEP Status health status module to monitor the application of PEP rules to interface sets on a 3D8140, 3D8250, or 3D9900. If PEP rules cannot be applied to interfaces in an interface set, this module generates an alert. That status data feeds into the health monitor.

To configure PEP Status health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **PEP Status**.

   The Health Policy Configuration - PEP Status monitor page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

Keysight, EX-1004
Page 1921

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 sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Power Supply Monitoring

Platform: DC/MDC

Component: None

Supported Platform: DC750, MDC750, DC1500, MDC 1500, DC3000, MDC3000, DC3500, 3D3500, MDC 3500, 3Dx800, 3D4500, 3D6500, 3D8140, 3D8250, 3D9900

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.

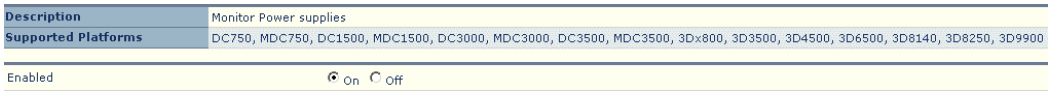
Note that the Hardware Alarm module can be used in addition to the Power Supply module to monitor power supply health on the 3Dx800 sensor models.

To configure Power Supply health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Power Supply**.

   The Health Policy Configuration - Power Supply monitor page appears.



2. Select **On** for the **Enabled** option to enable use of the module for health status testing.

Keysight, EX-1004
Page 1922

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 sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Process Status Monitoring

Platform: DC/MDC

Component: Any

Supported Platform: 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.

This module takes the place of the Data Correlator Process, Defense Center Status, eStreamer Status, IPS Process, and RNA Process modules. If you apply a health policy with Process Status enabled to an appliance running a version earlier than 4.10, the legacy modules are enabled and those modules report health status events for that appliance in place of the Process Status module.

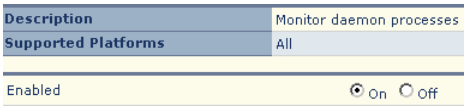
Keysight, EX-1004
Page 1923

To configure Process Status health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Process Status**.

   The Health Policy Configuration - Process 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 appliances if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring RNA Event Status Monitoring

Platform: DC/MDC

Component: RNA

Supported Platform: DC

Use the RNA Event Status module to monitor the health of the RNA process on a sensor from the Defense Center by generating alerts when too many seconds elapse between RNA 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.

Keysight, EX-1004
Page 1924

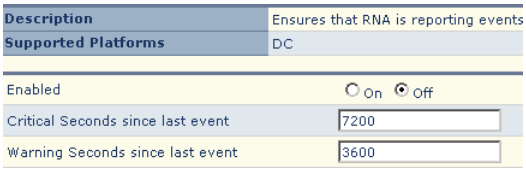
Note that the RNA Health module was renamed to the RNA Event Status module in 4.9.1 and that the supported platforms changed from 3D Sensor to Defense Center in 4.9.1.

To configure RNA Event Status module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **RNA Event Status**.

   The Health Policy Configuration - RNA Event Status page appears.



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 1928 for more information.

Configuring RNA Host Usage Monitoring

Platform: DC/MDC

Component: RNA

Supported Platform: RNA

Use the RNA Host License Limit health status module to set RNA Host shortage limits. If the number of remaining RNA Hosts on the monitored sensor falls below the Warning Hosts limit, the status classification for that module changes to Warning. If the number of remaining RNA Hosts on the monitored sensor falls

Keysight, EX-1004
Page 1925

below the Critical Hosts limit, the status classification for that module changes to Critical. That status data feeds into the health monitor.

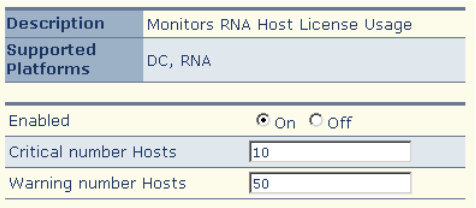
The maximum number of hosts you can set for either limit is 999, and the Critical limit must be higher than the Warning limit.

To configure RNA Host License Limit health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **RNA Host License Limit**.

   The Health Policy Configuration - RNA Host License Limit page appears.



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 sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Keysight, EX-1004
Page 1926

Configuring Time Synchronization Monitoring

Platform: DC/MDC

Component: None

Supported Platform: DC

Use the Time Synchronization Status module to detect when the time on a managed sensor 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: Maint/Admin

1. In the Health Policy Configuration page, select **Time Synchronization Status**.

   The Health Policy Configuration - Time Synchronization 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 appropriate sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Configuring Traffic Status Monitoring

Platform: DC/MDC

Component: IPS/RNA

Supported Platform: IPS, RNA

Use the Traffic Status health status module to detect whether a sensor receives traffic. If the Traffic Status module determines that a sensor does not receive

Keysight, EX-1004
Page 1927

traffic, the status classification for that module changes to Critical. That status data feeds into the health monitor.

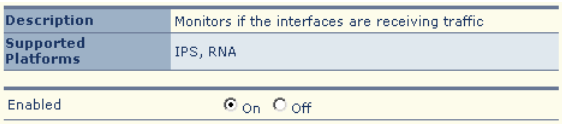
WARNING! If you enable the Traffic Status module on a sensor where there are unused interfaces that are included in an interface set associated with a detection engine, the module interprets the idleness of the port as a traffic failure and alerts on traffic status. To prevent alerting on idle interfaces, remove those interfaces from all interface sets associated with detection engines. For more information on managing interface sets, see Editing an Interface Set on page 239.

To configure Traffic Status health module settings:

Access: Maint/Admin

1. In the Health Policy Configuration page, select **Traffic Status**.

   The Health Policy Configuration - Traffic 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 appropriate sensors if you want your settings to take effect. See Applying Health Policies on page 1928 for more information.

Applying Health Policies

Platform: DC/MDC

Component: None

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

Keysight, EX-1004
Page 1928

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! Default health policies are not replicated between Defense Centers in a high availability pair. Each appliance uses the local default health policy configured for that appliance.

To apply a health policy:

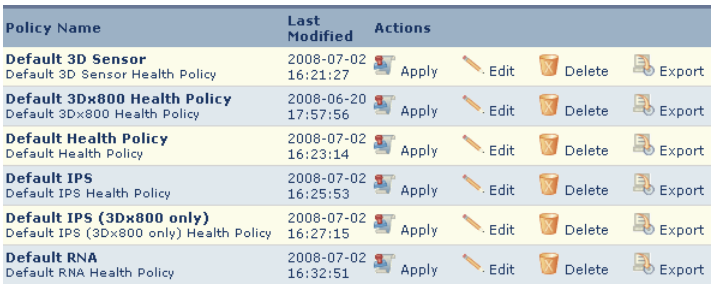
Access: Maint/Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

2. Click **Health Policy** in the health monitor toolbar.

   The Health Policy page appears.

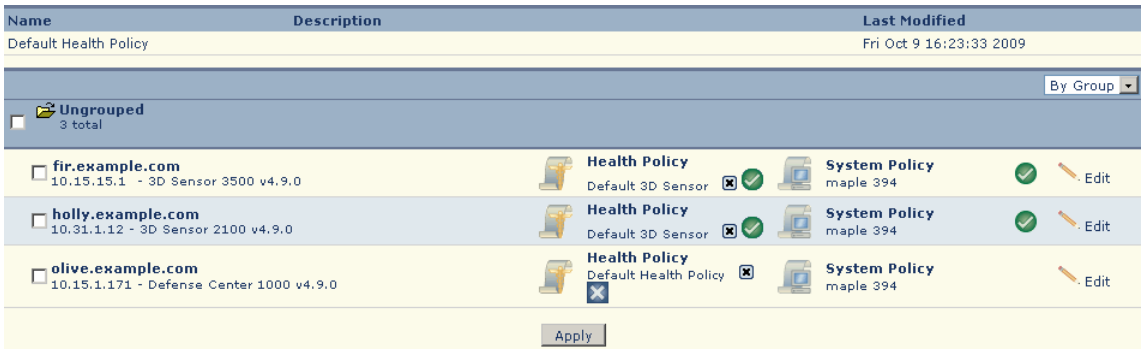


Policy Name	Last Modified	Actions
Default 3D Sensor Default 3D Sensor Health Policy	2008-07-02 16:21:27	Apply Edit Delete Export
Default 3Dx800 Health Policy Default 3Dx800 Health Policy	2008-06-20 17:57:56	Apply Edit Delete Export
Default Health Policy Default Health Policy	2008-07-02 16:23:14	Apply Edit Delete Export
Default IPS Default IPS Health Policy	2008-07-02 16:25:53	Apply Edit Delete Export
Default IPS (3Dx800 only) Default IPS (3Dx800 only) Health Policy	2008-07-02 16:27:15	Apply Edit Delete Export
Default RNA Default RNA Health Policy	2008-07-02 16:32:51	Apply Edit Delete Export

Keysight, EX-1004
Page 1929

3. Click **Apply** next to the policy you want to apply.

   The Health Policy Apply page appears.



> **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 sensor. Note that you can remove the currently applied policy by clicking the remove icon (✖).

4. Check the appliances where you want to apply the health policy.
5. 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:

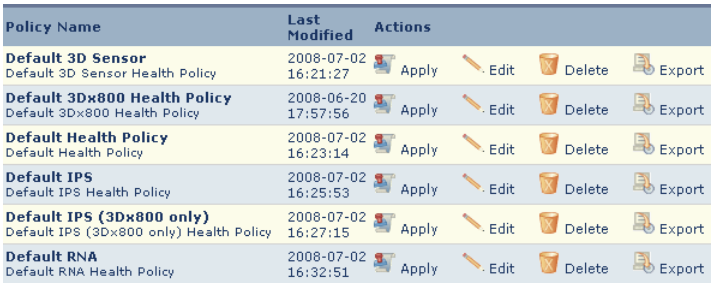
Access: Maint/Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

2. Click **Health Policy** in the health monitor toolbar.

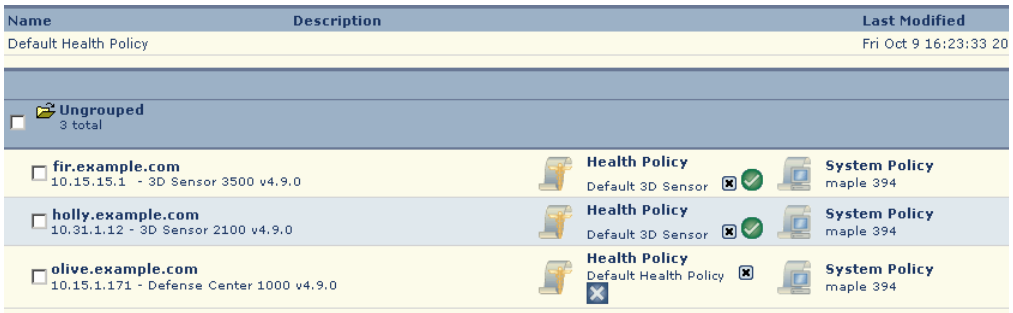
   The Health Policy page appears.



Keysight, EX-1004
Page 1930

3. Click **Apply** next to the policy you want to apply.

   The Health Policy Apply page appears.



4. 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

Platform: DC/MDC

Component: None

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 modules for various appliances are listed in the Health Modules Applicable to Appliances table.

Health Modules Applicable to Appliances

Module	Applicable Appliance
Appliance Heartbeat	Defense Center
Automatic Application Bypass Status	All 3D Sensors
CPU Temperature	3Dx800 Only
CPU Usage	All except 3D9900
Card Reset	All 3D Sensors except 3Dx800
Disk Usage	All

Keysight, EX-1004
Page 1931

Health Modules Applicable to Appliances (Continued)

Module	Applicable Appliance
Event Stream Status	Master Defense Center
Fan Alarm	3Dx800
Hardware Alarms	3Dx800, 3D8140, 3D8250, and 3D9900
Health Monitor Process	Defense Center
IPS Event Rate	3D Sensors with IPS
Link State Propagation	3D Sensors with IPS
Memory Usage	All
PEP Status	3D8140, 3D8250, 3D9900
Power Supply	DC750, MDC750, DC1500, MDC1500, DC3000, MDC3000, DC3500, MDC3500, 3Dx800, 3D3500, 3D4500, 3D6500, 3D8140, 3D8250, and 3D9900
Process Status	All
RNA Health	Defense Center
RNA Host License Limit	Defense Center
Time Synchronization Status	Defense Center
Traffic Status	3D Sensors with IPS, 3D Sensors with RNA

To edit a health policy:

Access: Maint/Admin

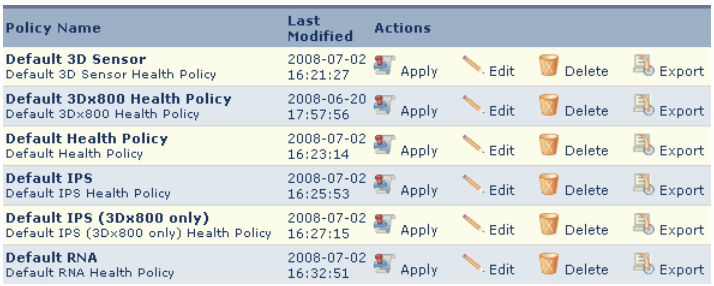
1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

Keysight, EX-1004
Page 1932

2. Click **Health Policy** in the health monitor toolbar.

   The Health Policy page appears.



Policy Name	Last Modified	Actions
Default 3D Sensor Default 3D Sensor Health Policy	2008-07-02 16:21:27	Apply Edit Delete Export
Default 3Dx800 Health Policy Default 3Dx800 Health Policy	2008-06-20 17:57:56	Apply Edit Delete Export
Default Health Policy Default Health Policy	2008-07-02 16:23:14	Apply Edit Delete Export
Default IPS Default IPS Health Policy	2008-07-02 16:25:53	Apply Edit Delete Export
Default IPS (3Dx800 only) Default IPS (3Dx800 only) Health Policy	2008-07-02 16:27:15	Apply Edit Delete Export
Default RNA Default RNA Health Policy	2008-07-02 16:32:51	Apply Edit Delete Export

3. Click **Edit** next to the policy you want to modify.

   The Health Policy Configuration page appears, with the Policy Run Time Interval settings selected.

4. Modify settings as needed, as described in the following sections:

   - Configuring Policy Run Time Intervals on page 1905
   - Configuring Appliance Heartbeat Monitoring on page 1906
   - Configuring Automatic Application Bypass Monitoring on page 1907
   - Configuring CPU Temperature Monitoring on page 1908
   - Configuring CPU Usage Monitoring on page 1909
   - Configuring Card Reset Monitoring on page 1911
   - Configuring Disk Usage Monitoring on page 1911
   - Configuring Event Stream Monitoring on page 1913
   - Configuring Fan Monitoring on page 1914
   - Configuring Hardware Monitoring on page 1915
   - Configuring Health Status Monitoring
   - Configuring IPS Event Rate Monitoring on page 1917
   - Configuring Link State Propagation Monitoring on page 1919
   - Configuring Memory Usage Monitoring on page 1920
   - Configuring PEP Status Monitoring on page 1921
   - Configuring Power Supply Monitoring on page 1922
   - Configuring Process Status Monitoring on page 1923
   - Configuring RNA Event Status Monitoring on page 1924
   - Configuring RNA Host Usage Monitoring on page 1925
   - Configuring Time Synchronization Monitoring on page 1927
   - Configuring Traffic Status Monitoring on page 1927

Keysight, EX-1004
Page 1933

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.
6. Reapply the policy to the appropriate appliances as described in Applying Health Policies on page 1928.

Deleting Health Policies

Platform: DC/MDC

Component: None

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 sensor, any health monitoring alerts in effect for the sensor remain active until you deactivate the underlying associated alert. For more information on deactivating alerts, see Activating and Deactivating Alerts on page 1293.

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 1902. For more information on applying health policies, see Applying Health Policies on page 1928.

To delete a health policy:

Access: Maint/Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.
2. Click **Health Policy** in the health monitor toolbar.

   The Health Policy page appears.
3. Click **Delete** next to the policy you want to delete.

   A message appears, indicating if the deletion was successful.

Keysight, EX-1004
Page 1934

Using the Health Monitor Blacklist

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 or Master Defense Center.

You can use the health monitor blacklist feature to disable health monitoring status reporting on an appliance, module, or detection engine. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a 3D Sensor 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 3D Sensor.

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, module, or detection engine 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 appliance no longer includes the 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 or detection engine. For example, when you run out of RNA host licenses on an appliance, you can blacklist the RNA Host License Limit status messages until you install a new license with more hosts.

Make sure to remove all unused sensing interfaces from any interface sets in use by a detection engine so health monitoring alerts do not generate for those interfaces.

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 1945.

A blacklist icon () and a notation are visible once you expand the view for a blacklisted or partially blacklisted appliance.

IMPORTANT! On a Defense Center, Health Monitor blacklist settings are system settings. Therefore if you blacklist a sensor, then delete it and later re-register it with the Defense Center, the blacklist settings remain persistent. The newly re-registered sensor remains blacklisted.

Keysight, EX-1004
Page 1935

Blacklisting Health Policies or Appliances

Platform: DC/MDC

Component: None

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. Once 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 sensor on one HA peer and not the other. You can also blacklist the HA peer to cause it to mark events generated by it and the sensors from which it receives health events as disabled.

TIP! You can blacklist 3D Sensors only from a Defense Center, not a Master Defense Center. You cannot blacklist intrusion agents.

To blacklist an entire health policy or group of appliances:

Access: Maint/Admin

1. Select **Operations > Monitoring > Health**.
   The Health Monitor page appears.
2. On the toolbar, click **Blacklist**.
   The Blacklist page appears.

Keysight, EX-1004
Page 1936

3. Use the drop-down list on the right to sort the list by group, policy, or model. (On a Master Defense Center, sort the list by group, manager, policy or model. Groups on a Defense Center are 3D Sensors. Groups on a Master Defense Center are appliances.)



> **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 sensor. Note that you can remove the currently applied policy by clicking the remove icon ( ).

4. To blacklist all appliances in a group, model, or policy category, select the category then click **Apply**. (On a Master Defense Center, to blacklist all appliances associated with a manager, select the manager then click **Apply**.)

   The page refreshes, now indicating the blacklisted state of the appliances.

Blacklisting an Appliance

If you need to set the events and health status for an individual appliance to disabled, you can blacklist the appliance. Once 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: Maint/Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

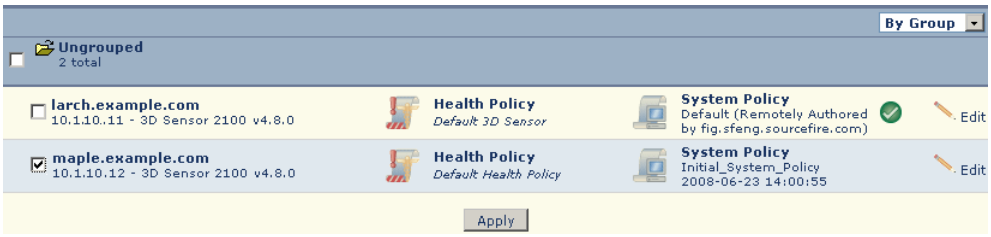
2. On the toolbar, click **Blacklist**.

   The Blacklist page appears.

3. Use the drop-down list on the right to sort the list by appliance group, model, or by policy. (On a Master Defense Center, sort the list by group, manager, policy or model.)

Keysight, EX-1004
Page 1937

4. To blacklist an individual appliance, select and expand a category folder, select the box next to the appropriate appliance, then click **Apply**.



The page refreshes then indicates the blacklisted state of the appliances. Click **Edit** and see Blacklisting a Health Policy Module on page 1938 to blacklist individual health policy modules.

Blacklisting a Health Policy Module

Platform: DC/MDC

Component: None

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.

For some modules, you can blacklist that module for a specific detection engine. For example, if you know you are going to disable the RNA detection engine on a sensor and do not want traffic status alerts to change the status for the sensor, you can blacklist the Traffic Status module for that detection engine.

Note that modules that allow you to select a specific detection engine have an arrow next to the module. When any part of a module is blacklisted, the line for that module appears in boldface type in the Defense Center web interface. In addition, the interface indicates the following information in parentheses after each module with detection engines: `number of blacklisted detection engines/maximum number of detection engines`.

Defense Center Only

Specific health policy modules operate for a Defense Center. When blacklisting modules for Defense Centers, only include the following modules:

- Appliance Heartbeat
- CPU Usage
- Disk Usage
- Health Monitor Process
- Memory Usage
- Time Synchronization Status
- Power Supply

Keysight, EX-1004
Page 1938

- Process Status
- RNA Host License Limit

Master Defense Center Only

Specific health policy modules operate for a Master Defense Center. When blacklisting modules for Master Defense Centers, only include the following modules:

- CPU Usage
- Defense Center Status
- Disk Usage
- Event Stream Status
- Memory Usage
- Power Supply
- Process Status

For details about applicable modules on all appliances, see the Health Modules Applicable to Appliances table on page 1931.

TIP! Once the blacklist settings take effect, the appliance shows as `Part Blacklisted` or `All Modules Blacklisted` in 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: Maint/Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

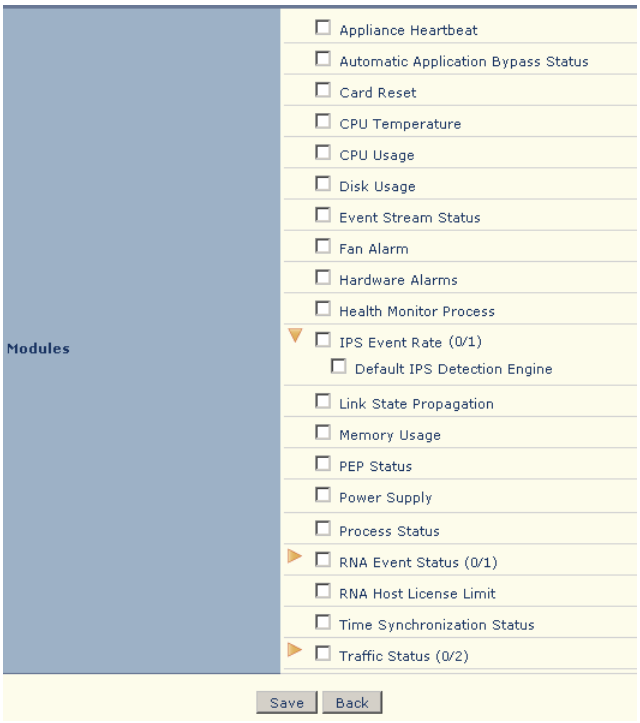
2. On the toolbar, click **Blacklist**.

   The Blacklist page appears.

Keysight, EX-1004
Page 1939

3. 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.



4. You have two options:
   - Select each module that you want to blacklist.
   - Expand the detection engine list by clicking on the arrow next to modules with detection engine lists, then select each detection engine for which you want to blacklist the module.
5. Click **Save**.

Configuring Health Monitor Alerts

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 with health event levels to cause that alert to trigger 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.

Keysight, EX-1004
Page 1940

For more information, see the following topics:

- Preparing to Create a Health Alert on page 1941
- Creating Health Monitor Alerts on page 1942
- Interpreting Health Monitor Alerts on page 1943
- Editing Health Monitor Alerts on page 1944
- Deleting Health Monitor Alerts on page 1945

Preparing to Create a Health Alert

Platform: DC/MDC

Component: None

If you want to create a health alert, you first need to create the underlying alert that you associate to the health alert. If you want to use email alerting, you also need to set up your email relay host in your system policy and re-apply that policy.

To prepare your system for alerting:

Access: Admin

1. If you plan to use email alerting:
   - Select **Operations > System Policy**.
   - Create a new policy or click **Edit** next to an existing one.
   - In the policy, click **Email Notification**.
   - Enter the name of the **Mail Relay Host**.
   - Click **Save Policy and Exit**.
   - Click **Apply** and apply the policy to the Defense Center where you plan to create the health alert.
2. Create email, SNMP, or syslog alerts you want to associate with health alerts:
   - For more information on creating syslog alerts, see Creating Syslog Alerts on page 1285.
   - For more information on creating email alerts, see Creating Email Alerts on page 1289.
   - For more information on creating SNMP alerts, see Creating SNMP Alerts on page 1290.

   Continue with Creating Health Monitor Alerts on page 1942.

Keysight, EX-1004
Page 1941

Creating Health Monitor Alerts

Platform: DC/MDC

Component: None

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert. You can use an existing alert or configure a new one specifically to report on system health. For more information on creating the alert, see Preparing to Create a Health Alert on page 1941. When the severity level occurs for the selected module, the associated 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:

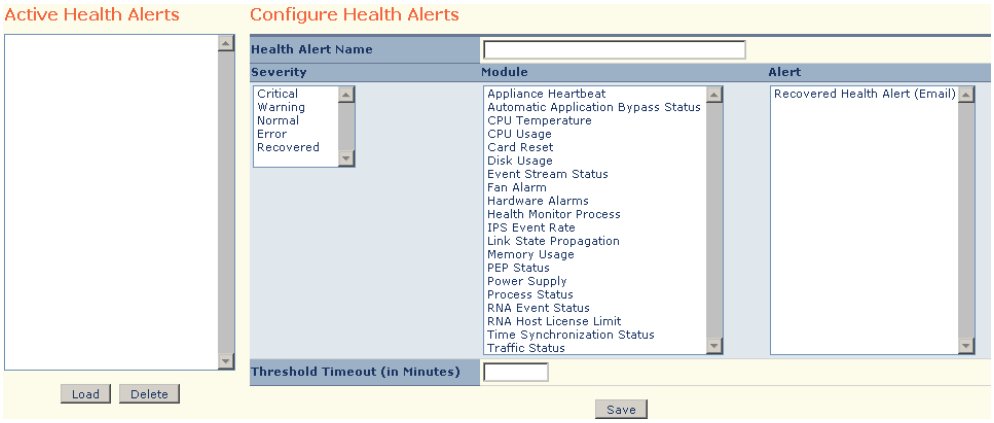
Access: Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

2. Click **Health Monitor Alerts** in the health monitor toolbar.

   The Health Monitor Alerts page appears.



3. Type a name for the health alert in the **Health Alert Name** field.
4. From the **Severity** list, select the severity level you want to use to trigger the alert.

Keysight, EX-1004
Page 1942

5. 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.

6. From the **Alert** list, select the alert which you want to trigger when the selected severity level is reached.

   TIP! Click **Alerts** in the toolbar to open the Alerts page. For more information on creating alerts, see Creating Alerts on page 1285.

7. In the **Threshold Timeout** field, type the number of minutes that should elapse before each threshold period ends and the threshold count resets.
8. 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

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.

Keysight, EX-1004
Page 1943

For more information on health modules, see Understanding Health Modules on page 1893.

Editing Health Monitor Alerts

Platform: DC/MDC

Component: None

You can edit existing health monitor alerts to change the severity level, health module, or alert associated with the health monitor alert.

To edit health monitor alerts:

Access: Admin

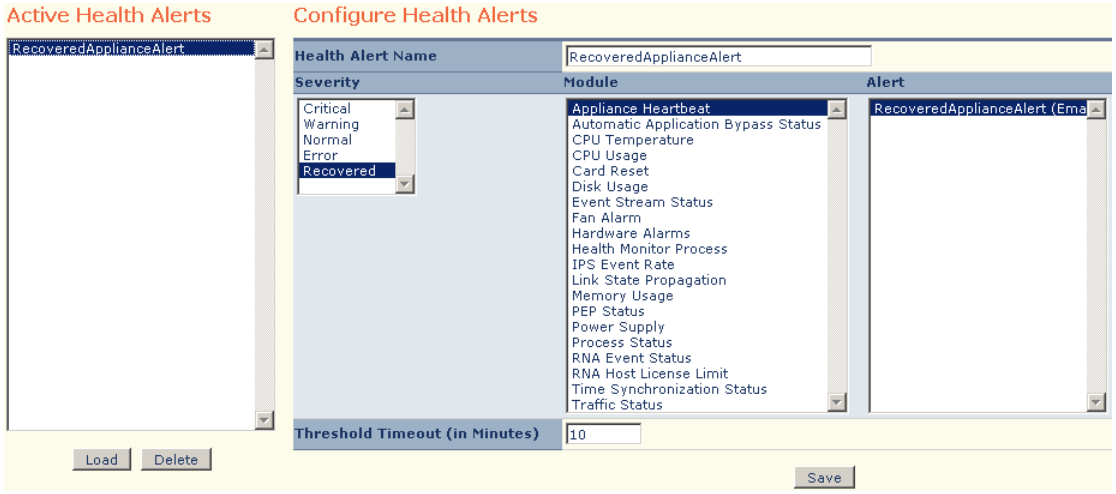
1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

2. Click **Health Monitor Alerts** in the health monitor toolbar.

   The Health Monitor Alerts page appears.

3. Select the alert you want to modify in the **Active Health Alerts** list.



4. Click **Load** to load the configured settings for the selected alert.
5. Modify settings as needed. For more information, see Creating Health Monitor Alerts on page 1942.
6. Click **Save** to save the modified health alert.

   A message appears, indicating if the alert configuration was successfully saved.

Keysight, EX-1004
Page 1944

Deleting Health Monitor Alerts

Platform: DC/MDC

Component: None

You can delete existing health monitor alerts.

IMPORTANT! Deleting a health monitor alert does not delete the associated alert. You must deactivate or delete the underlying alert to ensure that alerting does not continue. For more information on deactivating alerts, see Activating and Deactivating Alerts on page 1293. For more information on deleting alerts, see Deleting Alerts on page 1292.

To delete health monitor alerts:

Access: Admin

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

2. Click **Health Monitor Alerts** in the health monitor toolbar.

   The Health Monitor Alerts page appears.

3. Select the alert you want to delete in the **Active Health Alerts** list.

4. Click **Delete**.

   A message appears, indicating if the alert configuration was successfully deleted.

Using the Health Monitor

Platform: DC/MDC

Component: None

The Health Monitor page provides the compiled health status for all sensors 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.

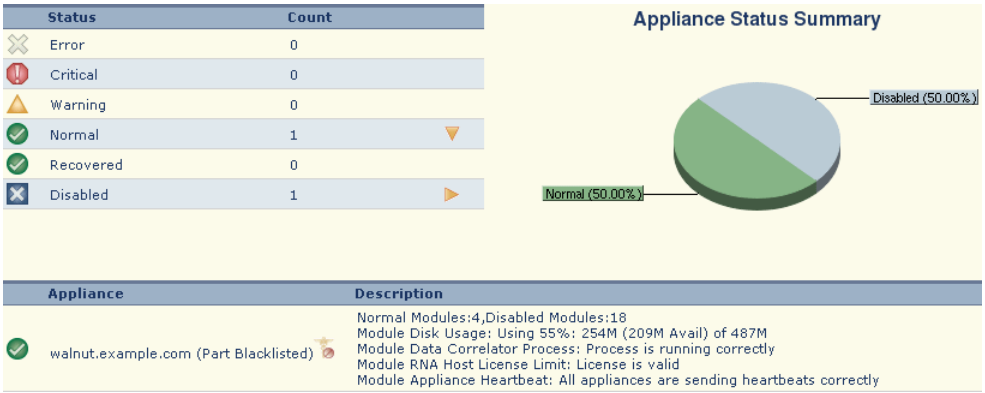
Keysight, EX-1004
Page 1945

To use the health monitor:

Access: Maint/Admin/Any Analyst except Restricted

1. Click **Health Monitor** on the toolbar.

   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 1947
- Using Appliance Health Monitors on page 1947
- Configuring Health Policies on page 1895
- Configuring Health Monitor Alerts on page 1940

Keysight, EX-1004
Page 1946

Interpreting Health Monitor Status

Available status categories, by severity, include Error, Critical, Warning, Normal, 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

Platform: DC/MDC

Component: None

The Appliance health monitor provides a detailed view of the health status of an appliance.

IMPORTANT! Your browser session will not be automatically timed out while you are viewing the Health Monitor page.

Keysight, EX-1004
Page 1947

To view the status summary for a specific appliance:

Access: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

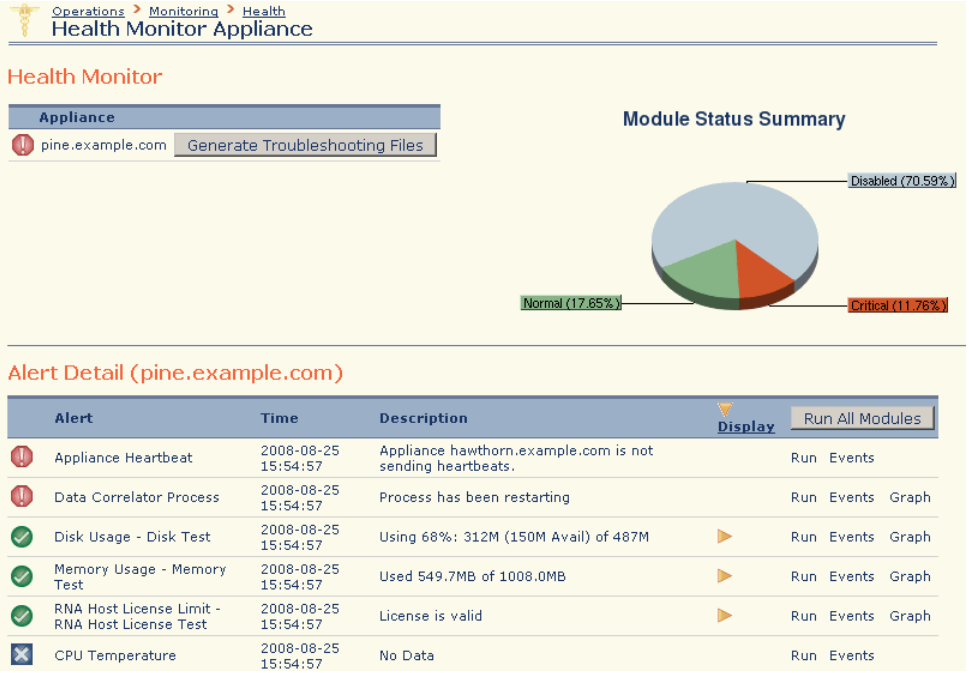
   The Health Monitor page appears.

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.



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:

   - Interpreting Appliance Health Monitor Status on page 1949
   - Viewing Alerts by Status on page 1950

Keysight, EX-1004
Page 1948

- Running All Modules for an Appliance on page 1950
- Running a Specific Health Module on page 1951
- Generating Health Module Alert Graphs on page 1953
- Using the Health Monitor to Troubleshoot on page 1954

Interpreting Appliance Health Monitor Status

Available status categories, by severity, include Error, Critical, Warning, Normal, and Disabled, as described in the Appliance Health Status Indicator table that follows.

Appliance Health Status Indicator

Status Level	Status Icon	Status Color	Description
Error		White	Indicates that the health monitoring module has failed 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 the health module on the appliance and the problem has not been corrected.
Warning		Yellow	Indicates that warning limits have been exceeded for the health module on the appliance and the problem has not been corrected.
Normal		Green	Indicates that the monitored item is running within the limits configured in the health policy applied to the appliance.
Recovered		Green	Indicates that the health for the monitored item is back within the limits configured in the health policy applied to the appliance.
Disabled		Blue	Indicates that a module is disabled or blacklisted, that the appliance does not have a health policy applied to it, or that the appliance is currently unreachable.

Keysight, EX-1004
Page 1949

Viewing Alerts by Status

Platform: DC/MDC

Component: None

You can show or hide categories of alerts by status.

To show alerts by status:

Access: Maint/Admin/Any Analyst except Restricted

- 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: Maint/Admin/Any Analyst except Restricted

- 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

Platform: DC/MDC

Component: None

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: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

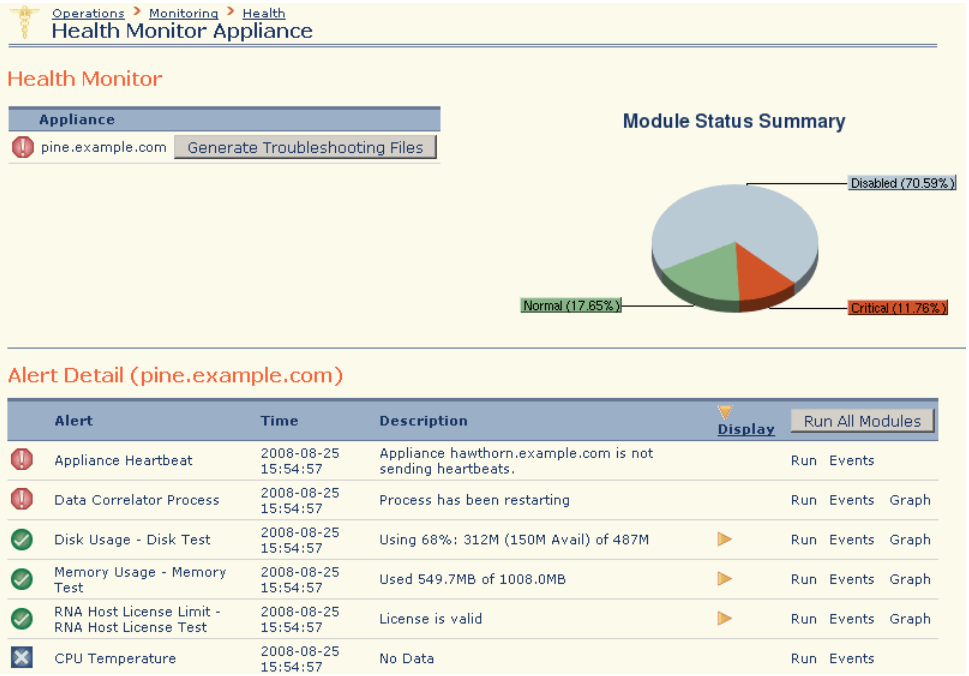
2. To expand the appliance list to show appliances with a particular status, click the arrow in that status row.

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.

Keysight, EX-1004
Page 1950

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.



4. 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 sensor name. You can also wait for the page to refresh again automatically.

Running a Specific Health Module

Platform: DC/MDC

Component: None

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.

Keysight, EX-1004
Page 1951

To run a specific health module:

Access: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   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.

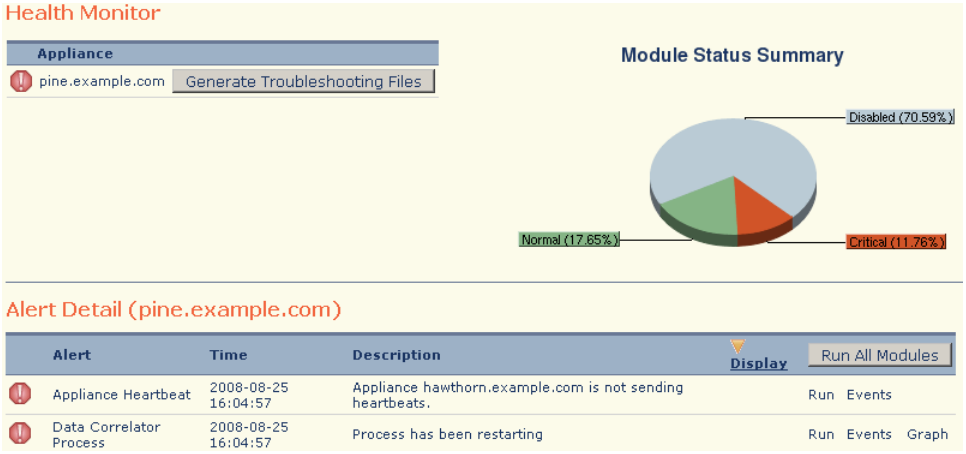
   TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the Appliance column of the appliance list, click the name of the appliance for which you want to view details in the health monitor toolbar.

   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.



Keysight, EX-1004
Page 1952

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 sensor name. You can also wait for the page to refresh automatically again.

Generating Health Module Alert Graphs

Platform: DC/MDC

Component: None

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: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   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 in the health monitor toolbar.

   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.

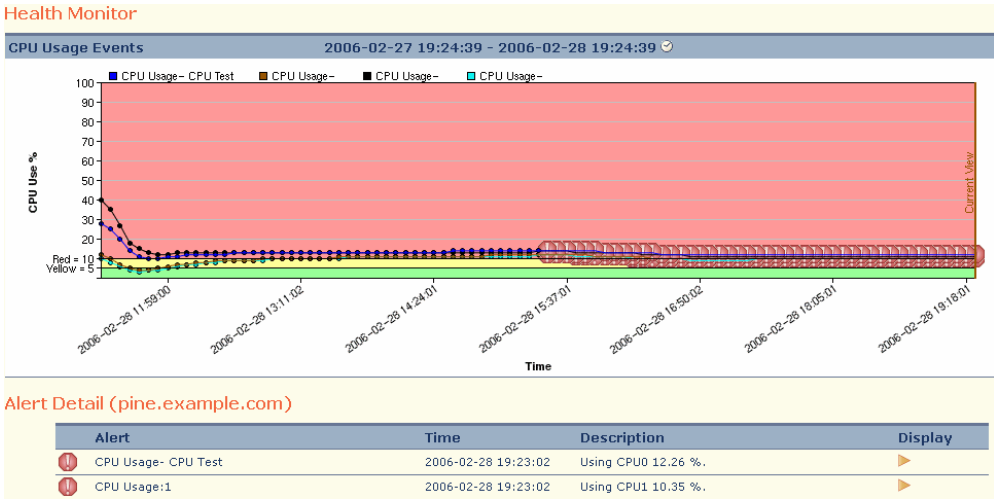
   The Alert Detail list expands to list the health alerts for the selected appliance for that status category.

Keysight, EX-1004
Page 1953

5. 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 1597 for more information.



Using the Health Monitor to Troubleshoot

Platform: DC/MDC

Component: None

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.

Keysight, EX-1004
Page 1954

Selectable Troubleshoot Options

This option...	Reports...
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 Engine Configuration, Policy, and Logs	configuration settings, data, and logs related to detection engines on the appliance
Interface and Network Related Data	configuration settings, data, and logs related to the network interfaces, interface sets, and network configuration of the appliance
RNA, RUA, VDB Data and Logs	configuration settings, data, and logs related to the current RNA and RUA 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

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 1955
- Downloading Troubleshooting Files on page 1957

Generating Appliance Troubleshooting Files

Platform: DC/MDC

Component: None

Use the following procedure to generate customized troubleshooting files that you can send to Sourcefire Support.

Keysight, EX-1004
Page 1955

To generate troubleshooting files:

Access: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   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.

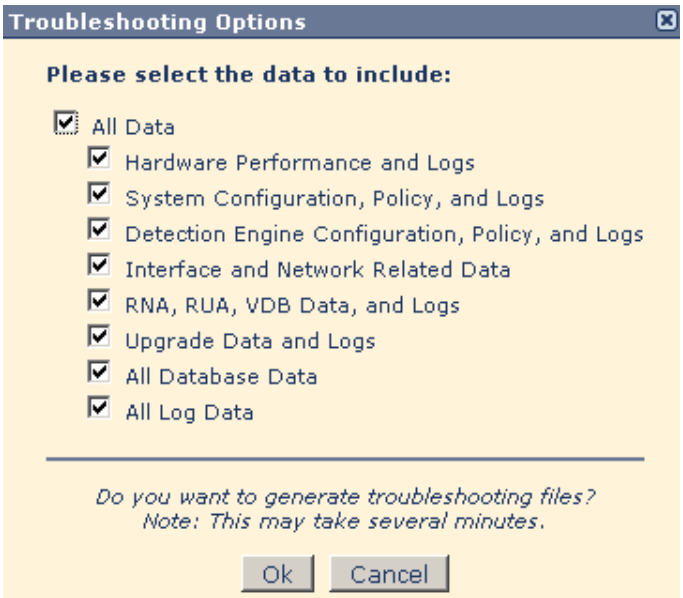
   TIP! If the arrow in the row for a status level points down, the appliance list for that status shows in the lower table. If the arrow points right, the appliance list is hidden.

3. In the Appliance column of the appliance list, click the name of the appliance for which you want to view details in the health monitor toolbar.

   The Health Monitor Appliance page appears.

4. Click **Generate Troubleshooting Files**.

   The Troubleshoot Options pop-up window appears.



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 1955.

6. Click **OK**.

   The Defense Center generates the troubleshooting files. You can monitor the file generation process in the task queue (**Operations > Monitoring > Task Status**).

Keysight, EX-1004
Page 1956

7. Continue with the procedure in the next section, Downloading Troubleshooting Files.

Downloading Troubleshooting Files

Platform: DC/MDC

Component: None

Use the following procedure to download copies of your generated troubleshooting files.

To download troubleshooting files:

Access: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Task Status**.

   The Task Status page appears.

2. Find the task that corresponds to 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

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 1958 for more information about these common procedures.

From the **Operations > Monitoring > Health** menu, you can view health events, and can search for specific events.

Keysight, EX-1004
Page 1957

See the following sections for more information about viewing events:

- Understanding Health Event Views on page 1958 describes the types of events that RNA generates.
- Viewing Health Events on page 1958 describes how to access and use the Event View page.
- Searching for Health Events on page 1965 describes how to search for specific events using the Event Search page.

Understanding Health Event Views

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 1893.

For more information about viewing and searching for health events, see the following sections:

- Viewing Health Events on page 1958
- Understanding the Health Events Table on page 1963
- Searching for Health Events on page 1965

Viewing Health Events

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 1958
- Viewing Health Events by Module and Appliance on page 1959
- Working with the Health Events Table View on page 1961
- Searching for Health Events on page 1965

Viewing All Health Events

Platform: DC/MDC

Component: None

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 1893.

When you access health events from the Health Monitor page on your Defense Center, you retrieve all health events for all managed appliances.

Keysight, EX-1004
Page 1958

To view all health events on all managed appliances:

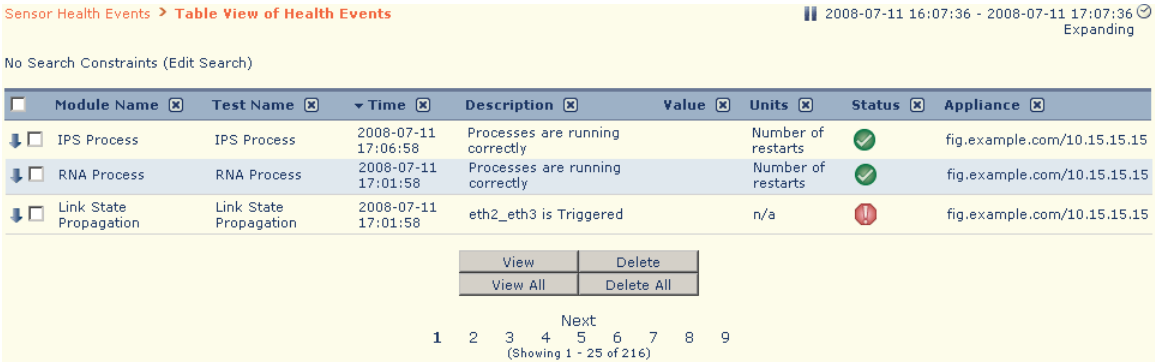
Access: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

2. In the toolbar, click **Health Events**.

   The Events page appears, containing all health events.



IMPORTANT! If no events appear, you may need to adjust the time range. See Setting Event Time Constraints on page 1597 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 1597.

Viewing Health Events by Module and Appliance

Platform: DC/MDC

Component: None

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: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

Keysight, EX-1004
Page 1959

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 in the health monitor toolbar.

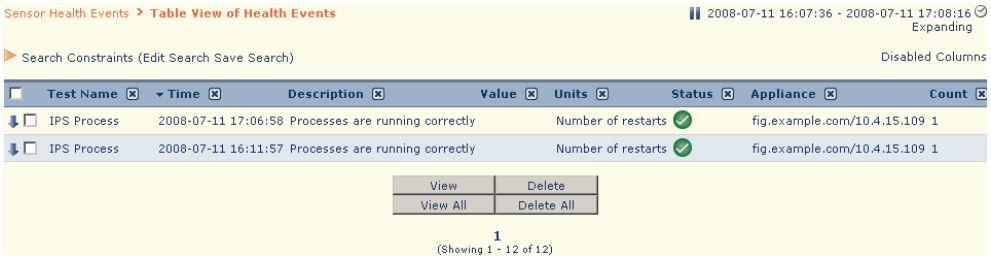
   The Health Monitor Appliance page appears.

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.



   If no events appear, you may need to adjust the time range. See Setting Event Time Constraints on page 1597 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.

Keysight, EX-1004
Page 1960

Working with the Health Events Table View

Platform: DC/MDC

Component: None

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 1963.
modify the time and date range for events listed in the Health table view	find more information in Setting Event Time Constraints on page 1597. 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 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 1612.
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 1613.
navigate to other event tables to view associated events	find more information in Navigating between Workflows on page 1614.
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 1615 for more information.
navigate to the bookmark management page	select **Analysis & Reporting > Bookmarks** or, from any event view, click **View Bookmarks**. See Using Bookmarks on page 1615 for more information.
generate a report based on data in the table view	click **Report Designer**. See Generating a Basic Report on page 1532 for more information.

Keysight, EX-1004
Page 1961

Health Event View Functions (Continued)

To...	You can...
select another health events workflow	click **Workflows** or select from the **Workflows** drop-down list in the toolbar. See Selecting Workflows on page 1590 for more information.
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 Sensors

For 3D9900 sensor models, hardware alarms generate in response to the events described in the Conditions Monitored for 3D9900 Sensors table. The triggering condition can be found in the message detail for the alert.

Conditions Monitored for 3D9900 Sensors

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	• 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.

Keysight, EX-1004
Page 1962

Conditions Monitored for 3D9900 Sensors (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.
`Psls` daemon	If the `Psls` 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.

Understanding the Health Events Table

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 1869.

Keysight, EX-1004
Page 1963

The fields in the health events table are described in the Health Event Fields table.

Health Event Fields

Field	Description
Module Name	The name of the health module that generated the event. For a list of health modules, see the Health Modules table on page 1893.
Test Name	The name of the test. This is typically the same as the module name.
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 sensor 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 sensor 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.
Sensor	The appliance where the health event was reported.

To display the table view of health events:

Access: Maint/Admin/Any Analyst except Restricted

1. Select **Operations > Monitoring > Health**.

   The Health Monitor page appears.

Keysight, EX-1004
Page 1964

2. On the toolbar, click **Health Events**.

   The table view appears. For information on working with health events, see Working with Health Events on page 1957.

> **TIP!** If you are using a custom workflow that does not include the table view of health events, click **Workflows**. On the Select Workflow page, click **Health Events**.

Searching for Health Events

Platform: DC/MDC

Component: None

You can use Event Search to search for specific network discovery events. You can create, save, and re-use event searches. When creating new searches or modifying default searches, there are a number of options you can configure. The Health Event Search Criteria table describes each search criterion you can specify.

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.

Keysight, EX-1004
Page 1965

Health Event Search Criteria (Continued)

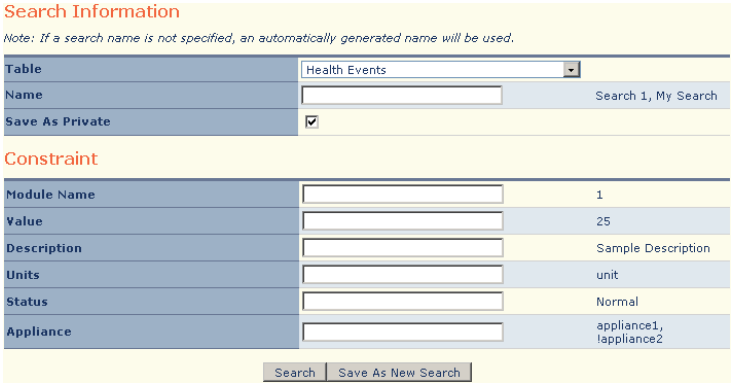
Search Field	Description
Units	Specify the units descriptor for the result obtained by the health test for the events you want to view. You can use an asterisk (*) in this field to create wildcard searches. For example, if you type `%` in the Units field, you retrieve all events for the Disk Usage modules, 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.
Status	Specify the status for the health events that you want to view. Valid status levels are Critical, Warning, Normal, Error, and Disabled. For example, type `Critical` to retrieve all health events that indicate a critical status.
Appliance	Specify the name of appliance.

To run and save health event searches:

Access: Any Analyst except Restricted/Admin

1. Select **Analysis & Reporting > Searches > Health Events**.

   The Search page appears.



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 created automatically when you save the search.

3. Enter your search criteria.

   See Health Event Search Criteria on page 1965 for more information about the values you can enter for search criteria.

Keysight, EX-1004
Page 1966

4. 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 save a search as a restriction for restricted data users, you **must** save it as a private search.

5. You have the following options:
   - Click **Search** to execute 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, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.
   - 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 1554
   - Deleting a Saved Search on page 1554

Keysight, EX-1004
Page 1967

Chapter 45
Auditing the System

You can audit activity on your system in two ways. The appliances that are a 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 1969 describes how to view and manage system audit information.
- Viewing the System Log on page 1981 describes how to view the system log, which contains system status messages.

TIP! Defense Centers and 3D Sensors with IPS also provide full-featured 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 Event Reports on page 1531.

Keysight, EX-1004
Page 1968

Managing Audit Records

Platform: DC/MDC/3D Sensor

Component: None

Defense Centers and 3D Sensors 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.

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.

For more information, see the following sections:

- Viewing Audit Records on page 1969
- Suppressing Audit Records on page 1973
- Understanding the Audit Log Table on page 1977
- Searching Audit Records on page 1978

Viewing Audit Records

Platform: DC/MDC/3D Sensor

Component: None

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 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 1618.

Keysight, EX-1004
Page 1969

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 1977.
modify the time range used when viewing audit records	find more information at Setting Event Time Constraints on page 1597. 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 events on the current workflow page	find more information in Sorting Table View Pages and Changing Their Layout on page 1611.
navigate within the current workflow page	find more information in Navigating to Other Pages in the Workflow on page 1613.
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 1593.
drill down to the next page in the workflow	use one of the following methods: • 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 checkboxes 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**. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.

Keysight, EX-1004
Page 1970

Audit Log Actions (Continued)

To...	You can...
constraining on a specific value	Click a value within a row. If you click a value on a drilldown page, you move to the next page and constrain on the value. 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. **TIP!** Table views always include "Table View" in the page name. For more information, see Constraining Events on page 1606.
delete audit records	use one of the following methods: • 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 **Workflows**. For more information, see Selecting Workflows on page 1590.
bookmark the current page so that you can quickly return to it	click **Bookmark This Page**. For more information, see Using Bookmarks on page 1615.
navigate to the bookmark management page	click **View Bookmarks**. For more information, see Using Bookmarks on page 1615.
generate a report based on the data in the current view	click **Report Designer**. For more information, see Generating a Basic Report on page 1532.

To view audit records:

Access: Admin

▶ Select **Operations > Monitoring > Audit**.

The first page of the default audit log workflow appears. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45. If no events appear, you may need to adjust the time range. For more information, see Setting Event Time Constraints on page 1597.

TIP! If you are using a custom workflow that does not include the table view of audit events, from the **Workflows** menu on the toolbar, select **Audit Log**.

Keysight, EX-1004
Page 1971

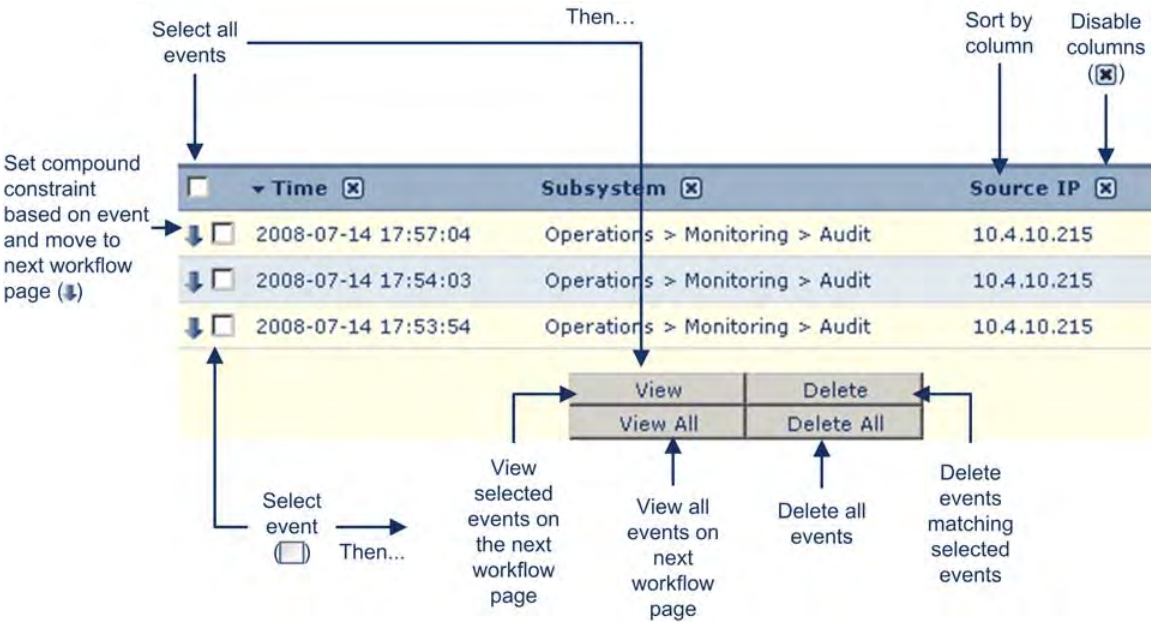
Working with Audit Events

Platform: DC/MDC/3D Sensor

Component: None

You can change the layout of the event view or constrain the events in the view by a field value.

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, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, use the Expand arrow (▶) to expand the search constraints, then click the column name under Disabled Columns.

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 1606.
- Using Compound Constraints on page 1609

Keysight, EX-1004
Page 1972

- Sorting Drill-down Workflow Pages on page 1612
- Understanding the Audit Log Table on page 1977

Suppressing Audit Records

Platform: DC/MDC/3D Sensor

Component: None

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 set up the suppression mechanism, 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.

Keysight, EX-1004
Page 1973

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 `AuditBlock.address` 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 `10.1.1` matches addresses from `10.1.1.0` through `10.1.1.255`.
Message	Create a file named `AuditBlock.message` and include, one per line, the message substrings that you want to suppress. Note that substrings are matched so that if you include `backup` in your file, all messages that include the word `backup` are suppressed.
Subsystem	Create a file named `AuditBlock.subsystem` and include, one per line, each subsystem that you want to suppress. Note that substrings are **not** matched. You must use exact strings. See the Subsystem Names table for a list of subsystems that are audited.
User	Create a file named `AuditBlock.user` 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 `IPSAnalyst` matches the usernames `IPSAnalyst1` and `IPSAnalyst2`.

IMPORTANT! 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.

Keysight, EX-1004
Page 1974

Subsystem Names

Name	Includes user interactions with...
Admin	Administrative features such as system and access configuration, time synchronization, back up and restore, sensor 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
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
Error	System-level errors
eStreamer	eStreamer configuration
EULA	Reviewing the end user license agreement
Events	Intrusion and RNA event views
Events Clipboard	Intrusion event clipboard
Events Reviewed	Reviewed intrusion events
Events Search	Any event search
Failed to install SEU *seu_id*	Installing SEUs
Header	Initial presentation of the user interface after a user logs in
Health	Health monitoring
Health Events	Health monitoring event views

Keysight, EX-1004
Page 1975

Subsystem Names (Continued)

Name	Includes user interactions with...
Help	Online help
High Availability	High availability feature
IDS Impact Flag	Impact flag configuration
IDS Policy	Intrusion policies
IDSPolicy > *policy_name* > Appliance > *det_engine_name*	Applying intrusion policies
IDSRule sid:*sig_id* rev:*rev_num*	Intrusion rules by SID
Incidents	Intrusion incidents
Insert Policy Apply Job	Applying policies
Install	Installing updates
Intrusion Events	Intrusion events
Login	Web interface login and logout functions
Menu	Any menu option
Object export > *obj_type* > *obj_name*	Importing objects of a specific type and name
Preferences	User preferences such as the time zone for a user account and individual event preferences
Policy	Any policy, including intrusion and OPSEC policies
Register	Registering sensors 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

Keysight, EX-1004
Page 1976

Subsystem Names (Continued)

Name	Includes user interactions with...
SEU Import Log	Viewing the SEU import log
SEU Install	Installing SEUs
Status	Syslog, as well as host and performance statistics
System	Various system-wide settings
System Policy > *policy_name*Appliance > *appliance_name*	Applying system policies
Task Queue	Viewing the task queue
Users	Creating and modifying user accounts

Understanding the Audit Log Table

Platform: DC/MDC/3D Sensor

Component: None

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, **Operations > Monitoring > Audit** is the menu path to view the audit log.
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.

Keysight, EX-1004
Page 1977

Audit Log Fields (Continued)

Field	Description
Source IP	IP address of 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.

Searching Audit Records

Platform: DC/MDC/3D Sensor

Component: None

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 re-use 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.	`Operations > 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.

Keysight, EX-1004
Page 1978

Audit Record Search Criteria (Continued)

Search Field	Description	Example
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 compliance 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 1556 for the syntax for entering time.	`> 2006-01-15 13:30:00` returns all audit records generated after January 15, 2006 at 1:30pm.
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.

For more information on searching, including how to load and delete saved searches, see Searching for Events on page 1550.

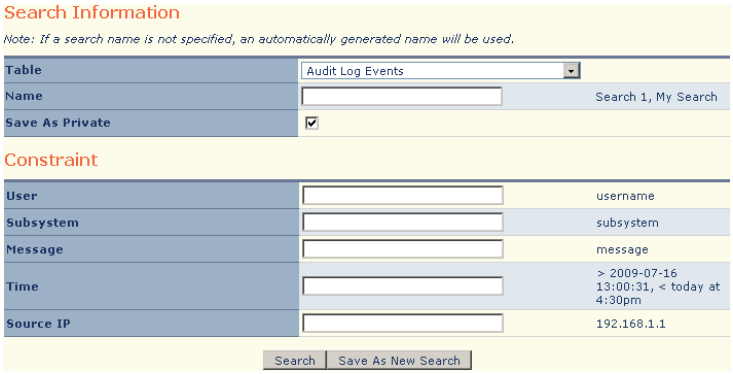
Keysight, EX-1004
Page 1979

To search for audit records:

Access: Admin

1. Select **Analysis & Reporting > Searches > Audit Log**.

   The Audit Log search page appears.



   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 web interface automatically creates one when you save it.

3. Enter your search criteria in the appropriate fields, as described in the Audit Record Search Criteria table. If you enter multiple criteria, the appliance 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 data 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 audit log workflow, constrained by the current time range. To use a different workflow, including a custom workflow, use the **Workflows** menu on the toolbar. For information on specifying a different default workflow, see Configuring Event View Settings on page 45.

Keysight, EX-1004
Page 1980

- 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

Platform: DC/MDC/3D Sensor

Component: None

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 sensors.

You can view system log messages for specific components by using the filter feature. For more information, see Filtering System Log Messages on page 1982.

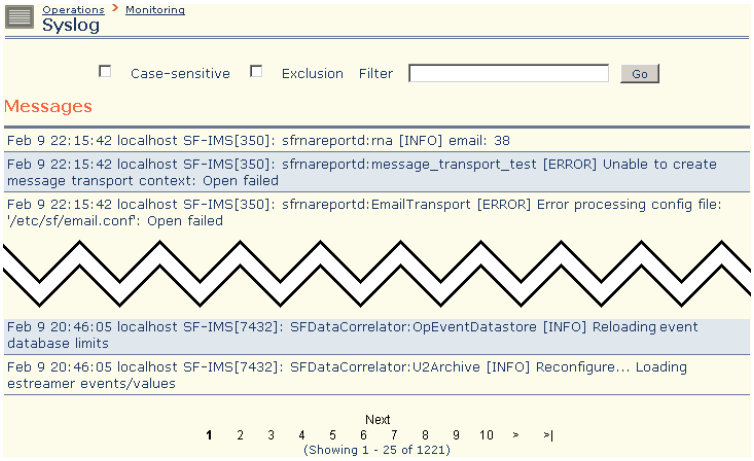
Keysight, EX-1004
Page 1981

To view the syslog:

Access: Maint/Admin

▶ Select **Operations > 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 sensor's syslog. You can find these messages by filtering on `lbim`.

Filtering System Log Messages

Platform: DC/MDC/3D Sensor

Component: None

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.

Keysight, EX-1004
Page 1982

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`, `bdmin`, `Cdmin`, and so on.
`[[:upper:]]`	Matches any uppercase alphabetic character	`[[:upper:]]dmin` matches `Admin`, `Bdmin`, `Cdmin`, and so on.
`[[:lower:]]`	Matches any lowercase alphabetic character	`[[:lower:]]dmin` matches `admin`, `bdmin`, `cdmin`, and so on.
`[[:digit:]]`	Matches any numeric character	`[[:digit:]]dmin` matches `0dmin`, `1dmin`, `2dmin`, and so on.
`[[:alnum:]]`	Matches any alphanumeric character	`[[:alnum:]]dmin` matches `1dmin`, `admin`, `2dmin`, `bdmin`, 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?`.

Keysight, EX-1004
Page 1983

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`
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: Maint/Admin

1. On the System Log page, enter a word or query in the **Filter** field.

   See the System Log Filter Syntax table on page 1983 and the System Log Filter Examples table on page 1984 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.

2. Optionally, to make your search case-sensitive, check **Case-sensitive**. (By default, filters are not case-sensitive.)
3. Optionally, check **Exclusion** to search for all system log messages that do **not** meet the criteria you entered.
4. Click **Go**.

   The messages that match the filter appear.

Keysight, EX-1004
Page 1984

Chapter 46
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, Sourcefire 3D System provides a mechanism for archiving data so that the Defense Center or 3D Sensor can be restored in case of disaster.

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 sensors. The configuration files include information that uniquely identifies a sensor and cannot be shared.

By default, system configuration files are saved in the backup file. You can also choose to back up the following, if applicable for the range of appliances in your deployment:

- the entire intrusion event database
- the entire RNA event database
- additional files that reside on the appliance

WARNING! If you applied any SEU updates, those updates are not backed up. You need to apply the latest SEU update after you restore.

Keysight, EX-1004
Page 1985

You can save backup files to the appliance or to your local computer. Additionally, if you are using a Series 2 Defense Center, you can use remote storage as detailed in Managing Remote Storage on page 1817.

WARNING! Never insert a USB drive into any USB port on a 3D9900 sensor. Additionally, remove any device with external storage (for example, a KVM switch with external storage) from a 3D9900 sensor before upgrading or restoring the sensor.

See the following sections for more information.

- See Creating Backup Files on page 1986 for information about backing up files from the appliance.
- See Creating Backup Profiles on page 1990 for information about creating backup profiles that you can use later as templates for creating backups.
- See Performing Sensor Backup with the Defense Center on page 1991 for information about backing up managed sensors with the Defense Center.
- See Uploading Backups from a Local Host on page 1992 for information about uploading backup files from a local host.
- See Restoring the Appliance from a Backup File on page 1993 for information about how to restore a backup file to the appliance.

Creating Backup Files

Platform: DC/MDC/3D Sensor

Component: None

To view and use existing system backups, go to the System 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 the saved configuration, if needed. You can choose to save the backup file on the appliance or on your local computer.

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 since web browsers do not support uploading files that large. On Series 2 Defense Centers, the backup file can be saved to a remote location; see Managing Remote Storage on page 1817.

IMPORTANT! When your backup task is collecting RNA events, data correlation is temporarily suspended.

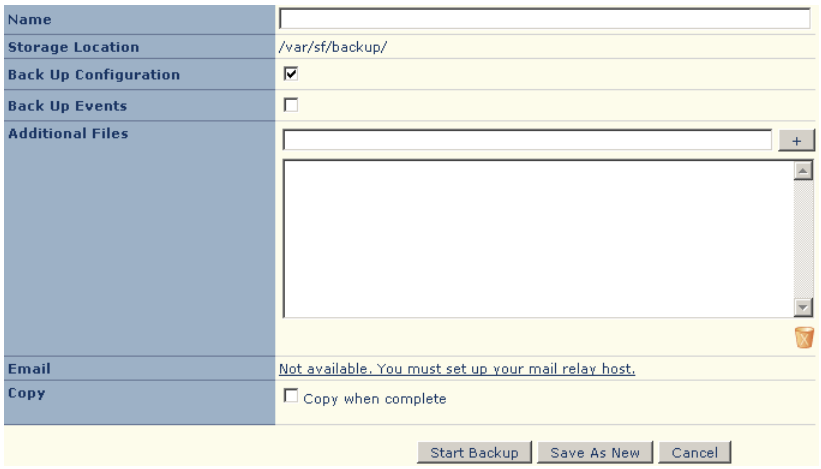
If you perform a backup and then delete reviewed intrusion events, your backup restores the deleted intrusion events but does not restore their reviewed status.

Keysight, EX-1004
Page 1986

You view those restored intrusion events under Intrusion Events, not under Reviewed Events. See Reviewing Intrusion Events on page 278.

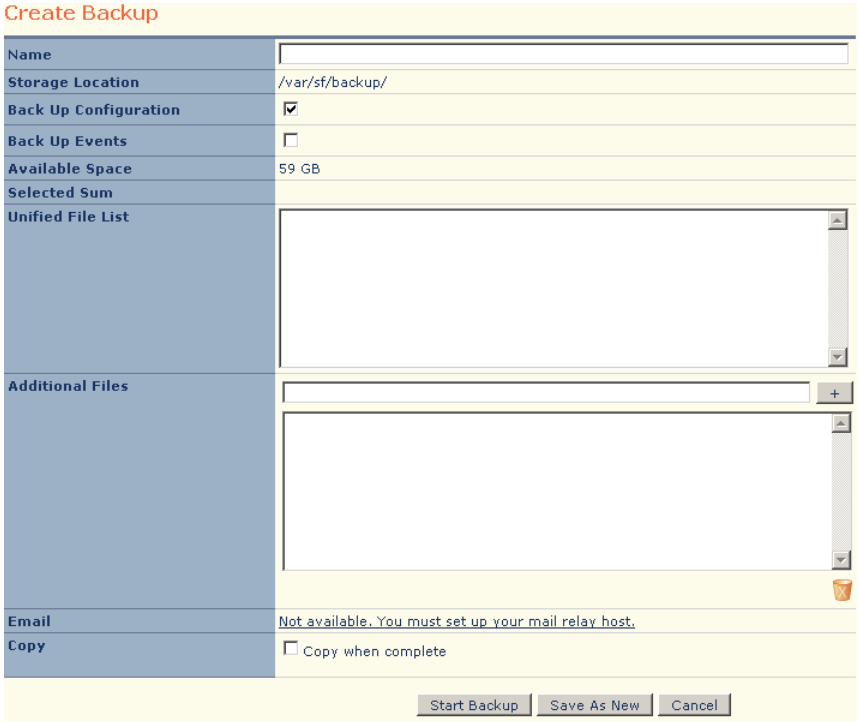
When you back up your 8000 Series sensor from the sensor itself, you back up the configuration only. Use the Defense Center that manages the 8000 Series sensor to perform a complete backup.

The Defense Center and Master Defense Center version of the page is shown below.



Keysight, EX-1004
Page 1987

For comparison, the 3D Sensor version of the page is shown below.



To create a backup file:

Access: Maint/Admin

1. Select **Operations > Tools > Backup/Restore**.

   The System Backup Management page appears.

2. Click **Sensor Backup** on the 3D Sensor toolbar or **Defense Center Backup** on athe Defense Center toolbar.

   The Backup page appears.

3. In the **Name** field, type a name for the backup file.

   You can use alphanumeric characters, punctuation, and spaces.

4. Requires: IPS or DC/MDC To archive the configuration, select **Back Up Configuration**.

5. Requires: IPS or DC/MDC To archive the entire event database, select **Back Up Events**.

6. Requires: IPS To archive individual intrusion event data files, select the files that you want to include from the **Unified File List**.

Keysight, EX-1004
Page 1988

7. Requires: IPS Ensure that the value of the compressed backup file in the **Selected Sum** field is less than the value in the **Available Space** field.

   TIP! The compressed value that appears in the Selected Sum field is a conservative estimate of the size of the compressed file. Often, the file will be smaller.

8. If you want to include an additional file in the backup, type the full path and file name in the **Additional Files** field and click the plus sign (**+**).

   You can repeat this step to add additional files.

9. 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! You must make sure that your mail relay host is configured as described in Configuring a Mail Relay Host and Notification Address on page 1751.

10. Optionally, to use secure copy (scp) to copy the backup archive to a different machine, select the **Copy when complete** check box and then type the following information in the accompanying text boxes:
    - the hostname or IP address of the machine where you want to copy the backup
    - the path to the directory where you want to copy the backup
    - the user name that you want to use to log into the remote machine
    - the password for that user name

    TIP! Sourcefire recommends that you periodically save backups to a remote location so that the appliance can be restored in case of system failure.

Keysight, EX-1004
Page 1989

11. 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 Series 2 Defense Centers, you can direct the backup file to a remote location; see Managing Remote Storage on page 1817.

      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 1993.
    - 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 **Operations > Tools > Backup & Restore** and then clicking **Backup Profiles**. See Creating Backup Profiles on page 1990 for more information.

Creating Backup Profiles

Platform: DC/MDC/3D Sensor

Component: None

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 are backing up the files on your appliance.

TIP! When you create a backup file as described in Creating Backup Files on page 1986, a backup profile is automatically created.

To create a backup profile:

Access: Maint/Admin

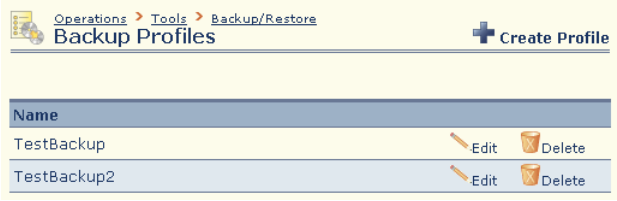
1. Select **Operations > Tools > Backup/Restore**.

   The System Backup Management page appears.

Keysight, EX-1004
Page 1990

2. Click **Backup Profiles** on the toolbar.

   The Backup Profiles page appears with a list of existing backup profiles.



TIP! You can click **Edit** to modify an existing profile or click **Delete** to delete a profile from the list.

3. Click **Create Profile**.

   The System 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 1986 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 includes your new profile in the list.

Performing Sensor Backup with the Defense Center

Platform: DC

Component: None

You can use the Defense Center to back up data on managed 3D Sensors. The default name for the backup file uses the name of the managed 3D Sensor.

TIP! If you use a backup file name containing spaces or punctuation characters, they change to underscores.

You **cannot** use remote backup and restore to manage data on Crossbeam-based software sensors, RNA Software for Red Hat Linux, 3Dx800 sensors, or Intrusion Agents.

Keysight, EX-1004
Page 1991

To back up a managed sensor:

Access: Maint/Admin

1. Select **Operations > Tools > Backup/Restore**.

   The System Backup Management page appears.

2. Click **Sensor Backup** on the toolbar.

   The Remote Backup page appears.



3. In the **Sensors** field, select the managed sensors that 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 the unified files are binary file that the Sourcefire 3D System uses to log event data.
5. To save the backup file on the Defense Center, select the **Retrieve to DC** check box. To save each sensor's backup file on the sensor itself, leave this check box unselected.
6. Click **Start Backup**.

   It can take several minutes to complete the backup. Check the task status for progress. A success messages appears and the backup task is set up. When the backup is complete, you can view the backup file on the Restoration Database page.

Uploading Backups from a Local Host

Platform: DC

Component: None

If you download a backup file to your local host using the download function described in the Backup Management table on page 1993, you can upload it to a Defense Center.

TIP! Uploading a backup larger than 4GB from your local host does not work 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 Series 2 Defense Centers, the backup file can be saved to and retrieved from a remote location; see Managing Remote Storage on page 1817.

Keysight, EX-1004
Page 1992

To upload a backup from your local host:

Access: Maint/Admin

1. Select **Operations > Tools > Backup/Restore**.

   The System Backup Management page appears.

2. Click **Upload Backup.**

   The Upload Backup page appears.

3. Click **Browse**, and navigate to the backup file.

   After you select the file to upload, click **Upload Backup**.

4. Click **Backup Management** on the toolbar to return to the System Backup Management page.

   The backup file is uploaded and appears in the backup list. After the Defense Center verifies the file integrity, refresh the System Backup Management page to reveal detailed file system information.

Restoring the Appliance from a Backup File

Platform: DC/MDC/3D Sensor

Component: None

You can restore the appliance from backup files using the System Backup Management page. After you complete the restoration process, you **must** apply the latest SEU.

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 top of the System Backup Management page. On Series 2 Defense Centers, select **Enable Remote Storage for Backups** to enable or disable remote storage at the top of the System Backup Management page. If you use remote storage, the protocol, backup system, and backup directory are listed at the top of the page. The Backup Management table describes each column and icon on the System Backup Management page.

Backup Management

Column	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

Keysight, EX-1004
Page 1993

Backup Management (Continued)

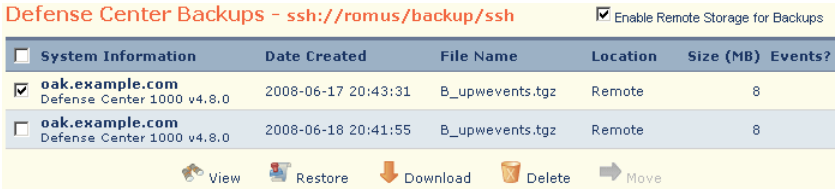
Column	Description
Location	The location of the backup file
Size (MB)	The size of the backup file, in megabytes
Events?	"Yes" indicates the backup includes event data.
View	Click with the backup file selected to view a list of the files included in the compressed backup file.
Restore	Click with the backup file selected to restore it on the appliance.
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 Series 2 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 **Operations > Tools > Backup/Restore**.

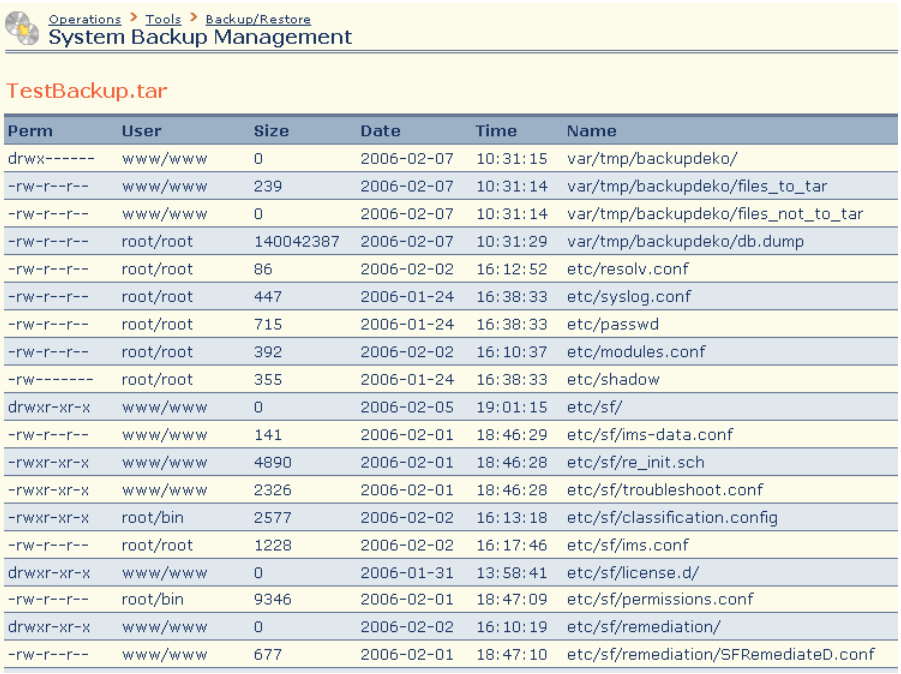
   The System Backup Management page appears. A Series 2 Defense Center version of the page is shown.



Keysight, EX-1004
Page 1994

2. To view the contents of a backup file, select the file and click **View**.

   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.



Operations > Tools > Backup/Restore
System Backup Management

TestBackup.tar

Perm	User	Size	Date	Time	Name
drwx------	www/www	0	2006-02-07	10:31:15	var/tmp/backupdeko/
-rw-r--r--	www/www	239	2006-02-07	10:31:14	var/tmp/backupdeko/files_to_tar
-rw-r--r--	www/www	0	2006-02-07	10:31:14	var/tmp/backupdeko/files_not_to_tar
-rw-r--r--	root/root	140042387	2006-02-07	10:31:29	var/tmp/backupdeko/db.dump
-rw-r--r--	root/root	86	2006-02-02	16:12:52	etc/resolv.conf
-rw-r--r--	root/root	447	2006-01-24	16:38:33	etc/syslog.conf
-rw-r--r--	root/root	715	2006-01-24	16:38:33	etc/passwd
-rw-r--r--	root/root	392	2006-02-02	16:10:37	etc/modules.conf
-rw-------	root/root	355	2006-01-24	16:38:33	etc/shadow
drwxr-xr-x	www/www	0	2006-02-05	19:01:15	etc/sf/
-rw-r--r--	www/www	141	2006-02-01	18:46:29	etc/sf/ims-data.conf
-rwxr-xr-x	www/www	4890	2006-02-01	18:46:28	etc/sf/re_init.sch
-rwxr-xr-x	www/www	2326	2006-02-01	18:46:28	etc/sf/troubleshoot.conf
-rwxr-xr-x	root/bin	2577	2006-02-02	16:13:18	etc/sf/classification.config
-rw-r--r--	root/root	1228	2006-02-02	16:17:46	etc/sf/ims.conf
drwxr-xr-x	www/www	0	2006-01-31	13:58:41	etc/sf/license.d/
-rw-r--r--	root/bin	9346	2006-02-01	18:47:09	etc/sf/permissions.conf
drwxr-xr-x	www/www	0	2006-02-02	16:10:19	etc/sf/remediation/
-rw-r--r--	www/www	677	2006-02-01	18:47:10	etc/sf/remediation/SFRemediateD.conf

3. On the toolbar, click **Backup Management** to return to the System Backup Management page.
4. Select the backup file that you want to restore and click **Restore**.

   The Restore Screen page appears.

   WARNING! This procedure overwrites all configuration files and, on the 3D Sensor, all event data.

5. On the Defense Center, to restore files, select either or both:
   - **Replace Configuration Data**
   - **Restore Event Data**
6. If you want to restore intrusion event data, select the files that you want to include from the Unified File List box.
7. Click **Restore** to begin the restoration.

   The appliance is restored using the backup file you specified.
8. Reboot the appliance.

Keysight, EX-1004
Page 1995

9. Apply the latest SEU to re-apply SEU rule and software updates.
10. Re-apply any intrusion, RNA detection, health, and system policies to the restored system.

Keysight, EX-1004
Page 1996

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:

- custom tables
- custom workflows
- dashboards
- health policies
- intrusion policies
- PEP policies
- RNA detection policies
- system policies
- third-party product mappings
- third-party vulnerability mappings
- RNA detectors

To import an exported object, both appliances must be running the same version of the Sourcefire 3D System. To import an exported intrusion policy, the SEU versions on both appliances must also match.

For more information, see the following sections:

- Exporting Objects on page 1998
- Importing Objects on page 2009

Keysight, EX-1004
Page 1997

Exporting Objects

Platform: DC/MDC/Standalone Sensor

Component: None

You can export a single object, or you can export several objects at once.

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, and then export a Defense Center system policy with authentication enabled, the authentication object is exported as well.

Note that depending on the number of objects being exported and the number of objects those objects reference, the export process may take several minutes.

For more information, see the following sections:

- Exporting a Custom Table on page 1999
- Exporting a Custom Workflow on page 1999
- Exporting a Dashboard on page 2000
- Exporting a Health Policy on page 2001
- Exporting an Intrusion Policy on page 2001
- Exporting a PEP Policy on page 2002
- Exporting an RNA Detection Policy on page 2003
- Exporting a System Policy on page 2003
- Exporting a Third-Party Product Mapping on page 2004
- Exporting a Third-Party Vulnerability Mapping on page 2005
- Exporting an RNA Detector on page 2005
- Exporting Multiple Objects on page 2006

Keysight, EX-1004
Page 1998

Exporting a Custom Table

Platform: DC

Component: RNA

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.

To export a custom table:

Access: Any RNA/Admin

1. Make sure that the Defense Center where you are exporting the custom table and the Defense Center where you plan to import the custom table are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Analysis & Reporting > Custom Tables**.

   The Custom Tables page appears.

3. Click **Export** next to the custom table you want to export.

4. Follow your web browser's prompts to save the exported package to your computer.

Exporting a Custom Workflow

Platform: DC/MDC/Standalone Sensor

Component: None

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 an appliance 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. For example, you cannot access a custom workflow based on RNA hosts that you created on the Defense Center and then imported onto a 3D Sensor or Master Defense Center.

To export a custom workflow:

Access: Any Analyst/Admin

1. Make sure that the appliance where you are exporting the custom workflow and the appliance where you plan to import the custom workflow are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

Keysight, EX-1004
Page 1999

2. Select **Analysis & Reporting > Custom Workflows**.

   The Custom Workflows page appears.

3. Click **Export** next to the custom workflow you want to export.
4. Follow your web browser's prompts to save the exported package to your computer.

Exporting a Dashboard

Platform: DC/MDC/Standalone Sensor

Component: None

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 example, a dashboard created on the Defense Center and imported onto a 3D Sensor or Master Defense Center may display some invalid, disabled widgets. For more information, see Understanding Widget Availability on page 63.

To export a dashboard:

Access: Any except Restricted or Database

1. Make sure that the appliance where you are exporting the dashboard and the appliance where you plan to import the dashboard are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Analysis & Reporting > Dashboards**.

   If you have a default dashboard defined, it appears; continue with the next step.

   If you do not have a default dashboard defined, the Dashboard List page appears; skip to step 4.

3. On the toolbar, click **Dashboards**.

   The Dashboard List page appears.

4. Click **Export** next to the dashboard you want to export.
5. Follow your web browser's prompts to save the exported package to your computer.

Keysight, EX-1004
Page 2000

Exporting a Health Policy

Platform: DC/MDC

Component: None

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.

To export a health policy:

Access: Maint/Admin

1. Make sure that the appliance where you are exporting the health policy and the appliance where you plan to import the health policy are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Operations** > **Monitoring** > **Health**.

   The Health Monitor page appears.

3. On the toolbar, click **Health Policy**.

   The Health Policy page appears.

4. Click **Export** next to the policy you want to export.

5. Follow your web browser's prompts to save the exported package to your computer.

Exporting an Intrusion Policy

Platform: DC/MDC/Standalone Sensor

Component: IPS

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; RNA 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

Keysight, EX-1004
Page 2001

policy on another appliance, you can edit the imported policy to suit your needs, including deleting, adding, and sharing 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 SEU version; see Importing SEUs and Rule Files on page 1848.

To export an intrusion policy:

Access: P&R Admin/Admin

1. Make sure that the appliance where you are exporting the intrusion policy and the appliance where you plan to import the policy are running the same version of the Sourcefire 3D System, as well as the same version of the SEU.

   If the versions of the Sourcefire 3D System and the SEU do not match, the import will fail.

2. Select **Policy & Response > IPS > Intrusion Policy**.

   The Intrusion Policy page appears.

3. Click **Export** next to the intrusion policy you want to export.

   Depending on the number of rules referenced by the policy you are exporting, the export process may take several minutes.

4. Follow your web browser's prompts to save the exported package to your computer.

Exporting a PEP Policy

Platform: DC/MDC

Component: None

PEP policies allow you to configure 3D9900 sensors to block, analyze, or pass traffic directly through the sensor with no further inspection by taking advantage of the hardware capabilities on those sensors.

To export a PEP Policy:

Access: P&R Admin/Admin

1. Make sure that the appliance where you are exporting the PEP policy and the appliance where you plan to import the PEP policy are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Policy & Response > PEP > Policy Management**.

   The PEP Policy Management page appears.

3. Click **Export** next to the policy you want to export.

Keysight, EX-1004
Page 2002

4. Follow your web browser's prompts to save the exported package to your computer.

Exporting an RNA Detection Policy

Platform: DC/MDC

Component: RNA

RNA detection policies control how RNA events and flow data are collected, which network segments are monitored by 3D Sensor with RNA and which are monitored with NetFlow-enabled devices, and whether traffic that travels from or to specific ports is excluded from monitoring.

To export an RNA detection policy:

Access: P&R Admin/Admin

1. Make sure that the Defense Center where you are exporting the detection policy and the Defense Center where you plan to import the detection policy are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Policy & Response** > **RNA** > **Detection Policy**.

   The Detection Policy page appears.

3. Click **Export** next to the policy you want to export.

4. Follow your web browser's prompts to save the exported package to your computer.

Exporting a System Policy

Platform: DC/MDC/Standalone Sensor

Component: None

A system policy controls the aspects of an appliance that are likely to be similar for other Sourcefire 3D System appliances in your deployment, including database event limits, time settings, login banners, and so on.

When you export a system policy from a Defense Center where external authentication is enabled, the Defense Center also exports the authentication objects on which the system policy depends. That is, if you set up authentication to an LDAP server on your Defense Center, and then export a Defense Center system policy with authentication enabled, the authentication object is exported as well.

Note that system policies on Defense Centers contain database settings that do not apply to 3D Sensors. If you export a system policy from a 3D Sensor and then import it onto a Defense Center, the database limits that you could not configure on the sensor are set to the default values on the Defense Center.

Keysight, EX-1004
Page 2003

To export a system policy:

Access: Admin

1. Make sure that the appliance where you are exporting the system policy and the appliance where you plan to import the system policy are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Operations > System Policy**.

   The System Policy page appears.

3. Click **Export** next to the system policy you want to export.

4. Follow your web browser's prompts to save the exported package to your computer.

Exporting a Third-Party Product Mapping

Platform: DC

Component: RNA

If you import data from a third-party application, you must map the RNA product to the third-party name to assign vulnerabilities and perform impact correlation using that data. Mapping the product associates RNA 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 1408.

To export a third-party product mapping:

Access: Admin

1. Make sure that the Defense Center where you are exporting the mapping and the Defense Center where you plan to import the mapping are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Operations** > **Tools** > **Import/Export**.

   The Import/Export page appears, including a list of the objects on the appliance. The third-party product mappings you can export are listed in the 3rd Party Map Sets section.



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.

Keysight, EX-1004
Page 2004

Exporting a Third-Party Vulnerability Mapping

Platform: DC

Component: RNA

To add vulnerability information from a third-party application to the VDB, you must map the third-party identification string for each imported vulnerability to any existing RNA, 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 1412.

To export a third-party vulnerability mapping:

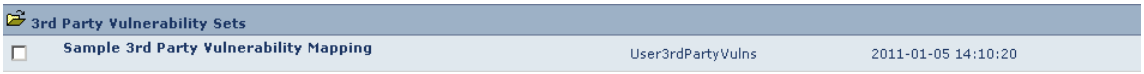
Access: Admin

1. Make sure that the Defense Center where you are exporting the mapping and the Defense Center where you plan to import the mapping are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Operations** > **Tools** > **Import/Export**.

   The Import/Export page appears, including a list of the objects on the appliance. The third-party vulnerability mappings you can export are listed in the 3rd Party Vulnerability Sets section.



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.

Exporting an RNA Detector

Platform: DC/MDC

Component: RNA

When RNA analyzes IP traffic, it uses RNA detectors to collect information about and then identify the commonly used services and client applications running on hosts on your network.

You can export two kinds of RNA detector:

- user-defined detectors
- individual add-on detectors provided by Sourcefire Professional Services

Keysight, EX-1004
Page 2005

For more information on detectors, see Working with Service and Application Detectors on page 1381.

To export an RNA detector:

Access: P&R Admin/Admin

1. Make sure that the Defense Center where you are exporting the detector and the Defense Center where you plan to import the detector are running the same version of the Sourcefire 3D System.

   If the versions of the Sourcefire 3D System do not match, the import will fail.

2. Select **Policy & Response > RNA > RNA Detectors**.

   The RNA Detectors page appears.

3. Select the check box next to the detector you want to export and click **Export**.

   Depending on how many detectors you have, the detector you want to export may not be 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 RNA Detectors on page 1394.

   TIP! To export multiple detectors at once, select the check boxes next to the appropriate detectors, then click **Export**. You can also select all detectors in the current filtered view by selecting the check box at the top of the page.

4. Follow your web browser's prompts to save the exported package to your computer.

Exporting Multiple Objects

Platform: DC/MDC/Standalone Sensor

Component: None

You can export several different objects at once (in a single package) using the Import/Export feature. When you later import the package onto another appliance, you choose which objects in the package to import.

The following table lists the objects that you can export from the various Sourcefire appliance types.

Objects with Import and Export Capability

Object	Platform	Component
Custom Tables	DC/MDC	None
Custom Workflows	DC/MDC/Standalone Sensor	None

Keysight, EX-1004
Page 2006

Objects with Import and Export Capability (Continued)

Object	Platform	Component
Dashboards	DC/MDC/Standalone Sensor	None
Health Policies	DC/MDC	None
Intrusion Policies	DC/MDC/Standalone Sensor	IPS
PEP Policies	DC/MDC	None
RNA Detection Policies	DC/MDC	RNA
System Policies	DC/MDC/Standalone Sensor	None
Third-Party Product Mappings	DC	RNA
Third-Party Vulnerability Mappings	DC	RNA
RNA Detectors	DC	RNA

Depending on the type of object you are exporting, you should keep the following points in mind:

- You must make sure that the appliance you are using to export an object is running the same version of the Sourcefire 3D System as the appliance you plan to use to import the exported object. For intrusion policies, the SEU versions on both appliances must also match. If the versions do not match, the import will fail.
- If you cannot view the table on which a custom workflow is based on your appliance, you can import the workflow but will not be able to view it.
- 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 3D Sensor or Master Defense Center may display some invalid, disabled widgets.
- 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.

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 SEU version; see Importing SEUs and Rule Files on page 1848.

Keysight, EX-1004
Page 2007

- When you export a system policy from a Defense Center where external authentication is enabled, the Defense Center also exports the authentication objects on which the system policy depends.

  Also note that if you export a system policy from a 3D Sensor and then import it onto a Defense Center, the database limits that you could not configure on the sensor are set to the default values on the Defense Center.
- You can export user-defined RNA detectors and individual add-on detectors provided by Sourcefire Professional Services.You cannot export any other kind of RNA detector.

For detailed information on exporting specific objects, see the following sections:

- Exporting a Custom Table on page 1999
- Exporting a Custom Workflow on page 1999
- Exporting a Dashboard on page 2000
- Exporting a Health Policy on page 2001
- Exporting an Intrusion Policy on page 2001
- Exporting a PEP Policy on page 2002
- Exporting an RNA Detection Policy on page 2003
- Exporting a System Policy on page 2003
- Exporting a Third-Party Product Mapping on page 2004
- Exporting a Third-Party Vulnerability Mapping on page 2005
- Exporting an RNA Detector on page 2005

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 multiple 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. For intrusion policies, you must also make sure that the SEU version matches.

   If the versions of the Sourcefire 3D System (and, for intrusion policies, the SEU version) do not match, the import will fail.

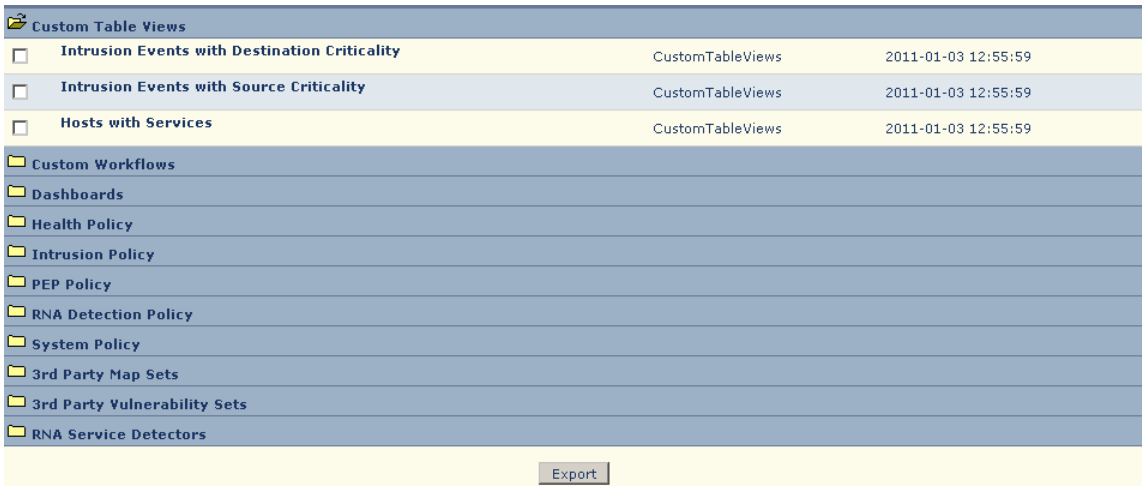
Keysight, EX-1004
Page 2008

2. Select **Operations** > **Tools** > **Import/Export**.

   The Import/Export page appears, including a list of the objects on the appliance.

   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

Platform: DC/MDC/Standalone Sensor

Component: None

After you export an object from another 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

Keysight, EX-1004
Page 2009

are using and on your user role. The following table lists the objects that you can import on the various Sourcefire appliance types.

Objects with Import and Export Capability

Object	Platform	Component
Custom Tables	DC/MDC	None
Custom Workflows	DC/MDC/Standalone Sensor	None
Dashboards	DC/MDC/Standalone Sensor	None
Health Policies	DC/MDC	None
Intrusion Policies	DC/MDC/Standalone Sensor	IPS
PEP Policies	DC/MDC	None
RNA Detection Policies	DC/MDC	RNA
System Policies	DC/MDC/Standalone Sensor	None
Third-Party Product Mappings	DC	RNA
Third-Party Vulnerability Mappings	DC	RNA
RNA Detectors	DC	RNA

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 are importing an object is running the same version of the Sourcefire 3D System as the appliance you used to export the object. For intrusion policies, the SEU versions on both appliances must also match. If the versions do not match, the import will fail.
- If your appliance does not allow you to view the table on which an custom workflow is based, you can import the workflow but will not be able to view it.
- 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 3D Sensor or Master Defense Center may display some invalid, disabled widgets.

Keysight, EX-1004
Page 2010

- 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 SEU version; see Importing SEUs and Rule Files on page 1848.

- 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.

  Also note that for a system policy exported from a 3D Sensor and then imported onto a Defense Center, the database limits that you could not configure on the sensor are set to the default values on the Defense Center.

Because 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.

WARNING! If you import default policies or rules to an appliance, you cannot delete them.

For information on using imported objects, see the following sections:

- Using Custom Tables on page 1561
- Using Custom Workflows on page 1618
- Working with Dashboards on page 92
- Applying Health Policies on page 1928
- Applying an Intrusion Policy on page 347
- Applying PEP Policies on page 1527
- Applying an RNA Detection Policy on page 919
- Applying a System Policy on page 1736

Keysight, EX-1004
Page 2011

- Mapping Third-Party Products on page 1408
- Mapping Third-Party Vulnerabilities on page 1412
- Activating and Deactivating RNA Detectors on page 1398

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. For intrusion policies, you must also make sure that the SEU version matches.

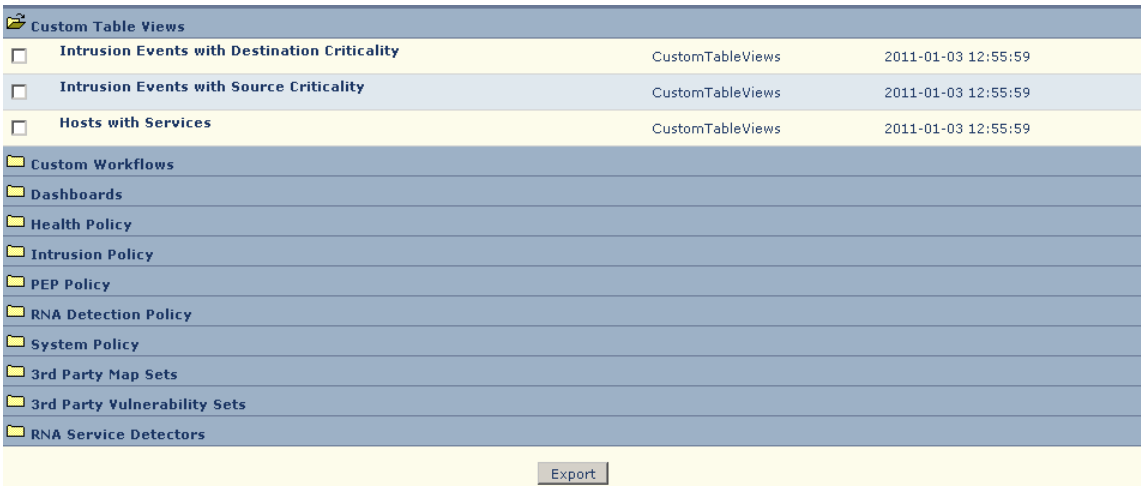
   If the version of the Sourcefire 3D System (and, for intrusion policies, the SEU version) do not match, the import will fail.

2. Export the objects you want to import; see Exporting Objects on page 1998.
3. On the appliance where you want to import the objects, select **Operations** > **Tools** > **Import/Export**.

   The Import/Export page appears.

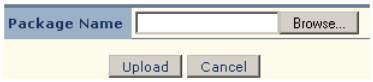
   TIP! You can click the collapse icon () next to an object type to collapse the list of objects. Click the expand folder icon () next to an object type to reveal objects.

   The Defense Center version of the page is shown below with some object types collapsed.



4. Click **Upload Package**.

   The Upload Package page appears.



Keysight, EX-1004
Page 2012

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 object and rule versions 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 SEU 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 SEU 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 occurs, 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 the objects you import do have previous revisions on your appliance, the Import Resolution page appears. Continue with step 8.



8. 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 existing if newer**.
   - To save the imported object as a new object, select **Import as new**, and, optionally, edit the object name.
9. Click **Import**.

   The objects are imported.

Keysight, EX-1004
Page 2013

Appendix B
Purging the RNA and RUA Databases

You can use the RNA/RUA Event Purge page to purge files from the RNA and RUA databases. Note that when you purge a database, the appropriate process (RNA or RUA) is restarted.

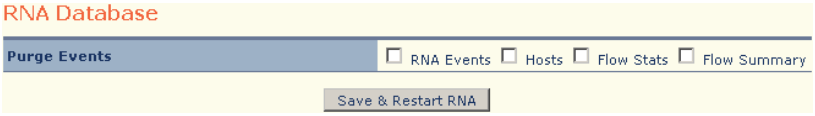
WARNING! Purging a database removes the data you specify from the Defense Center. After the data is deleted, it **cannot** be recovered.

To purge the RNA database:

Access: RNA/Admin

1. Select **Operations > Configuration > RNA/RUA Event Purge**.

   The RNA/RUA Event Purge page appears.



2. Under RNA Database, perform any or all of the following:
   - Select **RNA Events** to remove all network discovery events from the RNA database.
   - Select **Hosts** to remove all hosts from the RNA database.
   - Select **Flow Stats** to remove all flow data from the RNA database.
   - Select **Flow Summary** to remove all flow summary data from the RNA database.

Keysight, EX-1004
Page 2014

3. Click **Save & Restart RNA**.

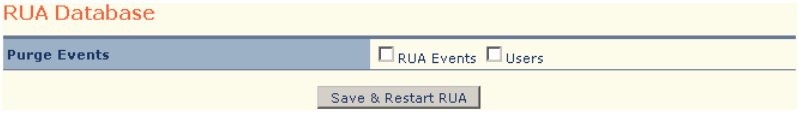
   The items are purged and RNA is restarted.

To purge the RUA database:

Access: RNA/Admin

1. Select **Operations > Configuration > RNA/RUA Event Purge**.

   The RNA/RUA Event Purge page appears.



2. Under RUA database, perform any or all of the following:
   - Select **RUA Events** to remove all RUA events from the RUA events database.
   - Select **Users** to remove all users from the RUA history database.
3. Click **Save & Restart RUA**.

   The items are purged and RUA is restarted.

Appendix C
Viewing the Status of Long-Running Tasks

When you perform long-running tasks, such as applying a policy or installing software, the status of these tasks is reported in the task queue. The task queue also reports when they are complete.

For more information, see the following sections:

- Viewing the Task Queue on page 2016
- Managing the Task Queue on page 2018

Viewing the Task Queue

Platform: DC/MDC/3D Sensor

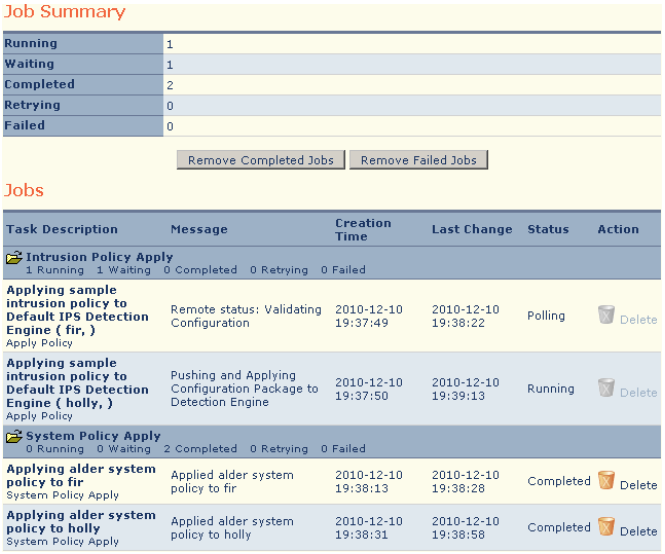
Component: None

When you perform long-running tasks, such as applying a policy or installing software, 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 account has Administrator access, you can also see the status of every task regardless of who initiated it.

Keysight, EX-1004
Page 2016

You view the task queue on the Task Status page, which automatically refreshes every 10 seconds. As an Administrator, you can see the status of every task, regardless of who initiated it.



The Job Summary section displays the state of the tasks listed on the page, as described in the following table.

Task Queue Task Types

Task Type	Description
Running	The number of tasks currently in progress.
Waiting	The number of tasks waiting for a in-progress task to complete before running.
Completed	The number of tasks that completed, regardless of whether they succeeded,
Retrying	The number of tasks that are automatically retrying. Note that not all tasks are permitted to try again.
Failed	The number of tasks that did not complete successfully.

The Jobs section provides information about each task, including a brief description, when the task was launched, the current status of the task, and when the status last changed. Tasks of the same type appear together.

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. You can also manually remove tasks from the queue; see Managing the Task Queue on page 2018.

Keysight, EX-1004
Page 2017

To view the task queue:

Access: Maint/P&R Admin/Admin

▶ 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 **Operations > Monitoring > Task Status**.

  The Task Status page appears.

For information on the actions you can perform on the Task Status page, see the next section, Managing the Task Queue.

Managing the Task Queue

Platform: DC/MDC/3D Sensor

Component: None

If you have Administrator, Maintenance, or Policy & Response Administrator access, there are several actions you can perform while viewing the task queue (see Viewing the Task Queue on page 2016), as described in the following table.

Task Queue Actions

To...	You can...
remove all completed tasks from the task queue	click **Remove Complete Jobs**.
remove all failed tasks from the task queue	click **Remove Failed Jobs**.
remove a single task from the task queue	click **Delete** 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 the view of tasks of the same type	click the collapse icon () next to the task type for the tasks you want to hide.
expand the view of tasks of the same type	click the expand folder icon () next to the task type for which you want to view individual tasks.

Keysight, EX-1004
Page 2018

Appendix D
Detected Operating Systems and Services

The RNA component of the Sourcefire 3D System detects a wide variety of operating systems, client applications, and services.

The specific operating systems, services, client applications, and HTTP payload content that RNA detects depend on both the version of the Sourcefire 3D System and the version of the vulnerability database (VDB) you have installed. This list may change over time as fingerprints are added and as company and product names change. 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 RNA detectors.

To further enhance RNA's detection capabilities, you can also create your own service detectors, as well as import individual add-on service and client application detectors provided by Sourcefire Professional Services.

The following sections provide a list of detected operating systems, services, client applications, and payloads:

- Detected Operating Systems on page 2020
- Detected Services on page 2027
- Detected Client Applications on page 2036
- Detected Payloads on page 2039

Keysight, EX-1004
Page 2019

Detected Operating Systems

Platform: DC

Component: RNA

RNA detects the operating systems of multiple types of devices. Note that the operating systems that RNA detects depend on the version of the VDB you have installed. This list may change over time as fingerprints are added and as company and product names change.

The operating systems that RNA detects are listed in the following sections:

- Desktops, Servers, and NAS Devices on page 2020
- Printers on page 2023
- Gaming Systems, Music, and Television on page 2024
- Telephones, Communications, VoIP, and Videoconferencing on page 2024
- Routers, Switches, Access Points, and Network Appliances on page 2026

Desktops, Servers, and NAS Devices

Platform: DC

Component: RNA

RNA detects the operating systems for the desktop systems and servers listed in the Detected Computer Operating Systems table.

Detected Computer Operating Systems

Operating System	Versions (if available)
Apple Mac OS	9.x
Apple Mac OS X Apple Mac OS X Server	10.x **IMPORTANT!** Because Mac OS X and OS X Server have identical fingerprints, RNA may identify a host running either operating system as running both. This does **not** affect vulnerability mappings or impact assessments.
BeOS	unknown
BSD	FreeBSD 3.1 - 7.x NetBSD 1.5.2, 1.6, 2.0 OpenBSD 2.4, 2.5, 2.6, 2.8, 3.x, 4.x

Keysight, EX-1004
Page 2020

Detected Computer Operating Systems (Continued)

Operating System	Versions (if available)
HP-UX	10.2 11.0, 11.11
HP VMS HP OpenVMS HP OpenVMS AXP	unknown
HP iLO Agent	unknown
IBM AIX	1.x 2.2.1 3.x 4.0 - 4.2, 4.3.1, 4.3.2 5.1 - 5.3 6.2
IBM AIX 5L	5.1, 5.2, 5.3
IBM OS/2	unknown
IBM z/OS	V1.R4
Linux	2.1.19 - 2.2.20 kernels * 2.4 - 2.6 kernels ** Debian 2.2 Knoppix Linspire PHLAK Red Hat Enterprise Linux 5.4 Red Hat Enterprise Linux AS 4.x SCO OpenLinux 1.3

Keysight, EX-1004
Page 2021

Detected Computer Operating Systems (Continued)

Operating System	Versions (if available)
Microsoft Windows	95 98, 98 SE Me NT 2000 and 2000 Professional, including SP3 and SP4 2000, 2003, and 2008 Server 2008, including R2 XP and XP Professional, including SP1 - SP3 Vista, including SP2 7
Novell Client	unknown
Novell NetWare	5.0, 5.1 6.0
PXE	unknown
Red Hat Anaconda Installer	unknown
SGI Irix	6.3, 6.5, 6.5.12
SunOS	5.6 - 5.10
Sun Solaris	2.6 - 2.8 7.0 - 10.x

*The Linux 2.1.19 - 2.2.20 kernels are used by a variety of operating systems, including but not limited to the following:

Operating Systems Identified as Linux 2.1.19-2.2.20

Mandrake 7.2, 8.0	SuSE 7.2
Red Hat 6.0, 7.0	Trustix 1.5

Keysight, EX-1004
Page 2022

**The Linux 2.4 - 2.6 kernels are used by a variety of operating systems, including but not limited to the following:

Operating Systems Identified as Linux 2.4-2.6

Knoppix 3.3	SuSE 7.3, 8.0 - 8.2
Mandrake 8.1, 8.2, 9.0, 9.1	Trustix 2.0
Red Hat 7.1 - 7.4, 8.0, 9.0	Ubuntu 6.0

Printers

Platform: DC

Component: RNA

RNA detects the printers listed in the Detected Printers table.

Detected Printers

Vendor and Type
Brother printers
Canon printers
Dell printers
HP printers, including LaserJet 4100 and 8550
Xerox printers

Keysight, EX-1004
Page 2023

Gaming Systems, Music, and Television

Platform: DC

Component: RNA

RNA detects the gaming systems, music devices, and television devices (such as DVRs) listed in the Detected Entertainment Systems table.

Detected Entertainment Systems

Vendor	Type
Amino	Aminet Set Top Box
Microsoft	Xbox
Nintendo	GameCube Wii
Replay	TV Converter
Slingbox	Slingbox
Sony	Playstation
TiVo	TiVo

Telephones, Communications, VoIP, and Videoconferencing

Platform: DC

Component: RNA

RNA detects the telecommunications devices listed in the Detected Communications Devices table.

Detected Communications Devices

Vendor	Name/Type	Versions (if available)
Avaya	IP phone	4602 SW (model 4602D02A)
BATM	VoIP adapter	unknown

Keysight, EX-1004
Page 2024

Detected Communications Devices (Continued)

Vendor	Name/Type	Versions (if available)
Cisco	IP phone	7905, 7905G 7910, version 3.x 7912G 7940 (version 7.1.1) 7960 (version 7.1.1)
Cisco	VoIP phone	CP-7940 CP-7960
Cisco	Linksys VoIP PAP adapter	unknown
Clipcom	VoIP WiFi phone	CPW-100E (version 1.1.12) CP-100E (version 1.1.60)
Mediatrix	VoIP adapter	unknown
Nokia	Internet Tablet	unknown
Nortel	CallPilot	unknown
Nortel	IP phone	unknown
Polycom	ViewStation	unknown
Sipura	VoIP adapter	unknown
Sunrocket	VoIP Gizmo	unknown
Tandberg	Videophone	unknown
UniData	IP Phone	unknown
Uniden	VoIP adapter	unknown

Keysight, EX-1004
Page 2025

Routers, Switches, Access Points, and Network Appliances

Platform: DC

Component: RNA

RNA detects the networking devices listed in the Detected Routers, Switches, and Access Points table.

Detected Routers, Switches, and Access Points

Vendor	Name/Type	Versions (if available)
A-V Tronics	2Wire Residential Gateway	unknown
Apple	AirPort	unknown
Bluesocket	BSC	unknown
Buffalo	wireless access point	unknown
Cisco	Aironet access point	340, 350 1130AG 1200, 1240AG 1300 3500 4500, 4800
Cisco	Catalyst switch	29xx, 2900 LRE 35xx 5500 6500
Cisco	IOS	11.x 12.x
Cisco	Linksys router	unknown
Compex	access point	unknown
Gemtek	wireless router	unknown
Hotway	LANDrive	unknown

Keysight, EX-1004
Page 2026

Detected Routers, Switches, and Access Points (Continued)

Vendor	Name/Type	Versions (if available)
Motorola	NIM100	unknown
Netgear	router	unknown
Nortel	Alteon Application Switch	AD4, 184
Nortel	Alteon Switched Firewall	5100
Nortel	Alteon Web Switch	2.8.11
Nortel	BayStack Switch	325 420, 425, 450
Nortel	Passport 8600 Routing Switch	8610.x
Nortel	Ethernet Routing Switch	8300
OEMed	wireless router	unknown

Detected Services

Platform: DC

Component: RNA

When RNA analyzes IP traffic, it uses RNA detectors to identify the commonly used services running on hosts on your network.

For more information, see the following sections:

- Service Detector Types on page 2028
- Special Considerations for Service Detection on page 2029
- Services Detected by RNA on page 2030

Keysight, EX-1004
Page 2027

Service Detector Types

Platform: DC

Component: RNA

There are a few different types of service detectors, as described in the following table.

RNA Service Detector Types

Detector Type	Description
internal	Service detectors that are delivered with updates to the Sourcefire 3D System.
add-on	Service detectors that are provided either via VDB updates or as individual importable detectors authored by Sourcefire Professional Services. Detectors provided through these delivery methods allow you to enhance RNA's detection capabilities with new pattern-based detectors without updating the Sourcefire 3D System as a whole. **IMPORTANT!** To improve performance, Sourcefire sometimes replaces add-on detectors with internal detectors delivered through updates to the Sourcefire 3D System.
well-known port	Service detectors, provided by Sourcefire, that are based on detection of network traffic on well-known ports.
user-defined	Service detectors that you create, based on a port, pattern, or both.
implied from application	The implied existence of a service based on the detection of a client application. If RNA can identify the client application used in a flow between a monitored host accessing a non-monitored server the Defense Center displays the service that corresponds with the application for that flow. Note that under normal conditions,RNA tracks only services on monitored networks, so flow events usually do not include service information for flows where a monitored host is accessing a non-monitored server.

For more information on working with service detectors, including information on creating user-defined service detectors, see Working with Service and Application Detectors on page 1381.

Keysight, EX-1004
Page 2028

Special Considerations for Service Detection

Platform: DC

Component: RNA

The following sections outline some important points you should keep in mind about RNA service detection.

Service Detectors Included in VDB and Sourcefire 3D System Updates

The services that RNA detects depend on both the version of the Sourcefire 3D System and the version of the VDB you have installed. 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 RNA detectors.

To further enhance RNA's detection capabilities, you can also create your own service detectors, as well as import individual add-on service detectors provided by Sourcefire Professional Services. For more information, see Working with Service and Application Detectors on page 1381.

Implied Service Detection from Client Application Detection

If RNA can identify the client application used in a flow between a monitored host accessing a non-monitored server, for that flow, the Defense Center displays the service that corresponds with the application. (Because RNA tracks only services on monitored networks, flow events usually do not include service information for flows where a monitored host is accessing a non-monitored server.)

There are several consequences of the implied detection of a service from the detection of a client application:

- Because RNA does not generate a New Service event for these services, the service does not appear in the Services table. In addition, you cannot trigger either RNA event alerts or compliance rules using the detection of these services as a criterion.
- Because the service is not associated with a host, you cannot view service details in host profiles, set the service's identity, or use the service information in host profile qualifications for traffic profiles or compliance rules. In addition, RNA does not associate vulnerabilities with hosts based on this type of service-in-flow detection.

You can, however, trigger compliance events on the service information in a flow. You can also use the service information in flow events to create flow trackers and traffic profiles.

Load Balancing RNA Detection Resources and Detecting VoIP Traffic

There can be unexpected consequences as a result of the load balancing that occurs when you assign more than one detection resource per RNA detection engine. For example, if you assign multiple detection resources to an RNA

Keysight, EX-1004
Page 2029

detection engine, RNA may not correctly report VoIP traffic, which uses both SIP (session initiation protocol) and RTP (real-time transport protocol). If the same detection resource does not process the SIP and then the RTP packets for the same VoIP session, the web interface will not report the RTP traffic. For more information on detection engines and detection resources, see Understanding Detection Resources and 3D Sensor Models on page 202.

Squid Proxy Detection

RNA positively identifies Squid service traffic when either:

- RNA detects a connection from a host on your monitored network to a Squid server where proxy authentication is enabled, or
- RNA 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, RNA 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

Services Detected by RNA

Platform: DC

Component: RNA

RNA detects the services listed in the following table. For each service, the table lists how it is identified within the Sourcefire 3D System and also provides a short description of the service and the type of service detector that detects it.

Detected Services

Service Detector	Description	Detector Type
aim	AOL Instant Messenger	internal
ajp	Apache JServ Protocol (AJP) communications between web servers and servlet repositories	add-on
applejuice	Applejuice peer-to-peer file sharing	add-on
battlefield	Battlefield multiplayer computer game	add-on
bgp	Border Gateway Protocol	internal

Keysight, EX-1004
Page 2030

Detected Services (Continued)

Service Detector	Description	Detector Type
bt	BitTorrent file sharing	internal
caicci	Common Communication Interface (CCI) communications between Computer Associates (CAI) systems	add-on
chargen	character generator (CHARGEN) service	well-known port
checkpoint	Check Point firewall management interface	add-on
cvspserver	Concurrent Versions System (CVS) server	well-known port
daytime	DAYTIME service	well-known port
dcerpc	Distributed Computing Environment/Remote Procedure Call (DCE/RPC) services	internal
dhcp	Dynamic Host Configuration Protocol (DHCP) services	internal
dhcpv6-server	Dynamic Host Configuration Protocol (DHCP) services	well-known port
directconnect	DirectConnect peer-to-peer file sharing	add-on
discard	DISCARD protocol	well-known port
domain	Domain Name Service (DNS) servers	internal
echo	echo protocol	well-known port
exec	remote execution services (exec, rexec)	internal
finger	Finger protocol	well-known port
font-service	X server-font renderer communication services	well-known port

Keysight, EX-1004
Page 2031

Detected Services (Continued)

Service Detector	Description	Detector Type
ftp ftp data	File Transfer Protocol	internal
ganglia	Ganglia (scalable distributed monitoring system for high-performance computing systems)	add-on
giop	General Inter-ORB Protocol (communication between object request brokers)	add-on
gnutella	Gnutella peer-to-peer file sharing	internal
gopher	Gopher protocol	well-known port
hostname	NIC Internet Hostname Server	well-known port
http	Hypertext Transfer (HTTP) Protocol	internal
ica	Citrix Independent Computing Architecture (ICA) remote desktop service	add-on
icabrowser	Citrix Independent Computing Architecture (ICA) Browser	add-on
icap	Internet Content Adaptation Protocol (proxy service for specific content)	add-on
ident	Ident protocol	well-known port
IKE	Internet Key Exchange protocol	add-on
imap imap3	Internet Message Access Protocol (IMAP) mail servers	internal
informix	Informix relational database management system	add-on
ircd	Internet Relay Chat (IRC) daemon	internal
ircu	Internet Relay Chat (IRC) server often used in customer-service chat applications	add-on

Keysight, EX-1004
Page 2032

Detected Services (Continued)

Service Detector	Description	Detector Type
jrmi	Java Remote Method Invocation (Java-based remote procedure call service)	add-on
kerberos	Kerberos protocol (TCP)	well-known port
kerberos	Kerberos protocol (UDP)	add-on
ldap	Lightweight Directory Access Protocol (LDAP) services	add-on
linuxconf	Linuxconf services	well-known port
login	login services (login, rlogin)	internal
lotusnotes	Lotus Notes (collaborative communication)	add-on
mq	Websphere Message Queue (MQ) distributed environment interprocess communications	add-on
MSN Messenger	MSN Messenger	implied from application
mysql	MySQL database services	internal
nameserver	Windows Internet Name Service (WINS)	well-known port
nessus	Nessus services	well-known port
netbios-dgm netbios-ns netbios-ssn	NETBIOS Datagram, Name, and Session services	internal
nntp	Network News Transfer Protocol	internal
ntp	Network Time Protocol	internal
oracle	Oracle Transparent Network Substrate (TNS) database connectivity services	internal
PC-Duo	Remote control software	add-on

Keysight, EX-1004
Page 2033

Detected Services (Continued)

Service Detector	Description	Detector Type
pop3	Post Office Protocol	internal
postgresql	PostgreSQL (object-relational database management system, or ORDBMS) services	well-known port
printer	Printer services	internal
radius radacct	Remote Authentication Dial-In User Service (RADIUS) client services	internal
rdp	Remote Desktop Protocol	internal
rsync	file synchronization services	internal
RTMP	Real Time Messaging Protocol	add-on
rtp	Real-time Transport Protocol	internal
rtsp	Real Time Streaming Protocol	well-known port
secondlife	Second Life (Linden Lab)	add-on
shell	shell services (shell, rshell)	internal
sip	Session Initiation Protocol (SIP) used for VoIP technology	internal
Skype	Skype	implied from application
smtp	Simple Mail Transport Protocol	internal
snmp snmp-trap	Simple Network Management Protocol (SNMP) services	internal
socks	SOCKS client-server proxy routing protocol	well-known port
SQL Server	Tabular Data Steam (TDS) protocol	add-on
Squid	Squid proxy server and web cache daemon	internal

Keysight, EX-1004
Page 2034

Detected Services (Continued)

Service Detector	Description	Detector Type
ssc	Symantec System Center (anti-virus software management)	add-on
ssh	Secure Shell (SSH) services	internal
ssl	Secure Sockets Layer (SSL) services	internal
STUN	Session Traversal Utilities for NAT	add-on
sunrpc	Open Network Computing Remote Procedure Call services (ONC RPC), which may include ToolTalk Server, cmsd, status, alis, rexd, mountd, nfs, portmap	internal
syslog	syslog client-server logging protocol	well-known port
talk ntalk	live text communication services	well-known port
telnet	RFC 854 Telnet services	internal
tftp	Trivial File Transfer Protocol (TFTP) services	internal
Timbuktu	Timbuktu remote desktop services	internal
tivoli	Tivoli (system management)	add-on
traceroute	traceroute services	well-known port
uc4	workload automation	add-on
unicenter	Unicenter workflow automation	add-on
vcom	virtual serial port driver	add-on
vnc vnc-server	Remote Frame Buffer (RFB) services, which may include remote desktop applications such as VNC or Bochs	internal
WebLogic	Oracle WebLogic	implied from application

Keysight, EX-1004
Page 2035

Detected Services (Continued)

Service Detector	Description	Detector Type
whois	WHOIS query-response protocol	well-known port
wx	WAN acceleration	add-on
x11	X Window System services	well-known port
Yahoo! Messenger	Yahoo! Messenger	implied from application

Detected Client Applications

Platform: DC

Component: RNA

RNA detects the client applications listed in the the Detected Client Applications table on page 2037. For each application detector, the table also lists its detector type.

There are two types of application detectors:

- internal application detectors, which are always on
- add-on application detectors, which are distributed out of band from the Sourcefire 3D System update path so that you do not have to install a full patch or upgrade to use new detectors, and which you can activate and deactivate

Note that the client applications that RNA detects depend on both the version of the Sourcefire 3D System and the version of the VDB you have installed. 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 RNA detectors. To further enhance RNA's detection capabilities, you can also import individual add-on client application detectors provided by Sourcefire Professional Services.

Keysight, EX-1004
Page 2036

For more information on working with application detectors, see Working with Service and Application Detectors on page 1381.

Detected Client Applications

Application Type	Client Applications	Detector Type
database	Oracle Database	internal
	Oracle WebLogic	add-on
	Tabular Data Steam (TDS)	add-on
desktop search	Google Desktop	internal
email	AOL software	internal
	Apple Mail	internal
	Eudora and Eudora Pro	internal
	Evolution	internal
	KMail	internal
	Lotus Notes	internal
	Outlook and Outlook Express	internal
	Mutt	internal
	SquirrelMail	internal
	Thunderbird	internal
entertainment	Google Earth	internal
	Station Launcher	internal
	World of Warcraft	internal
instant messenger	AOL Instant Messenger	internal
	ICQ	internal
	Microsoft MSN Messenger	internal
	Skype	internal
	Yahoo! Messenger	internal
media player	Windows Media Player	internal

Keysight, EX-1004
Page 2037

Detected Client Applications (Continued)

Application Type	Client Applications	Detector Type
mobile browser	Android browser	internal
	BlackBerry browser	internal
	Mobile Safari	internal
peer to peer	BitTorrent client and tracker	internal
	FastTrack	internal
	FlashGet	internal
	FreeCast	internal
	Gnutella	internal
	IceShare	internal
	Kazaa	internal
	Kugoo	internal
	PeerEnabler	internal
	PPLive	internal
	SoulSeek	internal
	Vuze	internal
	Yet ABC	internal
photo/video sharing	Picasa	internal

Keysight, EX-1004
Page 2038

Detected Client Applications (Continued)

Application Type	Client Applications	Detector Type
remote desktop	PC-Duo	add-on
	remote desktop client	internal
	Timbuktu	internal
	vnc	internal
Secure Shell (SSH)	Dropbear	internal
	lsh	internal
	OpenSSH	internal
	PuTTY	internal
VoIP	VoIP (RTP & SIP)	internal
	Yahoo! Voice	internal
web browser	Chrome	internal
	cURL	internal
	Firefox 2.x - 3.x	internal
	Konqueror	internal
	Internet Explorer 4.01 - 7.0	internal
	Safari	internal
	Wget	internal

Detected Payloads

Platform: DC

Component: RNA

RNA detects the payloads (that is, the content that users are accessing with their web browsers) listed in the Detected Payloads table.

Note that the payloads that RNA detects depend on both the version of the Sourcefire 3D System and the version of the VDB you have installed. You should

Keysight, EX-1004
Page 2039

carefully read the release notes for each Sourcefire 3D System update as well as the advisories for each VDB update for information on updated payload detectors.

Detected Payloads

Payload Type	Payloads
audio	MPEG WAV WMA generic (undetermined)
browser add-on	Google Toolbar Yahoo! Toolbar
collaboration	Zoho Wiki
content delivery	Limelight
file sharing and storage	4shared ADrive Megaupload RapidShare YouSendIt Zoho Share

Keysight, EX-1004
Page 2040

Detected Payloads (Continued)

Payload Type	Payloads
media streaming	Dailymotion Hulu Live365 Livestation Metacafe Move Networks MSN Napster Pandora PPStream SHOUTcast Radio Slingbox Soribada Ustream.tv XM Radio Online YouTube
online commerce	eBay Bid eBay Search eBay Watch
online conferencing	WebEx
online educational tool	Blackboard
online office application	Backpack
online scheduling	Google Calendar
real-time collaboration	WebDAV

Keysight, EX-1004
Page 2041

Detected Payloads (Continued)

Payload Type	Payloads
social bookmarking	Digg
	Fark
	Reddit
	Slashdot
	StumbleUpon
social networking	Bebo
	Blogger
	Facebook
	Friendster
	LinkedIn
	Livejournal
	Mixi
	MySpace
	Orkut
	Twitter
software update	Adobe Updater
	Microsoft Update
syndication	Atom
	RSS
tracking and statistics	Google Analytics
video	Flash Video
	MP4
	MPEG
	Shockwave
	QuickTime
	WMV
	generic (undetermined)

Keysight, EX-1004
Page 2042

Detected Payloads (Continued)

Payload Type	Payloads
web chat	AIM Express eBuddy ICQ2Go MySpace Chat Zoho Chat
website reputation	Google Safe Browsing
web games	PartyPoker Pogo Yahoo! Games
webmail	AOL Mail Gmail Verizon Email Windows Live Hotmail Yahoo! Mail Zoho Mail

Keysight, EX-1004
Page 2043

Appendix E
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.

There are numerous CLI modes. 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 account 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 login to the shell.

On 8000 Series sensors, you can assign command line permissions through the User Management page in the web interface. See Managing Users on page 1626

Keysight, EX-1004
Page 2044

for more information. On Virtual 3D Sensor, command line permissions are assigned through the CLI.

For information about logging into the command line, see Logging into the Appliance on page 37.

The following sections describe the CLI commands.

- Basic CLI Commands on page 2045
- Show Commands on page 2049
- Configuration Commands on page 2061

Basic CLI Commands

The basic CLI commands provide the ability to interact with the CLI. These commands do not affect the operation of the sensor. Basic commands are available to all CLI users.

The following sections describe the basic commands:

- configure password on page 2045
- end on page 2046
- exit on page 2046
- help on page 2046
- history on page 2047
- logout on page 2047
- ? (question mark) on page 2048
- ?? (double question marks) on page 2048

configure password

Initiates a password change for the current user. After issuing the command, the CLI prompts the user for their current (or old) password and then prompts the user to enter the new password twice.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
configure password
```


Keysight, EX-1004
Page 2045

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.)

Platforms

8000 Series Sensor, Virtual 3D Sensor

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.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
exit
```

Example

```
configure network ipv4> exit
configure network>
```

help

Displays an overview of the CLI syntax.

Keysight, EX-1004
Page 2046

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
help
```

Example

```
> help
```

history

Displays the command line history for the current session.

Platforms

8000 Series Sensor, Virtual 3D Sensor

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.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
logout
```

Example

```
> logout
```


Keysight, EX-1004
Page 2047

? (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.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
?
abbreviated_command ?
command [arguments] ?
```

Example

```
> ?
```

?? (double question marks)

Displays detailed context-sensitive help for CLI commands and parameters.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
??
abbreviated_command end??
command [arguments] ??
```

Example

```
> configure manager add ??
```


Keysight, EX-1004
Page 2048

Show Commands

Show commands provide information about the state of the sensor. These commands do not change the operational mode of the sensor 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` and `show users` commands.

The following sections describe the show command:

- alarms on page 2050
- clustering on page 2050
- cpu on page 2050
- detection-engines on page 2052
- disk on page 2052
- dns on page 2052
- failopen on page 2053
- gui on page 2053
- hostname on page 2053
- hyperthreading on page 2054
- interfaces on page 2054
- interface-sets on page 2054
- lcd on page 2055
- link-state on page 2055
- managers on page 2055
- memory on page 2056
- model on page 2056
- network on page 2056
- network-modules on page 2057
- ntp on page 2057
- portstats on page 2057
- processes on page 2058
- serial-number on page 2058
- summary on page 2059
- time on page 2059
- user on page 2059
- users on page 2060
- version on page 2061

Keysight, EX-1004
Page 2049

alarms

Displays the currently active (failed/down) hardware alarms on the 8000 Series sensors.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show alarms
```

Example

```
> show alarms
```

clustering

Shows the clustering configuration on 8000 Series sensors. The user must use the GUI to enable or disable clustering.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show clustering
```

Example

```
> show clustering
```

cpu

Displays the current CPU usage statistics appropriate for the platform for all CPUs on the sensor. For 8000 Series sensors, 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 3D Sensors, the following values are displayed:

- CPU

  Processor number

Keysight, EX-1004
Page 2050

- %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.

- %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. A softirq (software interrupt) is one of up to 32 enumerated software interrupts that can run on multiple CPUs at once.

- %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.

Platforms

8000 Series Sensor, Virtual 3D Sensor

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 an 8000 Series sensor, it is ignored because for this platform, utilization information can only be displayed for all processors.

Keysight, EX-1004
Page 2051

Example

```
> show cpu
```

detection-engines

Displays information about each of the detection engines configured in the sensor.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show detection-engines
```

Example

```
> show detection-engines
```

disk

Displays the current disk usage.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show disk
```

Example

```
> show disk
```

dns

Displays the current DNS server addresses and search domains.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show dns
```


Keysight, EX-1004
Page 2052

Example

```
> show dns
```

failopen

Lists the fail-open pairs in use and shows the status of those pairs, either normal or fail-open.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show failopen
```

Example

```
> show failopen
```

gui

Displays the current state of the web interface.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show gui
```

Example

```
> show gui
```

hostname

Displays the sensor's hostname and appliance UUID.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show hostname
```


Keysight, EX-1004
Page 2053

Example

```
> show hostname
```

hyperthreading

Displays the current value of the hyperthreading configuration.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show hyperthreading
```

Example

```
> show hyperthreading
```

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.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show interfaces [interface]
```

where *interface* is the specific interface for which you want the detailed information.

Example

```
> show interfaces
```

interface-sets

Displays data regarding interface sets configured on the sensor.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Keysight, EX-1004
Page 2054

Syntax

```
show interface-sets
```

Example

```
> show interface-sets
```

lcd

Displays whether the LCD hardware display is enabled or disabled.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show lcd
```

Example

```
> show lcd
```

link-state

Displays type, link, speed, duplex, and failopen state of the ports on the sensor.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show link-state
```

Example

```
> show link-state
```

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 sensor is registered to a high availability (HA) pair, information about both managing Defense Centers is displayed.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Keysight, EX-1004
Page 2055

Access

Basic

Syntax

```
show managers
```

Example

```
> show managers
```

memory

Displays the total memory, the memory in use, and the available memory for the sensor.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show memory
```

Example

```
> show memory
```

model

Displays the model information for the sensor.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show model
```

Example

```
> show model
```

network

Displays the IPv4 and IPv6 configuration of the management interface.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Keysight, EX-1004
Page 2056

Access

Basic

Syntax

```
show network
```

Example

```
> show network
```

network-modules

Displays all the installed modules and information about them.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show network-modules
```

Example

```
> show network-modules
```

ntp

Displays the ntp configuration.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show ntp
```

Example

```
> show ntp
```

portstats

Displays port statistics for all the installed ports on the 8000 Series sensor.

Platforms

8000 Series Sensor

Access

Keysight, EX-1004
Page 2057

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
```

processes

Displays the processes currently running on the sensor sorted by descending CPU usage.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show processes [flag] [filter]
```

where 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 to look for in the command or username. The header row is still displayed.

Example

```
> show processes
```

serial-number

Displays the chassis serial number.

Platforms

8000 Series Sensor

Access

Basic

Syntax

```
show serial-number
```

Example

```
> show serial-number
```


Keysight, EX-1004
Page 2058

summary

Displays a summary of the most commonly used information (version, interface, and detection-engine) about the sensor. For more detailed information, see show version on page 2061, show interfaces on page 2054, and detection-engines on page 2052.

Platforms

8000 Series Sensor, Virtual 3D Sensor

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.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show time
```

Example

```
> show time
```

user

Displays detailed configuration information for the specified users. The following values are displayed.

- Login, indicates the login name
- UID, indicates the numeric user ID
- Auth (Local or Remote), indicates how the user is authenticated
- Access (Basic or Config), indicates the user's privilege level
- Enabled (Enabled or Disabled), indicates whether the user is active
- Reset (Yes or No), indicates whether the user must change password at next login

Keysight, EX-1004
Page 2059

- Exp (Never or a number), indicates the number of days until the user's password must be changed
- Warn (N/A or a number), indicates the number of days a user is given to change their password before it expires
- Str (Yes or No), indicates whether the user's password must meet strength checking criteria
- Lock (Yes or No), indicates whether the user's account has been locked due to too many login failures
- Max (N/A or a number), indicates the maximum number of failed logins before the user's account is locked

Platforms

Virtual 3D Sensor

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

Displays detailed configuration information for all the local users.The following values are displayed.

- Login, indicates the login name
- UID, indicates the numeric user ID
- Auth (Local or Remote), indicates how the user is authenticated
- Access (Basic or Config), indicates the user's privilege level
- Enabled (Enabled or Disabled), indicates whether the user is active
- Reset (Yes or No), indicates whether the user must change password at next login
- Exp (Never or a number), indicates the number of days until the user's password must be changed
- Warn (N/A or a number), indicates the number of days a user is given to change their password before it expires.
- Str (Yes or No), indicates whether the user's password must meet strength checking criteria

Keysight, EX-1004
Page 2060

- Lock (Yes or No), indicates whether the user's account has been locked due to too many login failures
- Max (N/A or a number), indicates the maximum number of failed logins before the user's account is locked

Platforms

Virtual 3D Sensor

Access

Configuration

Syntax

```
show users
```

Example

```
> show users
```

version

Displays the product version and build. If the *detail* parameter is specified, displays the version of additional components.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Basic

Syntax

```
show version [detail]
```

Example

```
> show version
```

Configuration Commands

The configuration commands enable the user to configure and manage the system. These commands affect system operation, therefore, only users with configuration CLI access can issue these commands.

The following sections describe the configuration commands:

- gui on page 2062
- hyperthreading on page 2062
- lcd on page 2062
- manager on page 2063
- network on page 2064

Keysight, EX-1004
Page 2061

- user on page 2066
- system reboot on page 2071
- system restart on page 2071
- system shutdown on page 2072

gui

Enables or disables the Sensor web interface, including the Upgrade web interface.

Platforms

8000 Series Sensor

Access

Configuration

Syntax

```
configure gui {enable | disable}
```

Example

```
> configure gui disable
```

lcd

Enables or disables the LCD hardware display on the front of the sensor.

Platforms

8000 Series Sensor

Access

Configuration

Syntax

```
configure lcd {enable | disable}
```

Example

```
> configure lcd disable
```

hyperthreading

Turns hyperthreading on or off on 8000 Series Sensors.

Platforms

8000 Series Sensor

Access

Configuration

Keysight, EX-1004
Page 2062

Syntax

```
configure hyperthreading {enable | disable}
```

Example

```
> configure hyperthreading enable
```

manager

The configure manager commands configure the sensor's connection to its Defense Center.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Configuration

add

Configures the sensor to accept a connection from a managing Defense Center. This command works only if the sensor is not actively managed.

A unique alphanumeric registration key is always required to register a sensor to a Defense Center. In most cases, you must provide the hostname or the IP address along with the registration key. However, if the sensor and the Defense Center are separated by a NAT device, you must enter a unique NAT ID, along with the 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 sensor. 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 sensor to the Defense Center. *`nat_id`* is an optional alphanumeric string used during the registration process between the Defense Center and the sensor. It is required if the hostname is set to `DONTRESOLVE.`

Example

```
> configure manager add DONTRESOLVE abc123 efg456
```

delete

Removes the Defense Center connection information from the sensor. This command only works if the sensor is not actively managed.

Keysight, EX-1004
Page 2063

Syntax

```
configure manager delete
```

Example

```
> configure manager delete
```

network

The configure network commands configure the sensor's management interface.

Platforms

8000 Series Sensor, Virtual 3D Sensor

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 current list of DNS server with the list specified in the command.

Syntax

```
configure nework 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 sensor appliance.

Syntax

```
configure network hostname {name}
```

where *name* is the new hostname.

Example

```
> configure network hostname sfrocks
```


Keysight, EX-1004
Page 2064

ipv4 delete

Disables the IPv4 configuration of the sensor's management interface.

Syntax

```
configure network ipv4 delete
```

Example

```
> configure network ipv4 delete
```

ipv4 dhcp

Sets the IPv4 configuration of the sensor'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 sensor'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 sensor's management interface.

Syntax

```
configure network ipv6 delete
```

Example

```
> configure network ipv6 delete
```

ipv6 dhcp

Sets the IPv6 configuration of the sensor's management interface to DHCP. The management interface communicates with the DHCP server to obtain its configuration information.

Keysight, EX-1004
Page 2065

Syntax

```
configure network ipv6 dhcp
```

Example

```
> configure network ipv6 dhcp
```

ipv6 router

Sets the IPv6 configuration of the sensor'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 sensor's management interface.

Syntax

```
configure network ipv6 manual ip6addr/ip6prefix [ip6gw]
```

where *ip6addr/ip6prefix* is the IP address and subnet mask specified using CIDR notation and ip6*gw* is the IPv6 address of the default gateway.

Example

```
> configure network ipv6 manual
2001:DB8:3ffe:1900:4545:3:200:f8ff:fe21:67cf 64
```

user

The configure user commands manage the sensor's local user database.

Platforms

Virtual 3D Sensor

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
```


Keysight, EX-1004
Page 2067

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.

Keysight, EX-1004
Page 2068

Syntax

```
configure user password username
```

where *username* specifies the name of the user.

Example

```
> configure user pasword 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 users 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 file

The system file commands enable the user to manage the files in the common directory on the sensor.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Configuration

Keysight, EX-1004
Page 2069

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 and 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 and the file names are space-separated.

Example

```
> system file delete *
```

list

If no filenames 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 filenames.

Syntax

```
system file list {filenames filenames ...}
```

where filenames specifies the files to display and 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 ...
```


Keysight, EX-1004
Page 2070

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 and the file names are space-separated.

Example

```
> system file secure-copy 10.123.31.1 jdoe /tmp *
```

system generate-troubleshoot

Generates troubleshooting data for analysis by Sourcefire.

Platforms

8000 Series Sensor, Virtual 3D Sensor

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
```

system reboot

Reboots the sensor appliance.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Configuration

Syntax

```
system reboot
```

Example

```
> system reboot
```

system restart

Restarts the sensor application.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Keysight, EX-1004
Page 2071

Access

Configuration

Syntax

```
system restart
```

Example

```
> system restart
```

system shutdown

Shuts down the sensor.

Platforms

8000 Series Sensor, Virtual 3D Sensor

Access

Configuration

Syntax

```
system shutdown
```

Example

```
> system shutdown
```


Keysight, EX-1004
Page 2072

Appendix F
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 2.2.15	Apache License
Linux Kernel 2.6.32.24	GNU General Public License Version 2 (GPLv2)
Perl 5.10.1 and related modules	Perl Artistic License
SNORT 2.9	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:

```
/user/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.

Keysight, EX-1004
Page 2073

Glossary

3D Sensor An appliance-based sensor that, as part of the Sourcefire 3D System, can run the IPS component, the RNA component, the RUA component, or combinations of the components.

active detection The addition, to the network map, of data collected by active sources, such as host operating system and service information.

adaptive profile An intrusion policy profile that uses information from RNA host profiles to determine the operating system for the target host of a packet. Profiles within an intrusion policy then automatically adapt to cause the 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.

Administrator A type of user role that conveys rights to all Sourcefire 3D System functionality. Administrators can set up an appliance's network configuration, manage user accounts, and configure system policies and system settings. Users with the Administrator role also have the access rights provided to the Intrusion Event Analyst, RNA Event Analyst, Policy & Response Administrator, and Maintenance User roles.

advanced setting A preprocessor or other IPS component feature that requires specific expertise to configure. Advanced settings typically require little or no modification and are not common to every deployment. You enable and disable advanced settings and access their configuration pages from the Advanced Settings page.

Keysight, EX-1004
Page 2074

alert A message that notifies you when an intrusion event, health event, host input event, RNA event, RUA event, white list event, or compliance event is generated. You can send alerts to an external syslog server, a specific email address, or an SNMP trap server. See email alerting, SNMP alerting, and syslog alerting.

alert rule An intrusion rule that, when triggered, generates an intrusion event and logs the details of the packet that triggered the rule. Compare with pass rule and drop rule.

anomaly detection The detection of anomalous conditions in traffic rate or traffic content that indicates an attack.

audit log A record of user interactions with the web interface. The audit log comprises audit events.

audit event An event that describes a specific 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. You can view audit events in the audit log.

banner The first 256 bytes of the first packet detected by a service. A banner is collected only once, the first time a service is detected by RNA. Banners provide additional context to the information gathered by RNA.

base policy A selectable set of configurations that can be any one of the default intrusion policies provided by Sourcefire or a custom user layer.

base policy layer A built-in layer in an intrusion policy comprised of all of the default settings for the IPS component. The base policy selected for the intrusion policy determines the default settings in the base policy layer.

bit mask The notation used to identify which bits in an IP address correspond to the network address and subnet portions of the address.

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 range you selected, and any columns you disabled as well as any constraints you imposed. The bookmarks you create are available to all users with unrestricted analyst access.

bridge A network device that forwards traffic between network segments. RNA identifies bridges as network devices that communicate using Cisco Discovery Protocol (CDP) or Spanning Tree Protocol (STP). RNA may identify switches as bridges.

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 RNA Recommendations layer.

Keysight, EX-1004
Page 2075

Classless Inter-Domain Routing (CIDR) notation

A notation that defines IP address ranges by combining an IP address with a bit mask that signifies the subnet mask used to define the number of IP addresses in the specified range. For example, if you want to define the network described by 192.168.1.x with a subnet mask of 255.255.255.0, use 192.168.1.1/24, where 24 signifies the number of bits in the subnet mask.

client application

An application that runs on one host and relies on another host (a server) to perform some operation. For example, email clients are client applications that allow you to send and receive email. When RNA detects that a user on a host is using a specific client application 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 application.

client application event

Information that describes client application activity on monitored hosts. For each detected client application, RNA logs the IP address that used the application and when the application was last used, as well as the application name, version, and the number of times its use was detected.

clipboard

A holding area where you can copy up to 25,000 intrusion events that you can later add to incidents. The contents of the clipboard are sorted by the date and time that the events were generated.

clustering

A feature that allows you to increase the amount of traffic inspected on a network segment by connecting two fiber-based 3D9900 sensors in a clustered pair. When you establish a clustered pair configuration, you combine the 3D9900 sensors resources into a single, shared configuration.

complex condition

A complex way of qualifying compliance rules, flow trackers, host profile qualifications, and traffic profiles. A complex condition comprises at least two simple conditions, linked to each other with an AND or an OR operator.

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 event

An event generated by the Defense Center when a compliance rule triggers. You can search, view, and delete compliance events and can configure the number of compliance events saved in the database. Note that white list events, generated by white list violations, are a special kind of compliance event.

compliance policy

Describes the network activity that constitutes a security policy violation, using compliance rules and compliance white lists. You can specify responses to each rule or white list within a policy.

compliance rule

Along with compliance white lists, one of the ways you can specify criteria that network traffic must meet in order to violate a compliance policy. You can use the Defense Center to configure compliance rules to trigger (and generate a compliance event) when a specific intrusion event, RNA event, or flow event occurs, or when your network traffic deviates from your normal network traffic

Keysight, EX-1004
Page 2076

pattern as characterized in a traffic profile. You can constrain compliance rules with host profile qualifications, flow 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 compliance rule triggers.

compliance white list

Along with compliance rules, one of the ways you can specify criteria that network traffic must meet in order to violate a compliance policy. You can use the Defense Center to configure compliance white lists to specify which operating systems, services, client 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 a compliance white list is **not** associated with the white list of IP addresses that you can configure in certain remediations.

compliance white list event

See white list event.

compliance white list violation

See white list violation.

current identity

The operating system or service identity that RNA finds most likely to be correct, which is used to assign host vulnerability, to assess impact of an attack, to evaluate compliance rules written against operating system identifications, host profile qualifications, and compliance white lists, to display in the Hosts and Services table views in workflows and in the host profile, and to calculate the operating system and service statistics on the RNA Statistics page.

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 flow data table to examine flow data in a new context. Custom tables include Sourcefire-defined custom tables, which are custom tables delivered with the Defense Center.

custom workflow

A workflow that you create to meet the unique needs of your organization. Compare with predefined workflow and, on the Defense Center, saved custom workflow.

dashboard

A display that provides tabs of at-a-glance information about many aspects of the performance on your Sourcefire 3D System. You can configure as many dashboards as you need and decide which dashboard widgets appear on each tab to fit your system monitoring needs. The dashboard appears as the default home page for all user roles except the restricted event analyst roles.

Keysight, EX-1004
Page 2077

dashboard widget A dashboard widget provides status or performance information about a specific aspect of your Sourcefire 3D System. You can select which widgets to add to your dashboard.

data correlator A program that generates events and creates the network map on the Defense Center, using the data collected by RNA.

decoder A component of IPS that places sniffed packets into a format that can be understood by a preprocessor.

Defense Center A central management point that allows you to manage sensors and automatically aggregate the events they generate. You can also push policies created on the Defense Center and software updates to managed sensors. If you manage 3D Sensors with IPS and RNA with a Defense Center, the Defense Center correlates intrusion events with host vulnerabilities and assigns impact flags to the intrusion events. Impact correlation lets you focus on attacks most likely to affect high-priority hosts. The Defense Center also correlates intrusion information with user identity data from the RUA database.

defragmentation policy Describes how the IP defragmentation preprocessor (a component of IPS) 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 RNA 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.

detection engine The mechanism that is responsible for analyzing the traffic on the network segment where a sensor is connected. A detection engine has two main components: an interface set and a detection resource. RNA uses RNA detection engines, IPS uses IPS detection engines, and RUA uses RUA detection engines.

detection policy See RNA detection policy.

detection resource A portion of a sensor's computing resources used as part of a detection engine.

DNS cache Temporary storage of previously resolved IP addresses. Configuring DNS caching allows you to resolve those IP addresses without performing additional lookups. This can reduce the amount of traffic on your network and speed the display of event pages.

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.

Keysight, EX-1004
Page 2078

drop event An intrusion event generated when a drop rule triggers. On RNA compliance event views and IPS intrusion event views, 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, IPS drops the packet and generates an intrusion event (specifically, a drop event). Compare with alert rule and pass rule.

drop when inline An intrusion policy setting that determines whether IPS drops packets that trigger drop rules in an inline deployment.

dynamic rule state A rule state that is set for a specified period of time in response to a detected rate anomaly in traffic matching the rule.

email alerting The transmission of an alert as an email message.

eStreamer See Event Streamer.

event Information that is stored as an event. An event contains multiple fields that describe the activity that caused the event to be generated. IPS generates intrusion events, which also include drop events and preprocessor events. RNA generates network discovery events and flow events, as well as events that provide general information about your network topology: client application events, host events, host attributes, and service events. A vulnerability is also considered an RNA event. You can use the policy and response feature to configure your Defense Center to generate compliance events and white list events, as well as remediation status events. RUA generates RUA events when it detects user logins or user additions or deletions. In addition, every appliance generates records of user activity called audit events. The health monitor on the Defense Center also generates health events.

event analyst An event analyst examines event data collected by the Sourcefire 3D System. The Intrusion Event Analyst, RNA Event Analyst, and Restricted Event Analyst roles, or their read-only counterparts, can be assigned to a user to provide access to event analysis functionality.

Event Streamer Also known as eStreamer, a component of the Sourcefire 3D System that allows you to stream event data from a Defense Center or 3D Sensor to external client applications.

event suppression A feature that allows you to use suppress intrusion events when a specific IP address or range of IP addresses triggers a rule. Event suppression 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 events for the rules that are triggered by your mail server, so that you only see the events for legitimate attacks.

Keysight, EX-1004
Page 2079

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 view A workflow view containing a set of events. You can constrain the events included in an event view using an event search or using simple constraints or complex constraints.

export A method that you can use to transfer various configurations from appliance to appliance. You can export intrusion policies, RNA detection policies, system policies, health policies, dashboards, custom workflows and tables, and some RNA detectors. 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.

fail-open card A network interface card that allows network traffic to pass through a 3D Sensor that uses IPS detection engines that are deployed inline, even if the appliance itself fails or loses power.

feature license A license you can add to an appliance that enables additional features, including NetFlow, Intrusion Agents, Sourcefire 3D Sensor Software for X-Series, Sourcefire Virtual 3D Sensors, and the ability to monitor a number of hosts with RNA or users with RUA.

fingerprint An established definition that RNA compares against specific packet header values and other unique data from network traffic to identify a host's operating system. If RNA misidentifies or cannot identify a host's operating system, you can create a custom fingerprint that identifies the host.

flow data See flow event.

flow event An event generated when RNA detects that a connection between a monitored host and any other host is terminated. Flow events include information about the collected traffic, including the first packet of the transaction, the last packet of the transaction, the source IP address and port, the destination IP address and port, the number of packets and bytes sent and received by the monitored host, and the client application and URL involved in the transaction, if applicable.

flow summary Flow data aggregated over a five-minute interval. You can choose to store flow data only as flow summaries to save disk space.

Keysight, EX-1004
Page 2080

flow tracker One or more conditions that constrain a compliance rule so that after the rule's initial criteria are met, RNA begins tracking certain flows. The rule then triggers only if the tracked flows meet additional criteria.

gateway A device that acts as an entrance to and controls traffic within your organization's network. When you set up your 3D Sensor or Defense Center, you must specify the IP address of the gateway device for your network.

GID (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.

health alert An alert generated by the Defense Center or Master Defense Center when a specific health event occurs.

health event An event that is generated when one of the appliances in your deployment meets (or fails to meet) performance criteria specified in a health module. Health events indicate which module triggered the event and when the event was triggered.

health monitor A feature that continuously monitors the performance of the appliances in your deployment. The health monitor uses health modules to test various performance aspects of the appliances. You configure the health monitor using a health policy.

health monitor blacklist A blacklist that temporarily disables aspects of health monitoring to prevent the Defense Center from generating unnecessary health events. You can disable monitoring for a group of appliances, a single appliance, or a specific health module.

health module A test of a particular performance aspect of one of the appliances in your deployment. For example, you can monitor CPU usage or available disk space. You can configure health modules to generate health events and health alerts when the performance aspects they monitor reach a certain level.

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. The Defense Center and Master Defense Center are delivered with default health policies; you can modify them or create your own.

high availability A feature that allows you to designate redundant Defense Centers to manage groups of sensors. Event data streams from managed sensors to both Defense Centers and certain 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.

hop The trip a packet takes from one router or intermediate point to another in the network. RNA detects the number of network hops that exist between the

Keysight, EX-1004
Page 2081

sensors and the hosts they monitor, which provides you with information about the physical location the hosts on your network.

host A device that is connected to a network and has a unique IP address. To RNA, a host is any identified host that is not categorized as a bridge, router, NAT device, or load balancer.

host attribute A tool you can use to provide information about hosts detected by RNA and to classify them in ways that are important to your network environment. For example, you could create a host attribute that designates the physical location of each host on your network. You can use and configure the two predefined host attributes, host criticality and notes, as well as create your own host attributes. In addition, when you create a compliance white list, RNA automatically creates a host attribute that indicates the compliance of the host. You can use host attributes in compliance rules and compliance white lists, and you can search for hosts with specific host attribute values. You also can generate reports based on host attributes.

host criticality A host attribute that indicates the business criticality (importance) of any given host detected by RNA. You can use host criticality values when searching for hosts or when creating compliance rules and compliance white lists.

host event An event indicating that RNA has detected a host. RNA collects information about the hosts on monitored network segments. The information that RNA collects comprises that host's host profile.

host import input data Host input data imported using a command line utility or the host input API.

host input A feature that allows you to import host data from third-party applications to augment the information in the RNA network map using scripts or command-line files. You can also use the host input feature through the web interface to modify operating system or service identities or deleting services, protocols, host attributes, or client applications.

host input event An event that is created when a change is made to your network map using the host input feature.

host profile Collected information about a specific host detected by RNA. This includes general host information, such as its name and operating system, as well as its user history, host attributes, the protocols it uses, the services it is running, VLAN information, client applications running on the host, applicable white list violations, detected vulnerabilities, and any scan results for that host.

host profile qualification A constraint placed on a traffic profile or compliance rule. A host profile qualification within a compliance rule specifies that the Defense Center should generate a compliance event only if the host involved meets certain criteria. A host profile qualification within a traffic profile limits the hosts that are profiled.

Keysight, EX-1004
Page 2082

host 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.

host view The final page in workflows based on RNA events (with the exception of workflows based on vulnerabilities, which use the vulnerability detail page). The host view displays the host profiles of the hosts involved in the events you are viewing.

HTTP Inspection A preprocessor that decodes and normalizes URI data sent to and received from web servers on your network, detects and generates events against possible URI-encoding attacks, and makes the normalized data available for additional rule processing. This is important because HTTP traffic can be encoded in a variety of formats, making it difficult for IPS to inspect packets accurately.

identity conflict A conflict event that occurs when RNA reports a new passive identity that conflicts with the current active identity and previously-reported passive identities.

impact The qualification of each intrusion event on a Defense Center based on whether RNA deployed on the same segment detected a vulnerable service or open port on the target of the attack. If the targeted host is not vulnerable, the impact of the attack is low. However, if the targeted host is vulnerable, then the impact is high and you should act to mitigate the effects of the attack.

impact flag For intrusion events, an indicator of the correlation between intrusion data, RNA network discovery events, and vulnerability information. A red impact flag means that the host is *vulnerable* to the attack represented by the intrusion event, orange means it is *potentially vulnerable*, and so on. Intrusion events detected on network segments not monitored by RNA have gray impact flags; this 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 intrusion policies, RNA detection policies, system policies, health policies, dashboards, custom workflows and tables, and RNA detectors that you previously exported from another appliance of the same type.

inactive period An interval during which a compliance rule does not trigger. You can configure an inactive period to occur daily, weekly, or monthly and to begin at a specified time and last a specified number of minutes. For example, you might perform a nightly Nmap scan on your internal network to update host profile information. In that case, you could set a daily inactive period on the affected compliance rules for the time and duration of your scan so those rules do not trigger erroneously. See also snooze period.

incident One or more intrusion events that you suspect are involved in a possible violation of your security policy. The Sourcefire 3D System provides incident-handling

Keysight, EX-1004
Page 2083

features that you can use to collect and process information that is relevant to your investigation of the incident.

inline A type of interface set that allows you to deploy a 3D Sensor inline on a network. In this configuration, the IPS component can affect the traffic flow on the monitored network, including dropping malicious packets.

inline with fail open A type of interface set that allows you to use a compatible fail-open card that allows network traffic to continue flowing if the appliance fails for any reason.

interface set One or more sensing interfaces on a 3D Sensor that you can use to monitor network segments for one or more detection engines. You can use passive, inline, or inline with fail open interface sets.

internal authentication An authentication method that stores user credentials in a local database. 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 Agent Software that can be installed on certain Red Hat Linux, FreeBSD or Sun Solaris servers to transmit intrusion events generated by Snort to the Defense Center. You can use the Defense Center to aggregate event information from Intrusion Agents with data from 3D Sensors with IPS.

intrusion event A record of the network traffic that violated an intrusion policy. Intrusion event data includes the date, time, and the type of exploit, as well as other contextual information about the source of the attack and its target.

intrusion policy A variety of components that you can configure to inspect your network traffic for intrusions and 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; RNA 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.

intrusion rule A set of keywords and arguments that, when applied to captured network traffic, identify potential intrusions, policy violations, and security breaches. IPS compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in the rule, the rule triggers and generates an intrusion event. Intrusion rules include alert rules, drop rules, and pass rules.

IP address A 32-bit (IPv4) or 128-bit (IPv6) number, usually represented in dot notation (for example, 192.168.34.166), that identifies the host that sends or receives packets on the Internet or on the local network.

Keysight, EX-1004
Page 2084

IPS

A component of the Sourcefire 3D System, separately licensable on 3D Sensors, that provides intrusion detection and prevention capabilities. If you configure a 3D Sensor with IPS with an inline or inline with fail open interface set and use an intrusion policy that is set to drop when inline, you can alert on and drop malicious traffic. If instead you use the IPS component with a passive interface set or an intrusion policy that is not set to drop when inline, then the IPS component can only alert on malicious traffic and cannot affect the network traffic flow.

Intrusion Event Analyst

A user role that provides access to IPS analysis features, including intrusion event views, incidents, and reports. Intrusion Event Analysts see the main toolbar and IPS analysis-related options on the Analysis & Reporting and Operations menus. The Intrusion Event Analyst (Read Only) role provides read-only access to the same set of functions.

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.

link state propagation mode

An option you can enable for an inline interface set. With this option enabled, when one of the interfaces goes down the other interface in the set is automatically brought down within a few seconds. For copper fail-open NIMs, when the first interface comes back up the second interface comes up automatically. Link state propagation mode is also available for fiber interface cards, but that recovery is not automatic. To restore fiber interfaces you must reset the NIM. Crossbeam-based software sensors and 3D9800 sensors do not support this feature.

load balancer

A network device that distributes network traffic to optimize performance and resource use. RNA identifies network devices as load balancers if the TTL value changes from the client side, or if the TTL value changes more frequently than a typical boot time. RNA distinguishes between load balancers and NAT devices depending on what side the analyzed traffic is coming from: server (load balancer) or client (NAT device).

MAC address

Media Access Control address. A MAC address is a NIC's (network interface card's) unique hardware address. RNA detects the MAC addresses and hardware vendors of the NICs for the hosts and network devices on your network.

managed sensor

A 3D Sensor, Intrusion Agent, or software sensor configured and managed by a Defense Center.

Keysight, EX-1004
Page 2085

management interface The network interface that you use to administer the Defense Center or 3D Sensor. In most installations, the management interface is connected to an internal, protected network. Compare with sensing interface.

Maintenance User A user role that provides access to monitoring and maintenance features. Maintenance users see the main toolbar and maintenance-related options on the **Operations** top-level menu.

Master Defense Center A special-purpose appliance that is capable of aggregating intrusion events and compliance events from up to ten other Defense Centers. A Master Defense Center is also able to collect health status from its managed Defense Centers.

NAT device A network device that performs network address translation (NAT), most commonly to share a single internet connection among multiple hosts on a private network. RNA identifies network devices as NAT devices if the TTL value changes from the client side, or if the TTL value changes more frequently than a typical boot time. RNA distinguishes between load balancers and NAT devices depending on what side the analyzed traffic is coming from: server (load balancer) or client (NAT device).

Nessus An open source vulnerability scanner developed through the Nessus Project (http://www.nessus.org/) that uses Nessus plugins to test for vulnerabilities on the hosts that it scans.

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 data collected by RNA and to monitor networks not covered by 3D Sensors with RNA.

network device In the Sourcefire 3D System, a bridge, router, NAT device, or load balancer.

network discovery event A kind of RNA event that communicates the details of changes to the hosts on your monitored network. *New* events are generated for newly discovered network features, and *change* events are generated for any change in previously identified network assets. Settings in the system policy determine the types of network discovery events that are stored in the RNA database.

network map A detailed representation of your network generated by RNA. The network map allows you to view your network topology in terms of the hosts and network devices running on your network as well as their associated host attributes, services, and vulnerabilities.

Nmap An open source active scanner that you can use to detect operating systems and services running on a host. Running an Nmap scan adds the information detected to your network map.

Nmap scan A scan of a designated host or hosts to detect operating systems and services.

Keysight, EX-1004
Page 2086

NTP Network Time Protocol. NTP uses Coordinated Universal Time (UTC time) to synchronize the computer clocks in a network. You can synchronize the Defense Center's time with an NTP server. You can also configure a Defense Center as an NTP server so that managed sensors can synchronize time with it.

object, for import or export A policy or rule that is created on an appliance and can be exported from that appliance and imported by another appliance. Depending on the type of appliance and the components you are licensed to use, you can import and export some RNA detectors, custom table views, custom workflows, dashboards, system policies, intrusion policies, custom intrusion rules and rule classifications, RNA detection policies, and health policies.

operating system identity The operating system vendor and version details for an operating system on a host.

packet A unit of data routed between a source and a destination on a network. When data travels from one place to another on a network, the file is divided into chunks of an efficient size, called packets, for routing. The packets that comprise a single file may travel different routes through the network, but can be reassembled into the original file at the receiving end. If you are licensed for the IPS component, you can view the portion of a packet that was captured as part of an intrusion event.

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** generates 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 alert rule and drop rule.

passive A type of interface set that allows you to deploy a 3D Sensor passively on a network. In this configuration, the IPS component cannot affect the traffic flow.

passive detection The detection of host operating system and service information through analysis of traffic passively collected by RNA.

payload In an event, the content of http traffic detected by RNA, if available. Payload information is comprised of a payload type, which represents the general content type (for example, audio or video) as well as a payload, which represents the specific type of content (for example, WMV or QuickTime).

PCRE Perl-compatible regular expression. You can search packet payloads for content using PCREs. This is useful if you want to search 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.

Keysight, EX-1004
Page 2087

PEP	A technology based on the hardware capabilities of 3D9900 interface sets that allows you to use a PEP policy for advanced traffic management.
PEP policy	The criteria used when determining if a PEP-capable interface set should block, analyze, or send traffic directly through the sensor with no further inspection.
policy	A mechanism for applying settings to an appliance or detection engine. See intrusion policy, defragmentation policy, compliance policy, RNA detection policy, health policy, PEP policy, system policy, and security policy.
policy and response	A feature you can use to build a compliance policy that responds in real-time to threats on your network. In addition, the remediation component of policy and response provides a flexible API that allows you to create and upload your own custom remediation modules to respond to policy violations.
Policy & Response Administrator	A user role that provides access to rules and policy configuration. Policy & Response Administrators have access to the main toolbar and rule and policy-related options on the Policy & Response and Operations menus.
policy violation	A security breach, attack, exploit, or other misuse of your network as detected by a compliance policy.
port	The endpoint of a logical connection on a TCP or UDP network. Each port on a host has a number, which identifies the type of port. Many services have default ports; for example, HTTP traffic typically uses port 80. TCP and UDP use port numbers to separate data transmissions on the same network interface on the same host. With IPS, when you tune your intrusion policy, you can define, in both variables and rules, specific port numbers, such as ports susceptible to shell code exploits, HTTP (or web server) ports, and database server ports. This lets you specify the level of granularity of inspection so that rules execute against ports appropriate to your network needs.
portscan	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 services are running on these ports.
predefined table	A database table delivered with the Sourcefire 3D System. You can use the web interface to view the event information in the predefined tables. Predefined tables cannot be modified. Compare with custom table.
predefined workflow	A workflow delivered with the Sourcefire 3D System. You cannot modify predefined workflows. Compare with custom workflow and saved custom workflow.
preprocessor	A feature of IPS that normalizes traffic and helps identify network layer and transport layer protocol anomalies by identifying inappropriate header options,

Keysight, EX-1004
Page 2088

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 detection engine can analyze; these preprocessors are called data normalization preprocessors, or application-layer protocol preprocessors. Normalizing application-layer protocol encoding allows the detection engine to effectively apply the same content-related 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 — A rule associated with a detection option of one of the preprocessors or with the portscan flow detector included in the IPS component of the Sourcefire 3D System. 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 terms that is 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.

RADIUS authentication — 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 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 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 compliance policy, associate them with compliance rules and compliance white lists so that when they trigger, the Defense Center launches the remediation. This not only can automatically mitigate attacks when you are not immediately available to address them, but also can ensure that your system remains compliant with your organization's security policy. The Defense Center ships with predefined remediation modules: three that are designed for particular firewalls and routers, one that lets you perform Nmap scans, and one that lets you set host attributes. You also can use a flexible API to create custom remediations.

remediation status event — An event generated when a remediation is launched.

Keysight, EX-1004
Page 2089

replace rule When using an inline IPS detection engine, you use the `replace` keyword in a custom standard text rule to replace a specific string with exactly the same number of characters. This allows you to replace the content of malicious packets with benign alternatives. Only the first instance of the content found by the rule is replaced. The sensor automatically updates the packet checksum so that the destination host can receive the packet without error.

report profile A template for an event report. You can create and save custom report profiles. You can then manually run reports based on the profiles, or schedule the Sourcefire 3D System to generate reports automatically. You can use report profiles to add your company logo to reports, define the set of events that appear, specify the amount of detail, and specify the report's output file format.

response A reaction to a compliance policy violation—either an alert or a remediation.

Restricted Event Analyst A user role that can provide access to the same features as Intrusion Event Analyst or RNA Event Analyst access. You can restrict access by only allowing access to those events that match specified search criteria or you can turn off access for an entire category of events. Restricted event analyst users see only the main toolbar and analysis-related options on the Analysis & Reporting and Operations menus. The Restricted Event Analyst (Read Only) role provides read-only access to the same set of functions.

RNA A component of the Sourcefire 3D System that is installed by default on the 3D Sensor and that passively analyzes your network traffic to provide you with a complete, persistent view of your network. RNA identifies new and changed hosts on your network as well as tracks the sessions involving monitored hosts. For each detected host, RNA discovers the services and client applications that they use as well as vulnerabilities to which the host is susceptible. Note that you must add a feature license in the form of an RNA host license on the Defense Center that manages the sensor before you can view RNA data.

RNA detection policy A policy that you apply to RNA detection engines that specifies the kinds of data RNA collects, as well as the network segments each RNA detection engine or NetFlow-enabled device monitors.

RNA detector An RNA detector provides RNA with the information needed to identify non-standard services, including the port used by service traffic, a pattern within the traffic, or both the port and the pattern. The Sourcefire 3D System is delivered with many internal RNA detectors, or you can create your own. In addition, Sourcefire may deliver additional RNA detectors that you can add to the Sourcefire 3D System via vulnerability database updates or via the Import/Export feature.

RNA event An event generated by RNA. RNA events include network discovery events, which communicate the details of changes to the hosts on your monitored network, and flow events, which are records of sessions involving monitored hosts. RNA events also include client application events, host events, host

Keysight, EX-1004
Page 2090

attributes, and service events, which provide general information about your network topology. A vulnerability is also considered an RNA event.

RNA Event Analyst — A user role that provides access to RNA analysis features, including event views, network maps, host profiles, services, vulnerabilities, client applications, and reports. RNA Event Analysts see the main toolbar and RNA analysis-related options on the Analysis & Reporting and Operations menus. The RNA Analyst (Read Only) role provides read-only access to the same set of functions.

RNA Recommendations layer — A built-in layer in an intrusion policy that exists when you choose to allow IPS to modify the rule states of shared object rules and standard text rules to the states recommended by the RNA recommended rules features. You cannot manually modify or remove this layer. IPS removes or restores the layer when you decide to not use or use, respectively, recommendations for a policy.

RNA recommended rules — A feature that recommends which rules should be enabled or disabled in your intrusion policy, based on information from your RNA 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 RNA Recommendations layer.

router — A network device, located at a gateway, that forwards packets between networks. RNA identifies network devices as routers if they are detected communicating using CDP or if they meet other qualifying criteria. For example, if multiple hosts appear to be using the same MAC address, that MAC address is often identified as the router to which the multiple hosts are connected.

RUA — Real-time User Awareness, also called RUA, allows your organization to correlate threat, endpoint, and network intelligence with user identity information.

RUA Agent — An RUA Agent is an agent you install on a Microsoft Active Directory server to monitor users as they log into the network or when they authenticate against Active Directory credentials for any other reason.

RUA event — An event generated by RUA in response to a detected user login or the addition or deletion of a user from the RUA database. RUA events are stored in the RUA database.

RUA user — A user detected by RUA whose user identity data is stored in the RUA database.

rule — A construct that provides criteria against which network traffic is examined. Rules can detect a variety of intrusions, attacks, exploits, and suspicious traffic. See compliance rule, intrusion rule, alert rule, pass rule, drop rule, and preprocessor rule.

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

Keysight, EX-1004
Page 2091

not used. A drop rule can drop packets that trigger the rule in an inline deployment.

saved custom workflow — A custom workflow that is based on a custom table and delivered with the Defense Center. Unlike predefined workflows, you can modify saved custom workflows.

scheduled task — An administrative task that you can schedule to run once or at recurring intervals. Depending on the appliance where you are creating the task, you can schedule tasks to run backups, apply an intrusion policy, generate reports, download and install SEUs, manage RNA recommended rules, run Nmap scans, download and install software and vulnerability database updates, and push downloaded updates to managed sensors.

security policy — 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. For example, your AUP might forbid the use of instant messaging client applications.

sensing interface — A network interface on a sensor that you use to monitor a network segment. You can connect sensing interfaces to your network in various ways. How you plan to deploy your detection engines (passively or inline) affects how you connect them to your network. Compare with management interface.

sensitive data — A preprocessor that detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. This can be particularly useful for detecting accidental data leaks that can occur, for example, when an employee emails themselves a list of credit card numbers to work with at home.

sensor group — On the Defense Center, a logical group that can contain or more managed sensors so you can more readily manage them. For example, you can easily apply a system policy to, or install updates on, multiple sensors at once.

Series 1 appliance — The first series of Sourcefire appliance models, including the following models: 3D500 PW, 3D1000 NH, 3D2000 NH, 3D2100 NH, 3D3000 JR, DC1000 JR, DC3000 JR, and the 3Dx800 sensor models.

Series 2 appliance — The second series of Sourcefire appliance models, including the following models: 3D500 PB, 3D1000 PB, 3D2000 PB, 3D2100 FR, 3D2500 FR, 3D3500 FR, 3D4500 FR, DC1000 AL, DC3000 AL, and MDC3000 AL. All appliances currently shipping from Sourcefire, with the exception of the 3Dx800 models, are Series 2 appliances.

service — Work performed by a server. NTP, SSH, HTTP, and AIM are examples of services.

service event — An event indicating that RNA has detected a service running on a specific host. RNA collects information about all services run by hosts on monitored network

Keysight, EX-1004
Page 2092

segments. The information that RNA collects includes the name of the service, the protocol used by the service, the IP address of the host running a service, and the port on which the service is running.

service identity — The service type, vendor, and version details for a service on a host.

SEU (Security Enhancement Update) — An as-needed product update that contains new and updated standard text rules and shared object rules. In addition, SEUs can provide your Defense Centers and 3D Sensors with an updated version of Snort, as well as new and updated features such as preprocessors and decoders.

shared layer — An intrusion policy layer than 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 — 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 a SID (Signature ID, also called Snort ID) 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.

simple condition — A single constraint placed on a compliance rule, flow tracker, host profile qualification, or traffic profile. You can link simple conditions with other simple or complex conditions using AND or OR operators.

simple constraint — A simple constraint sets a single constraint on the events retrieved in an event view or event search.

SNMP alerting — The transmission of an alert as an SNMP trap. Each event SNMP trap contains information identifying the server's name, the sensor's IP address, and the event data.

SNMP trap — A message sent by a network device on UDP port 162 using the simple network management protocol (SNMP) when errors or specific events occur on the network. See also SNMP alerting.

Keysight, EX-1004
Page 2093

snooze period An interval specified in seconds, minutes, or hours after a compliance 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 elapsed, 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 IPS detection engines use Snort to test packets against decoders, preprocessors, and intrusion rules.

Sourcefire-defined custom table A table that is delivered with the Defense Center that contains fields from two or more predefined tables.

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 existing standard text rules provided by Sourcefire. A standard text rule has a GID (generator ID) of 1.

stateful inspection A preprocessor that makes sure that only packets that are part of a TCP session established with a legitimate three-way handshake between a client and server can generate intrusion events. This allows analysts to focus on these events rather than the volume of events caused by denial of service (DoS) attacks like stick or snot.

stream reassembly A preprocessor that IPS uses to collect and reassemble all of the packets that are part of a TCP session's server-to-client or client-to-server communication stream. Stream reassembly allows the detection engine to inspect the stream as a single entity rather than only the individual packets, which allows the detection engine to identify stream-based attacks.

subnet detection A feature that allows RNA to automatically determine the closest subnets to each RNA detection engine and then make recommendations about which detection engines should be the reporting detection engines for specific subnets.

subnet mask A bit mask used to identify which bits in an IP address correspond to the network address, and which correspond to the subnet portion of the address.

suppression See event suppression.

switch A multiport bridge.

syslog A logging system, also called the system log, used by many operating systems. You can configure the Defense Center or 3D Sensor to perform syslog alerting.

syslog alerting The transmission of an alert as a message to an external syslog. All syslog messages include both a facility and a priority level. The facility indicates the

Keysight, EX-1004
Page 2094

subsystem (for example, FTP, NTP, or MAIL) that created the message and the priority defines the importance of the message.

system policy Settings that are likely to be similar for multiple appliances in a deployment, such as access configuration, authentication profiles, database limits, DNS cache settings, the mail relay host, a notification address for database prune messages, language selection (English or Japanese), login banner, RNA settings, and time synchronization settings. You can configure a system policy on a Defense Center and then apply the policy to the Defense Center and its managed sensors.

system settings Settings that are specific to a single appliance, such as appliance name, IP address, time settings, licensing, and remote management settings. You can also use the system settings pages to shut down or reboot an appliance and to restart its software.

table view A type of workflow page that displays event information. Table views include a column for each of the fields in the database. For example, the table view of intrusion events includes columns such as Time, Priority, Impact Flag, Source IP, Destination IP. As another example, the table view of RNA network discovery events includes such columns as Time, Event, IP Address, MAC Address, and so on. Generally, you 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 the next to last page in predefined workflows; advancing from the table view leads to the packet view (for workflows based on intrusion events), the host view (for workflows based on RNA events), the vulnerability detail page (for workflows based on vulnerabilities), or the user identity view (for workflows based on RUA events).

tap mode A setting for an inline interface set on a 3D3800, 3D5800, 3D9800, or 3D9900 sensor where a copy of each packet is sent to the sensor and the network traffic flow is undisturbed instead of the packet flow passing through the sensor. Because you are working with copies of packets rather than the packets themselves, you cannot use drop rules or replace rules as you can with a sensor that is deployed in the packet stream.

task queue A queue of jobs that the Defense Center or 3D Sensor needs to perform. When you apply a policy, install software, 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.

three-way handshake The process two hosts use to establish a TCP/IP connection. A three-way handshake occurs when the originating host sends a SYN (synchronization) packet to the destination host. The destination then sends its own SYN packet and an ACK (acknowledgement) packet. The originator then returns an ACK which acknowledges the SYN/ACK packets the destination sent. With IPS, you can configure intrusion rules evaluate the data in established TCP sessions only

Keysight, EX-1004
Page 2095

(where a three-way handshake has occurred) or on all traffic, including orphaned packets.

thresholding See event thresholding.

traffic profile A profile of the traffic on your network, based on flow data collected by RNA 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 using criteria based on the data in flow events. Then, you can use the policy and response feature to detect abnormal network traffic by evaluating new traffic against an existing profile.

transparent inline mode A setting that allows 3D Sensors configured with inline interface sets that forward packets regardless of whether they contain MAC addresses that are valid for the monitored network.

unidentified host A host whose operating system cannot be identified because RNA has not yet gathered enough information about the host. Compare with unknown host.

unknown host A host whose traffic has been analyzed by RNA, but whose operating system does not match any known fingerprints. Compare with unidentified host.

Unified file A binary file format that the Sourcefire 3D System uses to log event data.

user access 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.

user identity view The user identity view provides details on RUA users and a host history with a graphic representation of the last twenty-four hours of the user's activity.

user input data Host input data added through the Sourcefire 3D System user interface by setting or modifying an identity.

user layer A layer in an intrusion policy where you can modify settings in the policy.

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 pre-configured 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. You can use a variable within an intrusion policy or a specific IPS detection engine.

Keysight, EX-1004
Page 2096

VLAN

A virtual local area network. VLANs map hosts not by geographic location, but by some other criterion (such as by department or primary use). This is useful if you want to separate hosts into small, logical network segments. A host's host profile shows any VLAN information associated with the host. VLAN information is included in intrusion events (as the innermost VLAN tag in the packet that triggered the event). You can filter intrusion policies by VLAN, or target compliance white lists by VLAN.

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 RNA has detected on your entire monitored network. If you deem that a host or hosts is no longer vulnerable to a specific compromise, you can deactivate, or mark as invalid, a specific vulnerability.

vulnerability database

A database of known vulnerabilities to which hosts may be susceptible. The database includes such technical details as vulnerability title and identification number, technical details, whether any exploits are known to take advantage of the vulnerability, known solutions, and so on. RNA correlates the operating system and services detected on each host with the vulnerability database to help you determine whether a particular host increases your risk of network compromise.

vulnerability detail page

A page in a workflow that provides information about a specific vulnerability, including technical details and known solutions. The vulnerability detail page is the final page in workflows based on vulnerabilities.

white list

Either a compliance white list *or* a list of IP addresses that you can configure within a remediation to exempt the IP addresses from some kind of action. For example, you could configure a firewall-based remediation to block all hosts that trigger a specific compliance rule, with the exception of hosts specified in a white list.

white list event

An event generated when RNA detects that a valid target host has become non-compliant with a compliance white list. For example, you can configure the Defense Center to generate a white list event when RNA detects a new non-compliant service running on a target host. Note that a white list event is a special kind of compliance event.

white list violation

A white list violation is an event that occurs when RNA generates an event that indicates that a host is out of compliance. The Sourcefire 3D System includes workflows that allow you to view each of the individual white list violations, as well as the number of violations per host.

whois

A mechanism for finding contact and registration information for IP addresses. If your Defense Center or 3D Sensor is connected to the Internet, you can use the

Keysight, EX-1004
Page 2097

web interface to look up information about an IP address using the whois feature, which uses ARIN's (American Registry for Internet Numbers) WHOIS service.

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 (which could be, depending upon the type of analysis you are performing, a packet view, host view, vulnerability detail page, or user identity view). IPS provides two categories of workflows: predefined workflows and custom workflows. The Defense Center provides three categories of workflows: predefined workflows, custom workflows, and saved custom workflows.

Keysight, EX-1004
Page 2098

Index

Symbols

%U encoding (HTTP Inspect option) 571
$AIM_SERVERS 373
$DNS_SERVERS 373
$EXTERNAL_NET 373
$HOME_NET 373
$HTTP_PORTS 373
$HTTP_SERVERS 373
$ORACLE_PORTS 373
$SHELLCODE_PORTS 373
$SMTP_SERVERS 374
$SNMP_SERVERS 374
$SQL_SERVERS 374
$TELNET_SERVERS 374

Numerics

3D Sensors
  adding to a Defense Center 122
  deleting 126
  deleting 3Dx800s 133
  disabling communications 144
  health policy 1897
  host name 143, 153
  management concepts 104
  managing 103, 118
  managing 3Dx800s 131
  resetting communications 134
  resetting management 127
  restarting 143
  sensor attributes 139
  sensor information 141
  Sensors page 120
  stopping 143
  time sync 145
  unregistering 162
  updating 1838, 1840
3D9900s
  clustering 250
  hardware alert details 1962
3Dx800s
  deleting 133
  health policy 1897, 1901
  managing 112, 131
  resetting communications 134
3Dx900s
  PEP policies 1510
3-way handshake timeout (stream option) 636

Keysight, EX-1004
Page 2099

A

access list 1738
accessing the appliance 37, 40
ack (rule keyword) 818
action (firewall rule) 1311, 1314, 1317
active responses
  configure interface 847
  configure reset attempts 847
  intiating with drop rules 629
  intiating with keywords 844
active scanning 1418
  analyzing results 1451
  monitoring scans 1444
  Nmap 1419, 1428
  scan results 1443
  searching for scan results 1449
adaptive profiles 697
  configuring 702
  intrusion rules 700
  preprocessors 698
  understanding 698
add client application (RNA event type) 1227
add host (RNA event type) 1011, 1227
add protocol (RNA event type) 1011, 1227
add scan result (RNA event type) 1011, 1227
add service (RNA event type) 1011, 1227
adding sensors to a Defense Center 122
Additional Information (vulnerability detail) 982
additional MAC detected (RNA event type) 1006, 1223
Admin access 1683
advanced settings (intrusion policy)
  automatically enabling 490
  layers 504
  modifying 477
alerting
  compliance responses 1282
  email alerts 756, 1289
  impact flag alerting 1295
  intrusion events 458
  OPSEC alerts 415
  RNA events 1293
  SNMP alerts 415, 748, 1290
  syslog alerts 752, 754, 1285
any, in variables and intrusion rules 400
appliance groups 191
  creating 191
  deleting 193
  editing 192
appliance heartbeat monitoring 1893, 1906
appliance information 141, 1777
appliance status widget 69
Application 229
application layer decoders 513
ARP 258
ASCII encoding (HTTP Inspect option) 570
asn1 (rule keyword) 829
asynchronous network (stream option) 637
asynchronous routing and interface sets 234
audit log
  intrusion policy 346
  time window 47
audit log settings 1740
auditing
  audit records 1969
  field descriptions 1977
  introduction 1968
  searching 1978
  understanding 1977
  viewing 1969
auditing, predefined workflows 1586
authentication objects 1631
  creating 1642
  deleting 1674
  editing 1661
  LDAP 1642, 1644
  RADIUS 1662
authentication profiles 1741
Auto MDI/MDIX 1805
Automatic Application Bypass 229
automatic application bypass monitoring 1907
automatically enabling IPS features 490
available exploits (vulnerability detail) 982, 1060

B

Back Orifice
  detecting 648
  generator ID 489
backup and restore 1985
  remote backups 1991
  scheduling backups 1702
backup files
  creating 1986
  location 1990
  restoring 1993
backup profiles 1990
bare byte UTF-8 (HTTP Inspect option) 571
base (intrusion) policy
  allowing SEU to modify 363
  selecting 364
base 36 encoding (HTTP Inspect option) 571

Keysight, EX-1004
Page 2100

base64_data (rule keyword) 859
base64_decode (rule keyword) 859
blacklists
  health monitoring 1935, 1938
  Master Defense Center 196
  system settings 1776, 1815
bookmarks 1615
both ports (stream reassembly option) 640
both services (stream reassembly option) 640
browser requirements 37
Bugtraq ID
  network map 932
  vulnerability detail 981, 1060
bypass mode for fail open fiber interfaces 246, 247
byte_extract (rule keyword) 840
byte_jump (rule keyword) 794
byte_test (rule keyword) 797

C

card reset monitoring 1911
Check Point OPSEC SAM 1298
  adding the OPSEC instance 1302
  certificate creation 739
  clearing responses 745
  configuring the sensor 739
  creating the OPSEC application 736, 1299
  disabling responses 735
  filter options 464
  firewall objects 745
  firewall responses 743
  intrusion event responses 460, 735
  OPSEC block 1306, 1310, 1313, 1316
  viewing OPSEC debug messages 747
  viewing SAM client information 746
checksum
  verification 602
CIDR notation 59
Cisco IOS routers
  adding a Cisco IOS instance 1320
  configuring remediations 1319
  IOS block 1322, 1324, 1325, 1326
Cisco PIX firewalls
  adding a Cisco PIX instance 1329
  Cisco PIX block 1331, 1333
  remediations 1327
classtype (rule keyword) 772
client application timeout (RNA event type) 1006, 1223
client application timeout (RNA settings) 1760
client application update (RNA event type) 1006, 1223
client applications
  deleting from host profile 972
  detected by RNA 2036
  field descriptions 1053
  host profiles 969, 970
  introduction 1050
  payloads 971, 1053
  searching 1054
  understanding 1053
  viewing 1051
  white lists 1180
client fingerprints 1364
client ports (stream reassembly option) 639
client requirements 37
client services (stream reassembly option) 639
clipboard
  clipboard reports 316
  copying events 286, 291
  deleting intrusion events 318
  introduction 316
clustering 3D9900s 250
comments (rule editor) 867
comments (rules in an intrusion policy) 417, 466
committing intrusion policy 345
communications channel 118, 1807
comparison report, intrusion policy 357
comparison view, intrusion policy 355
compliance events
  aggregating 168
  dashboard widgets 69
  field descriptions 1274
  introduction 1271
  predefined workflows 1585
  searching 1276
  understanding 1274
  viewing 1272
compliance policies 1212
  activating and deactivating 1270
  adding responses 1267
  adding rules 1265
  adding white lists 1265
  creating 1262
  creating email alerts 1289
  creating SNMP alerts 1290
  creating syslog alerts 1285
  deleting 1271
  editing 1270
  introduction 1212
  managing 1269
  remediations 1296
  response groups 1351
  responses 1281

Keysight, EX-1004
Page 2101

setting rule and white list priorities 1266
compliance responses 1281, 1293
activating and deactivating alerts 1293
adding to rules and white lists 1267
creating alerts 1282
creating email alerts 1289
creating SNMP alerts 1290
creating syslog alerts 1285
deleting or editing alerts 1292
introduction 1281
response groups 1351
compliance rules 1212
adding to compliance policies 1265
basic information 1216
building rules 1251
conditions 1255, 1258
creating 1214
creating rule groups 1261
deleting 1261
deleting rule groups 1262
editing 1260
flow data triggers 1218
flow trackers 1237
host input event triggers 1227
host profile qualification triggers 1234, 1249
host profile qualifications 1233
inactive periods 1250
introduction 1212
intrusion event triggers 1221
managing 1260
RNA event triggers 1223
RUA 1248
RUA event triggers 1226
rule groups 1261
snooze periods 1250
traffic profile triggers 1229
triggers 1216
compliance white lists, *see* white lists
compound constraints 1609
conditions (in compliance rules) 1255, 1258
confidence 1015
detected services 961, 964
host profile 953
host view 1024
consecutive small segments (stream option) 635
constraining regular expressions 728
content (rule keyword) 776
context menus 53
conventions
platform requirements 56
correlator process 1880
CPU temperature monitoring 1893, 1908
CPU usage monitoring 1893, 1909
criticality 986
host view 1023
Crossbeam-based sensors 37, 115, 116
deploying inline 348
IPS 265
CSV reports 1547
current sessions widget 71
current time 90
custom
fingerprints 1363
login banner 1755
OS mappings 966, 968, 1367, 1374, 1410, 1411
report footers 1548
RNA detectors 1381
tables 1561
topologies 936
workflows 1587, 1618
custom analysis widget
configuring 74
enabling 1745, 1746
understanding 71
custom tables
and workflows 1569
creating 1565
deleting 1568
editing 1568
introduction 1561
searching 1569
table combinations 1563
understanding 1562
custom topology 936
activating 942
creating 937
deleting 943
editing 942
importing 939, 940
managing 942
manually editing 940
topology information 939
CVE ID
network map 932
vulnerability detail 981
cvs (rule keyword) 839

Keysight, EX-1004
Page 2102

D

dashboards 61, 92
  adding widgets 99
  custom dashboards 92
  default dashboard 52, 61
  deleting 102
  home page 62
  modifying 96
  properties 97
  settings 1745, 1746
  tabs 98, 99
  viewing 94
  widgets 62, 66
data link layer 258, 300
database
  limits 1747
  purging 2014
datasets (for flow graphs) 1082, 1113
date published (vulnerability detail) 981, 1060
DCE/RPC preprocessor 514
  generator ID 490
  global options 515
  RPC over HTTP transport 521
  target-based policies 517
  target-based policy options 522
  transports 518
decoders
  decoding packets 257, 621
  FTP Telnet 538
  HTTP Inspect 560
  RNA detectors 1381
  RPC 578
decoy portscan 653
default detection engines 205
Defense Center 103
  adding sensors 122
  deleting sensors 126
  disabling communications 144
  managing 3D Sensors 118
  managing with a Master Defense Center 165
  resetting management 127
  restarting managed sensors 143
  time sync 145
  updating 1835
Defense Center groups 191
  creating 191
  deleting 193
  editing 192
  managing 191
Defense Center, introduction 34
delete client application (RNA event type) 1011, 1227
delete host/network (RNA event type) 1011, 1227
delete protocol (RNA event type) 1011, 1227
delete service (RNA event type) 1011, 1227
depth content option (rule keyword) 783
derived fingerprints 1359
description (vulnerability detail) 982
destination IPs
  rules 766
destination ports
  rules 768
detect connectionless DCE/RPC traffic (DCE/RPC option) 526
detected services 958
detection engine groups 211
  creating 211
  deleting 213
detection engines
  and IPS 202
  assign variable values 215
  creating 206
  creating variables 217
  default 205
  definition 199
  deleting 211
  deleting variables 218
  detection engine groups 211
  detection engine types 207
  editing 208
  in an intrusion policy 368
  in intrusion policies 350
  in searches 1559
  introduction 198
  managing 206
  resetting variables 218
  RNA detection policy 905
  understanding 199
  variables 213
detection_filter (rule keyword) 666, 849
DHCP 1799, 1802
  IP address changed (RNA event type) 1006, 1223
  IP address reassigned (RNA event type) 1006, 1223
differentiated services (DS)
  in rules 815
directional operators, in rules 769
directory transversal (HTTP Inspect option) 572
disable pipeline decoding (HTTP Inspect option) 573
disabling
  communications between appliances 144
  high availability 162
disk usage monitoring 1893, 1911
disk usage widget 82
distance content option (rule keyword) 781
distributed portscan 653

Keysight, EX-1004
Page 2103

DNS
  configuring the cache 1750
  primary server 1801
  secondary server 1801
  tertiary server 1801
DNS preprocessor
  configuring 535
  experimental options 535
  generator ID 490
  name server responses 531
  obsolete resource record types 534
  overflows in RData text fields 534
  resource record inspection 532
documentation
  conventions 55
  resources 55
domain name 1801
double encoding (HTTP Inspect option) 571
downloading patches 982, 983
drill-down pages 1574
  flow data 1105
  intrusion events 280
drop new hosts (RNA settings) 1760
drop rules, definition 761
dsize (rule keyword) 837
dynamic rule states 413
  understanding 453

E

EAP, EAPOL 258
email alerting 756
  configuring 759
  creating 1289
email notification 1751
email relay host 1751
enable defragmentation (DCE/RPC option) 516
Enabling Fail-Safe 230
enabling features automatically 490
eStreamer 37
Ethernet 258
event aggregation 166
  compliance events 168
  intrusion events 167
  limitations 168
event classification 772
event database limits 1747
event graphs 273
event preferences 45
event statistics
  intrusion events 272
  RNA events 996
event stream monitoring 1913
event stream process monitoring 1893
event summary
  event breakdown (RNA) 998
  intrusion events 269
  OS breakdown (RNA) 1001
  protocol breakdown (RNA) 1000
  service breakdown (RNA 1000
  statistics summary (RNA) 997
event view time window 1597
events
  acknowledging 278
  copying to the clipboard 286, 291
  deleting 286, 291
  event graphs 273
  generating 260
  host profile 950
  intrusion event summary 269
  IPS and RNA correlation 307
  reports 1531
  restoring events from backup files 1995
  searching 1550
  time window 47, 1597
excluding ports from RNA monitoring 907
expanding time window 47
experimental DNS options 535
experimental TCP options 623
exporting
  custom tables 1999
  custom workflows 1999
  dashboards 2000
  flow data 1115
  health policies 2001
  introduction 1997
  intrusion policies 2001
  multiple objects 2006
  policies 1998
  RNA detection policies 2003
  RNA detectors 2005
  system policies 2003
  third-party product mappings 2004
  third-party vulnerability mappings 2005
extended ASCII encoding (HTTP Inspect option) 573
external alerting
  email alerting 756
  SNMP alerting 748
  syslog alerting 752
external authentication 1631, 1662
External Database access 1683

Keysight, EX-1004
Page 2104

F

failed logins 1681
fan monitoring 1893, 1914
fast_pattern content option (rule keyword) 789
FDDI 258
feature licenses
  adding 1788
  viewing 1790
file_data (rule keyword) 858
filter configuration 187
filtering intrusion rules
  in a policy 417
  keywords (in a policy) 419
  keywords (rule editor) 873
  rule editor 872, 876
  strings (in a policy) 431
  strings (rule editor) 875
fingerprints
  activating 1377
  client fingerprints 1364
  deactivating 1378
  deleting 1378
  derived 1359
  editing 1379
  introduction 1363
  managing 1377
  OS mappings 966, 968, 1367, 1374, 1410, 1411
  server fingerprints 1370
firewall object 745, 1312, 1315, 1318
fixes (vulnerability detail) 982, 983
flags
  impact 307
  rule keyword 819
flow (detected services) 959, 971
flow (rule keyword) 820
flow data
  aggregated flow data 1101
  changing graph types 1109
  custom flow data workflows 1621
  datasets 1082
  definition 884
  detaching graphs 1114
  exporting 1115
  field descriptions 1089, 1091
  filtering 893
  flow summary page 1116
  graphs 1078, 1079, 1084, 1100, 1109
  introduction 1071
  line graphs 1108
  payloads 1090
  predefined workflows 1582
  searching 1094
  selecting datasets 1113
  tables 1086
  understanding 1072
  viewing 1093
  workflows 1103, 1105
flow filtering 893
  consequences 1074
  limitation with NetFlow 895
flow summaries 1116
  field descriptions 1091
  flow summary page 1116
  from external responders 1077
  graph types 1116
  limitations 1077
  long-running flows 1076
  understanding 1076
flow tracker events syntax 1247
flow trackers 1237
  syntax 1244
flowbits (rule keyword) 851
flush factor (stream option) 635
fragbits (rule keyword) 813
fragmentation, in IP datagrams 613
fragoffset (rule keyword) 839
frame information (packet view) 299
Frames 229
FTP decoder 538
  generator ID 490

G

generating events 260
generator IDs (GIDs) 488
global host profiles 1161, 1175
global policy management 170
global thresholding 708
graphs
  changing flow data graphs 1109
  detaching flow data graphs 1114
  flow data 1078, 1100
  health monitoring 1953
  intrusion events 273
  performance statistics 1884
groups
  appliance groups 191
  compliance response groups 1351
  compliance rule groups 1261
  Defense Center groups 191
  detection engines 211
  RNA detectors 1402
  sensor groups 137

Keysight, EX-1004
Page 2105

H

hardware alerts for 3Dx900s 1962
hardware monitoring 1893, 1915
health alerts 1940, 1941
  creating 1942
  deleting 1945
  editing 1944
  understanding 1943
health events 1957
  field descriptions 1964
  predefined workflows 1586
  searching 1965
  table view 1963
  understanding 1958
  viewing 1958
health modules 1893
  blacklisting 1938
  running all modules 1950
  running specific modules 1951
health monitor
  blacklist 1776, 1935, 1938
health monitor process monitoring 1894
health monitoring
  alerts 1940
  appliance monitoring 1947
  blacklists 1815
  configuring 1895
  creating alerts 1942
  default health policy 1899
  deleting alerts 1945
  editing alerts 1944
  events 1957
  graphs 1953
  health modules 1893
  health monitor 1945
  health policies 1892
  introduction 1890
  link state propagation 1919
  power supplies 1922
  running health modules 1950, 1951
  status icons 1949
  status indicators 1947
  time window 47
  understanding 1891
  understanding alerts 1943
health policies 1892
  applying 1928
  configuring 1895
  creating 1902
  defaults 1896
  deleting 1934
  editing 1931
  exporting 2001
health status monitoring 1916
high availability 117, 151, 152, 153, 155
  configuration guidelines 157
  disabling 162
  health monitoring blacklist 1936
  monitoring 160
  pausing communications 163
  restarting communications 163
  setting up 158
  shared configurations 153
  shared policies 154
  understanding 156
hits (detected services) 964
home page 51
  dashboards 62
hops
  hops change (RNA event type) 1006
  host profile 949
  host view 1023
host attribute add (RNA event type) 1011
host attribute delete (RNA event type) 1011
host attribute delete value (RNA event type) 1011, 1227
host attribute set value (RNA event type) 1011, 1227
host attribute update (RNA event type) 1011
host attributes
  assigning attribute values 974
  attribute types 988
  built-in host attributes 985
  creating 989
  definition 882
  deleting 991
  editing 991
  field descriptions 1036
  host profiles 973
  introduction 1033
  network map 935
  predefined 985
  remediations 1340
  restrictions 988
  searching 1038
  set value (RNA event type) 1011, 1227
  setting attributes 1036
  understanding 1035
  user-defined 987
  viewing 1034
host criticality 986
  host view 1023
host deleted (RNA event type) 1007, 1223
host dropped (RNA event type) 1007
host input 1406
  custom product mappings 1414
  event types 1010
  events 1004, 1012

Keysight, EX-1004
Page 2106

patch tools 1410
service identification 1414
third-party tools 1408
vulnerabilities 1412
host profile qualifications
compliance rules 1233
in traffic profiles 1124
syntax 1125
host profiles
activating vulnerabilities 982, 984
basic information 948
client applications 969, 970
creating white lists 977
definition 882
host attributes 973
host criticality 986
host protocols 974
identity conflicts 956, 967
impact qualification 981
introduction 944
network map 926
notes 986
operating system 951
predefined host attributes 985
scan results 992
scanning a host 992
services 958
shared host profiles 1163
third-party vulnerabilities 978
user history 973
viewing 948
VLAN tags 972
vulnerabilities 978
vulnerability details 980
white list violations 976
white lists 1160, 1175
host protocols 974
host statistics 1870
host statistics, IPS 271
host timeout (RNA event type) 1007, 1223
host timeout, RNA settings 1760
host type (host view) 1023
host type changed router/bridge (RNA event type) 1007, 1223
hostname 1801
hosts
current user (host profile) 950
host attributes 1033
host criticality 1023
host name (host profile) 949
host type (host profile) 950
identity conflicts 956, 967
introduction 1019
IP address 1023
last seen 950, 1023
MAC addresses 1025
misidentification 1357
NETBIOS name 1023
new host (RNA event type) 1008, 1223
searching 1027
source type 1024
traffic profiles 1025
viewing 1020
viewing on the network map 927
VLAN ID 1023
white list 1026
whois information 1027
HTML reports 1547
hts-Out 1823
HTTP decoder
generator ID 489
HTTP Inspect
configuring 574
decoding 560
generator ID 489
normalization encoding options 570
normalization options 562, 563, 564
http_client_body content option (rule keyword) 788
http_cookie content option (rule keyword) 787
http_header content option (rule keyword) 786
http_method content option (rule keyword) 786
http_raw_cookie content option (rule keyword) 787
http_raw_header content option (rule keyword) 786
http_raw_uri content option (rule keyword) 786
http_status_code content option (rule keyword) 788
http_status_encode (rule keyword) 855
http_status_message content option (rule keyword) 788
http_uri content option (rule keyword) 785

I

ICMP 258
header values 816
icmp_all (resp keyword) 846
icmp_host (resp keyword) 845
icmp_id (rule keyword) 817
icmp_net (resp keyword) 845
icmp_port (resp keyword) 846
icmp_seq (rule keyword) 817
icode (rule keyword) 818
ID, rule keyword 814

Keysight, EX-1004
Page 2107

identity
current identity 1360
identity conflicts 1362
identity conflict (RNA event type) 1007, 1223
identity conflicts
resolving 956, 967
identity timeout (RNA event type) 1008, 1223
IEEE 802.11 258
IIS backslash (HTTP Inspect option) 572
IIS encoding (HTTP Inspect option) 571
impact (vulnerability detail) 981
impact flags 307
alerting 1295
definition 885
descriptions 307
impact qualification 981
impact qualification 982
import log 1859, 1861, 1865
importing
introduction 1997
policies 2009
inactive periods 1250, 1251
incidents
common processes 321
creating 324
default incident types 324
definition 320
editing 326
incident handling basics 320
incident types 329
introduction 319
reports 327
inline normalization 604
preprocessor 604
Inspect Traffic During Policy Apply 347
integer (host attribute type) 988, 989
interface sets 222
asynchronous routing 234
creating 230, 241
definition 200, 222
deleting 244
editing 239
forcing into bypass mode 247
interface status widget 70
interface traffic widget 83
introduction 198
link state propagation mode 228
multiple inline interface sets 234
PEP policies 1514
tap mode 227
transparent inline mode 226
types 200, 222, 225
interface status widget 70
interface traffic widget 83
Internet Control Message Protocol, *see* ICMP
Internet Protocol, *see* IP
introduction 32
intrusion agents 37, 111
licenses 1787
intrusion detection and prevention, *see* IPS
intrusion event responses 261
Check Point OPSEC SAM responses 460, 735
email alerting 756
introduction 734
SNMP alerting 748
syslog alerting 752, 754
intrusion events
aggregating 167
analyzing 261
clipboard 316
constraining events 288
drill-down pages 280
event statistics 272
field descriptions 276
impact flags 307
introduction 267
intrusion event responses 734
packet view 282, 290
portscan events 658
predefined workflows 1576
reviewing 278
searching 309
suppression 447
table view of events 281
thresholding and suppression 439
understanding 276
unreviewing 278
viewing 274, 284
widget 83
workflows 280
intrusion policies
adaptive profiles 697
advanced settings 477
alerting 458
applying 347
base policy (allowing SEUs to modify) 363
base policy (selecting) 364
benefits 265
committing 345
comparing 354, 357
creating 338
default policies 362
detection challenges 485
detection engines 350, 368
drop rule behavior (setting) 358
dynamic rule states 452

Keysight, EX-1004
Page 2108

editing 341
event filtering 439
exporting 2001
importing local intrusion rules 1857
importing SEUs 363, 1848, 1851
introduction 331, 477
layers 494
managing 335
navigation panel 344
normalizing inline traffic 604
planning 333
predefined variables 372
preferences 1753
preprocessors 477, 484, 713
quick apply 349
rate-based attack prevention 452
report 350
RNA recommended rules 468
rules 403
Rules intrusion policy page 497
scheduling policy applications 1719
setting rule state 435
suppression 447
target 368
thresholding 439
troubleshooting options 492
variables 371, 374
viewing rules 406
VLANs and subnetworks 393
intrusion rules
adaptive profiles 700
alerting 458
comments 867
comments (in a policy) 417, 466
creating 861
directional operators 769
dynamic rule states 413, 452
editing 864
filtering (in a policy) 417
filtering (rule editor) 872
ICMP header values 816
importing local rule files 1857
importing SEUs 363, 1848, 1851
introduction 761
IP header values 812
keywords 769
managing 403
OPSEC alerts 415
parts of a rule 762
replacing content 792
RNA recommended rules 468
rule actions 765
rule details 410
rule headers 764
searching for 869
setting rule state 435
SNMP alerting 415
sorting (in a policy) 409
specifying IPs 766
specifying ports 768
suppression 413, 447
TCP header values 818
thresholding 412, 439
invalid IP options (packet decoder) 622
invalid RFC delimiters (HTTP Inspect option) 572
IP (Internet Protocol) 258
IP addresses
excluding 400
host view 1023
in rules 766
in variables and rules 399
syntax for searches and reports 1557
IP defragmentation 613
exploits 614
generator ID 490
target-based policies 615
IP header values 812
IP layer (packet view) 301
IP_Proto (rule keyword) 815
IPopts (rule keyword) 814
IPS
analyzing intrusion events 261
architecture 256
deployments 262
introduction 253
intrusion event responses 261
intrusion events 267
intrusion policies 331, 477
preprocessors 484
IPS analyst access 1683
IPS event rate monitoring 1894, 1913, 1917
IPv4 packet filters
PEP policies 1516
IPv4 PEP rules
PEP policies 1520
IPv6 packet filters
PEP policies 1518
IPv6 PEP rules
PEP policies 1522
IPV6 Teredo tunneling
detecting 622
isdataat (rule keyword) 837
ISDN for Linux 258
itype (rule keyword) 817

Keysight, EX-1004
Page 2109

J

jumbo frames 229

K

keyword
- ack 818
- asn1 829
- base64_data 859
- base64_decode 859
- byte_extract 840
- byte_jump 794
- byte_test 797
- classtype 772
- content 776
- content depth option 783
- content distance option 781
- content fast_pattern option 789
- content http_client_body option 788
- content http_cookie option 787
- content http_header option 786
- content http_method option 786
- content http_raw_cookie option 787
- content http_raw_header option 786
- content http_raw_uri option 786
- content http_status_code option 788
- content http_status_message option 788
- content http_uri option 785
- content nocase option 778
- content offset option 782
- content rawbytes option 779
- content within option 781
- cvs 839
- detection_filter 849
- dsize 837
- file_data 858
- flags 819
- flow 820
- flowbits 851
- fragbits 813
- fragoffset 839
- http_encode 855
- icmp_id 817
- icmp_seq 817
- icode 818
- ID 814
- IP_Proto 815
- IPopts 814
- isdataat 837
- itype 817
- metadata 810
- msg 772
- pcre 800
- priority 772
- react 846
- reference 776
- resp 845
- rpc 828
- sameip 839
- seq 822
- ssl_state 826
- ssl_version 827
- stream_reassemble 824
- stream_size 823
- tag 850
- threshold (deprecated) 850
- tos 815
- ttl 815
- urilen 830
- window 822

keywords, in rules 769

L

language 1754
last seen (host view) 950, 1023
last successful login 43
latency thresholding
- configuring 719, 724
- troubleshooting options 493
- understanding 716

layer (host profile) 975
layers
- adding 508
- advanced settings 504
- configuring 507
- RNA recommendations 499
- rule settings (removing) 499
- sharing 496
- working with 494

LDAP 1628, 1631, 1656
- and RUA 1456, 1464, 1478
- authentication profiles 1741
- authentication server 1644
- examples 1657
- group access settings 1651
- logging in 40
- managing user accounts 1680
- shell access 1654

Keysight, EX-1004
Page 2110

legacy reassembly (stream option) 637
licenses
adding 1788
feature licenses 1787
managing 1780
NetFlow 1790, 1793
product licensing widget 87
requesting 1788
RNA hosts 1791
RUA users 1781, 1792
verifying 1786
viewing 1790
lights-out management 1823
line graphs (flow data) 1108
link mode, editing 1803
link state propagation 228
monitoring 1919
link state propagation monitoring 1894
Linux SLL (cooked mode) 258
list (host attribute type) 988, 990
listeners 1116
load balancers
exclusion from monitoring 892
load balancers and RNA 892
local time 90
local vs. remote policies 107
logging into the appliance 37
using LDAP 40
logging level (firewall rule) 1312, 1315, 1318
logging multiple packets per stream 714
logging out of the appliance 42
login banner 1755

M

MAC addresses
additional MAC detected (RNA event type) 1006, 1223
host profile 949
host view 1025
information change (RNA event type) 1008, 1224
mail relay host 1751
Maintenance access 1683
management interface 1801
management virtual network 1807
managing
3D Sensors 103
3Dx800s 112, 131, 133
Defense Centers 165
managed sensor system settings 139
sensors 104, 118
time settings 145
managing rules 403
Master Defense Center 165
adding a Defense Center 174, 178
appliance groups 191
deleting a Defense Center 181
editing settings for the Defense Center 185
event aggregation 166
filter configuration 187
global policy management 170
health policies 172
intrusion policies 171
managing Defense Centers 165
policy limitations 173
RNA detection policies 171
system policies 172
system settings 193
updating 1835
maximum active responses (TCP stream option) 630
maximum chunk size (HTTP Inspect option) 573
maximum fragment size (DCE/RPC option) 516
maximum response seconds (TCP stream option) 631
maximum TCP window (stream option) 634
maximum transmission unit (MTU) 613
memory cap reached (DCE/RPC option) 517
memory usage monitoring 1894, 1920
metacharacters 802
metadata (rule keyword) 810
metadata, in rules 810
Microsoft
%U encoding (HTTP Inspect option) 571
Active Directory 1456, 1457, 1464
IIS encoding (HTTP Inspect option) 571
Microsoft Active Directory 1632
monitored networks, configuring 889, 891
msg (rule keyword) 772
msg react keyword) 846
MTU 613
multiple inline interface sets 234
multi-rule search engine 259
multi-slash obfuscation (HTTP Inspect option) 572

N

NAT
and RNA 892
excluding devices from monitoring 892
working in NAT environments 117

Keysight, EX-1004
Page 2111

Nessus scans
deleting scan targets 1443
editing scan targets 1442
importing scan results 1445
scan targets 1442
NetBIOS name
host profile 949
host view 1023
NetBIOS name change (RNA event type) 1008, 1223
NetFlow
adding devices 1151, 1816
information in NetFlow records 884
licenses 1787, 1790, 1793
limitations 882, 885, 894
prerequisites 887, 1146
RNA detection policies 891, 895, 909, 1153
system settings options 1777
Network (TCP stream option) 633
network address translation, *see* NAT
network behavioral analysis 1118
network compliance widget 85
network devices (on the network map) 928
network gateway 1801
network interfaces 1803
network layer 258, 301
network map
Bugtraq ID 932
custom topology 936
CVE ID 932
definition 882
host attributes 935
hosts 926
introduction 923
network devices 928
RNA ID 932
services 930
SIDs 932
understanding 924
viewing hosts 927
vulnerabilities 932
network settings
configuring 1800
using DHCP 1799, 1802
using static 1803
network surveys (white lists) 1168
Networks (DCE/RPC option) 522
networks to monitor 905, 909
new client application (RNA event type) 1008, 1223
new host (RNA event type) 1008, 1223
new network protocol (RNA event type) 1008, 1224
new operating system (RNA event type) 1008
new TCP service (RNA event type) 1009, 1224
Nmap
adding an Nmap instance 1334
introduction 1419
managing 1437
on-demand scans 1440
remediations 1333, 1335, 1431, 1439
sample scanning profiles 1424
scan instances 1428, 1437
scan targets 1430
setting up 1428
Nmap scans, scheduling 1722
no_alert_large_fragments (RPC decoder) 579
nocase content option (rule keyword) 778
non-RFC characters (HTTP Inspect option) 573
normalizing inline traffic 604
notes in the host profile 986
notification
email alerting 756
SNMP alerting 748
syslog alerting 752

O

obsolete TCP options 624
offset content option (rule keyword) 782
on-demand scans 1440
OpenLDAP 1456, 1457, 1464, 1481, 1632, 1639, 1649
operating system identities
editing 955
resolving conflicts 956
viewing 954
operating systems, detected by RNA 2020
OPSEC alerting
intrusion rules 415
OPSEC, *see* Check Point OPSEC SAM
order of execution, of preprocessors 486
OS (host view) 1024
OS breakdown 1001
OS custom fingerprints 1363
OS information update (RNA event type) 1224
OS mappings 966, 968, 1367, 1374, 1410, 1411
OS names (host view) 1024
OS vendor (host view) 1024
OS version (host view) 1024
OS, TCP policy option (stream) 634
overlap limit (stream option) 634

Keysight, EX-1004
Page 2112

P

packet bytes (packet view) 306
packet decoder
  capturing packets 257
  experimental TCP options 623
  generator ID 489
  invalid IP options 622
  invalid TCP options 624
  obsolete TCP options 624
  other protocol header anomalies 624
  T/TCP options 624
  Teredo traffic (detecting) 622
  understanding 621
packet latency thresholding 716
  settings 718
packet size performance boost, TCP stream option 636
packet type performance boost (TCP stream option) 630
packet view 290
  copying events to the clipboard 286, 291
  datalink layer 300
  definition 282
  deleting events 286, 291
  event information 293
  frame information 299
  network layer 301
  packet bytes 306
  portscans 660
  rule actions 295
  setting thresholds 297
  suppression options 298
  TCP packets 304
  transport layer 303
  UDP packets 305
packets
  capturing 257
  data link layer 258
  decoding 257, 621
  packet bytes 306
  processing 258
password settings 1681
passwords
  changing 1689
  failed logins 1681
  force reset 1682
  password options 1681
  strength check options 1682
passwords, changing 43
patches (for vulnerabilities) 982, 983
payloads
  detected by RNA 2039
  detection 1006
  in client applications 971, 1053
  in flow trackers 1246
  in flows 1090
  in services 960, 965, 1044
  payload types 2039
PCRE 802
  character classes 803
  examples 808
  in intrusion rules 800
  metacharacters 802
  modifier options 804
pcre (rule keyword) 800
PDF reports 1547
peer manager 159
PEP
  applying policies 1527
  creating policies 1513
  deleting policies 1529
  editing policies 1526
  introduction 1510
  IPv4 packet filters 1516
  IPv4 PEP rules 1520
  IPv6 packet filters 1518
  IPv6 PEP rules 1522
  managing traffic 1510
  target interface sets 1514
  understanding 1511
  viewing active policy details 1529
  working with policies 1525
PEP status monitoring 1894
performance boost
  UDP option (stream) 646
performance monitor 726
performance statistics 1883, 1886
Perl-compatible regular expressions 800, 802
  options 804
platform requirements conventions 56
policy (DCE/RPC option) 522
Policy & Response Admin access 1684
Policy by VLAN or Network 396
port exclusion
  limitation with NetFlow 895
ports
  detected services 959, 960
  exclusion from RNA monitoring 893
  in variables 398
  in variables and intrusion rules 401
ports ignoring small segments (stream option) 636
portscans 651
  configuring 655
  decoy portscan 653
  distributed portscan 653

Keysight, EX-1004
Page 2113

- generator ID 489
- packet view 660
- portscan events 658
- portscan types 653
- portsweep 653
- sensitivity levels 654

portsweep 653
power supply monitoring 1894, 1922
PPP 258
PPPOE 258
predefined host attributes 985
preferences 43
- changing passwords 43
- event preferences 45
- home page 51
- intrusion policies 1753
- time zone 50
- widgets 66

preprocessor
- events 488
- execution order 486

preprocessor rules, setting rule states 435
preprocessors
- adaptive profiles 698
- application layer 513
- Back Orifice detection 648
- DCE/RPC 514
- DNS 531
- FTP Telnet 538
- generator IDs 488
- inline normalization 604
- introduction 477, 713
- intrusion policies 484
- IP defragmentation 613
- order of execution 486
- performance monitor 726
- processing packets 258
- reading events 488
- sensitive data detection 677
- SSH preprocessor 590
- SSL preprocessor 596

primary Defense Center 156, 157, 158
priority (rule keyword) 772
processing packets 258
product licensing widget 87
product updates widget 88
profile conditions (in traffic profiles) 1130
profiling time window (PTW) 1118, 1127
- options 1127

prohibit packet transfer to the Defense Center 141
protocol breakdown 1000
protocols
- detected services 959
- host profile 975
- white lists 1182

PTW 1118, 1127
purging the RNA/RUA databases 2014

R

RADIUS 1662
- connection settings 1663
- creating authentication objects 1662
- custom attributes 1669
- editing authentication objects 1674
- examples 1671
- shell access 1668
- testing authentication 1670
- user roles 1665

rate-based attack prevention
- configuring 673
- dynamic rule states 452
- intrusion rules 413
- policywide setting 662
- thresholding 668
- with detection_filter keyword 666
- with other filters 666
- with thresholding 669

rawbytes content option (rule keyword) 779
react (rule keyword) 846
read only access 1683
Real-time Network Awareness, *see* RNA
Real-time User Awareness, *see* RUA
reassembly threshold (DCE/RPC option) 516
reference (rule keyword) 776
refresh interval 1604
registration ID 159
registration key 117, 159, 176, 179
regular expressions 728
regular expressions, in intrusion rules 800
remediation events 1343
- field descriptions 1347
- searching 1349
- understanding 1347
- viewing 1344

remediations 1296
- adding a Cisco IOS instance 1320
- Cisco IOS routers 1319
- Cisco PIX 1327
- deleting modules 1297
- installing new modules 1297
- Nmap remediations 1431, 1439
- status events 1343

Keysight, EX-1004
Page 2114

remote (vulnerability detail) 981, 1060
remote backups 1991
remote management 1807, 1810
remote reports 1536
remote storage 1817
  local storage 1817
  NFS 1818
  SMG 1821
  SSH 1820
  system settings options 1777
remote storage, reports 1537
remote vs. local policies 107
replace rules 792
report designer 1542
report profiles
  creating 1542
  deleting 1549
  introduction 1538
  predefined 1538
  viewing generated reports 1534
reports
  clipboard reports 316
  creating report profiles 1542
  deleting 1535
  deleting report profiles 1549
  downloading 1535
  event graphs 273
  footers 1548
  from event views 1531
  incidents 327
  introduction 1531
  intrusion policy, comparing 354, 357
  intrusion policy, using 350
  predefined profiles 1538
  remote reports 1536
  remote storage 1537
  report profiles 1538
  scheduling reports 1721
  viewing generated reports 1534
requesting feature licenses 1788
require TCP handshake (stream option) 636
Requires conventions 56
reset_dest (resp keyword) 845
reset_source (resp keyword) 845
resetting management 127
resp (rule keyword) 845
response groups 1351
  activating and deactivating 1353
  creating 1352
  deleting 1353
  editing 1353
restarting the appliance 1806
restoring from backup files 1993
Restricted event analyst access 1684, 1685
right-click menus 53
RNA 35
  active detection 1359
  analyzing RNA events 1004
  and Crossbeam 878
  client applications 1050
  compliance events 1271
  compliance policies 1212
  compliance policy remediations 1296
  compliance responses 1281
  compliance rules 1212
  data collection 880
  detected client applications 2036
  detected operating systems 2020
  detected payloads 2039
  detected services 2027
  detecting custom services 1381
  enhancing detection 1355, 1356
  event statistics 996
  excluding ports 907
  flow data 884, 1071
  host attributes 882, 987
  host profiles 882, 944
  impact flags 885
  introduction 35, 878, 1140
  load balancers and NAT devices 892
  misidentifications 1357, 1358
  network map 882, 923
  passive detection 1359
  performance statistics 1886, 1887, 1889
  predefined workflows 1580
  purging the database 2014
  remediation events 1343
  RNA detection policies 889
  RNA event alerts 1293
  RNA events 883, 994
  RNA ID 932, 981
  RNA settings in system policies 1756
  services 1040
  subnet detection (system policy) 1757
  system policy settings 888
  third-party vulnerabilities 1064
  traffic profiles 884, 1071
  understanding 879, 1141
  vulnerabilities 886, 1057
  vulnerability ID 1060, 1067
  white lists 1156
rna (executable name) 1881
RNA analyst access 1683
RNA detection policies
  applying 919
  configuration strategies 889

Keysight, EX-1004
Page 2115

creating 900, 1152
deleting 921
editing 921
exporting 2003
introduction 889
NetFlow options 909
subnet detection 911
RNA detectors 1381
activating 1398
creating 1387
deleting 1400
editing 1399
exporting 2005
filtering 1395
groups 1402
internal detectors 1401
navigating 1397
sorting 1395
RNA events
analyzing 1004
definition 883
event breakdown 998
event statistics 996
event types 1006, 1011
field descriptions 1014
introduction 994
OS breakdown 1001
protocol breakdown 1000
searching 1016
service breakdown 1000
statistics summary 997
viewing 1012
workflows 1002
RNA events, purging the database 2014
RNA for Red Hat Linux 114
RNA health monitoring 1894, 1924
RNA host usage monitoring 1894, 1925
RNA ID 932
vulnerability detail 981
RNA recommended rules 468
configuring 472
generating recommendations 472
layer 503
scheduling 1725
RNA subnet detection settings 1757
rpc (rule keyword) 828
RPC decoder 578
generator ID 489
RPC proxy traffic only (DCE/RPC option) 523
RSS feed widget 89
rst_both (resp keyword) 845
RUA 1452
and 3D Sensors 1459, 1465, 1487
and Defense Centers 1460
and LDAP 1456, 1464, 1478
compliance rules 1248
configuring 1474
correlating user identity 1468
deployment scenarios 1454, 1463, 1466
event searches 1506
examples 1468, 1469, 1470, 1471
host and user histories 1467
information retrieved 1466
introduction 1452
licenses 1475, 1787, 1792
limitations 1473
login types 1489
predefined workflows 1584
purging RUA events 2014
RUA agents 1458, 1464, 1483
RUA events 1467, 1500, 1501, 1505, 1506
RUA users 1467, 1489, 1490, 1493
system policy settings 1764
understanding 1453
user identity 1490
user searches 1497
RUA users 1489, 1490, 1493
detailed information 1495
searching 1497
rule categories 427
rule classifications
exporting 2001
intrusion event searches 314
rule comments
in an intrusion policy 417, 466
rule comments (in an intrusion policy) 417
rule details 410
rule editor 861
rule headers 764, 769
building 764
specifying IPs 766
specifying ports 768
types 765
rule latency thresholding 720
settings 723
rule states
changing 435
rules
categories 427
creating 861
destination IPs 766
directional operators 769
filtering (in an intrusion policy) 417
keywords 769
managing in an intrusion policy 403
metadata 810

Keysight, EX-1004
Page 2116

setting rule state 435
source ports 768
specifying IPs 766
tuning 266
Rules intrusion policy page 497

S

SAM, see Check Point OPSEC SAM
sameip (rule keyword) 839
saved searches 1554
deleting 1554, 1555
loading 1554
scan results in host profiles 992
scanning 1418
analyzing results 1451
monitoring scans 1444
Nmap 1419, 1428
scan results 1443
searching for scan results 1449
scheduling tasks
backups 1702
deleting 1730
editing 1730
introduction 1699
intrusion policy applications 1719
Nmap scans 1722
recurring tasks 1700
reports 1721
RNA recommended rules 1725
SEU imports 1716
software downloads 1705
software installs 1708
software pushes 1706
software updates 1704
using the calendar 1727
using the task list 1729
VDB downloads 1711
VDB installs 1714
VDB pushes 1712
VDB updates 1710
viewing 1727
searching
active scan results 1449
audit records 1978
CIDR notation 59
client applications 1054
compliance events 1276
custom tables 1569
events 1550, 1551
flow data 1094
health events 1965
host attributes 1038
hosts 1027
intrusion events 309
intrusion rules 869
loading a saved search 1554
remediation events 1349
RNA events 1016
RUA events 1506
RUA user identity 1497
services 1045
SEU import logs 1865
specifying detection engines 1559
specifying ports 1558
third-party vulnerabilities 1068
time constraints 1556
vulnerabilities 1062
white list violations 1208
white lists 1201
wildcards 1555
secondary Defense Center 156, 157, 158
SecurID 39, 41
sensitive data
configuring custom data types 693
configuring detection 683
data type options 680
defining custom data patterns 690
detecting 677
detecting in FTP traffic (special case) 689
global detection options 678
predefined data types 681
sensitive data, monitoring RNA services 686
sensor groups
creating 137
deleting 139
editing 138
managing 137
sensor list 120
seq (rule keyword) 822
server fingerprints 1370
server ports (stream reassembly option) 640
server services (stream reassembly option) 640
service breakdown 1000
service detectors 1381
service identities
editing 965
resolving conflicts 967
service timeout (RNA settings) 1760
services
details 962
detected by RNA 2027
detected services 960

Keysight, EX-1004
Page 2117

field descriptions 1043
host profiles 958
introduction 1040
misidentified 1358
network map 930
payloads 960, 965, 1044
RNA detectors 1381
searching 1045
service identity 965
source type 1045
understanding 1043
viewing 1041, 1043
viewing from the network map 931
white lists 1178
serving time to managed sensors 1771
session logouts 39
sessions, current sessions widget 71
set operating system definition (RNA event type) 1011, 1227
set service definition (RNA event type) 1012, 1227
SEUs 363, 1848
automatic one-time imports 1853
detail view 1863
import log 1859, 1861
manual one-time imports 1851
recurring Imports 1855
searching import logs 1865
SEUs, scheduling imports 1716
shared host profiles 1189
adding 1184
creating 1190
deleting 1195
editing 1192
shared object rules 761
generator ID 489
shell access 1654, 1668
shutting down the appliance 1806
SIDs
mapped to vulnerabilities 1060
network map 932
vulnerability detail 981
sliding time window 47
SLIP 258
small segment size (stream option) 636
SMB invalid shares (DCE/RPC option) 523
SMB maximum AndX chain (DCE/RPC option) 523
SMTP decoder, generator ID 490
SNMP alerting 748
configuring 749
creating 1290
intrusion rules 415
SNMP decoding 582
SNMP polling 1766
snooze periods 1250, 1251
Snort ID, *see* SID
software sensors 110
software updates 1831
installing 1833
scheduling downloads 1705
scheduling installs 1708
scheduling pushes 1706
scheduling updates 1704
solution (vulnerability detail) 982
source (detected services) 961
source IPs
in rules 766
source ports
in rules 768
source type
host profiles 953
services 964
SSH preprocessor 590
configuring 594
options 591
SSL preprocessor 596
configuring 599
understanding 597
SSL rule keywords 825
ssl_state (rule keyword) 826
ssl_version (rule keyword) 827
standard text rules 761, 861
generator ID 489
stateful inspection
anomalies (stream option) 635
static time window 47
statistics
events 269
host 271
performance 1883, 1886
status, tasks 2016
store events only on Defense Center 141
stream preprocessor
configuring TCP stream preprocessing 640
configuring UDP session tracking 645
generator ID 490
state-related TCP exploits 628
stream reassembly options 638
target-based policies 631
TCP global options 630
TCP policy options 633
TCP session tracking 627
UDP session tracking 645
stream reassembly 637
generator ID (stream) 490
reassembly options (stream) 638
stream-based attacks 638

Keysight, EX-1004
Page 2118

stream_reassemble (rule keyword) 824
stream_size (rule keyword) 823
subnet detection 897
  auto-detect 891
  identifying network segments 891
  limitation with NetFlow 895
  managing 911
  manually applying recommendations 915, 916
  manually generating recommendations 915
  process 898
  recommendations 881
subnet detection, system policy settings 1757
subnet recommendations
  automatically generating 918
  evaluating 916
Sun directory server 1456, 1457, 1464, 1632
suppression
  configuring 447
  creating 448
  deleting 450
  in the packet view 298
  introduction 439
suspended standard text rule (GID) 490
syslog
  alerting 754
  filer examples 1984
  filter syntax 1982
  priority levels 753
  severity levels 1287
  syslog facilities 753, 1286
  viewing 1981
system daemons 1878
system default variables 372
system load widget 90
system management
  backup and restore 1985
  IPS performance statistics 1883
  RNA performance statistics 1886
  software updates 1831
  system policies 1732
  system settings 1775
  uninstalling software updates 1845
  VDB updates 1831
system monitoring
  disk usage 1874
  host statistics 1870
  introduction 1869
  system processes 1875
  system status 1874
system monitoring, system load widget 90
system policies 1732
  access list 1738
  applying 1736
  audit log 1740
  authentication profiles 1741
  creating 1733
  custom login banner 1755
  dashboard settings 1745, 1746
  database limits 1747
  deleting 1737
  detection policy settings 1757
  DNS cache 1750
  editing 1735
  exporting 2003
  language 1754
  mail relay host 1751
  multiple fingerprint settings 1762
  RNA data storage 1759
  RNA settings 1756
  RUA settings 1764
  service vulnerability mapping 1773
  time sync 1768
  vulnerability impact assessment settings 1762
system policy
  SNMP polling 1766
system processes
  understanding 1877
  viewing 1875
system settings 1775
  appliance information 1777
  introduction 1775
  licenses 1780
  managing with a Defense Center 139
  Master Defense Center 193
  NetFlow- enabled devices 1151
  NetFlow-enabled devices 1816
  network settings 1800
  remote management 1807, 1810
  remote storage 1817
  restarting the appliance 1806
  shutting down the appliance 1806
  time sync 1813
system time widget 90
system utilities and executables 1880

T

T/TCP options 624
tab obfuscation (HTTP Inspect option) 572
tab URI delimiter (HTTP Inspect option) 573
table view of events 1574
  features 1596
  intrusion events 281

Keysight, EX-1004
Page 2119

tag (rule keyword) 850
tagged packets
  generator ID 489
  tag keyword 850
tap mode 227
target (intrusion policy) 368
task queue 2016
  managing 2018
  viewing 2016
TCP 258
  experimental options 623
  invalid options 624
  new TCP service (RNA event type) 1009, 1224
  obsolete options 624
  port closed (RNA event type) 1009, 1224
  port timeout (RNA event type) 1009, 1224
  service information update (RNA event type) 1009, 1224
  session tracking (stream) 627
  T/TCP options 624
TCP header values 818
TCP service confidence update (RNA event type) 1009
TCP services (RNA event type) 1009, 1224
technical description (vulnerability detail) 982
Telnet decoder 538
  generator ID 490
Teredo traffic, detecting (packet decoder) 622
testing authentication 1656
text (host attribute type) 988
third-party product mappings
  exporting 2004
third-party tools 1408
  host identification 1357
  patch tools 1410
  service identification 1414
  vulnerabilities 1358, 1412
third-party vulnerabilities
  exporting mappings 2005
threshold (deprecated rule keyword) 850
thresholding
  adding a new threshold 442
  by rule 439
  configuring 439
  deleting 445
  global 705
  in the packet view 297
  introduction 439
  options 440
  packet latency thresholding 716
  policywide options 708
  rule latency thresholding 720
  with dynamic rule states 668
  with rate-based attack prevention 669
time constraints (in searches) 1556
time sync 1768, 1771
  Master Defense Center 196
  serving time 1771
  setting manually 1813
  setting time on a managed sensor 145
time sync, system time widget 90
time synchronization monitoring 1894, 1927
time window 47, 1598
  changing the default 1603
  configuring 1597
  default setting 47
  pausing 1606
time zone 50
timeout
  defragmentation option 617
  DNS cache option 1751
  host timeout 1007, 1223
  session logout 42
  session logouts 39
  TCP policy option (stream) 634
  TCP port timeout (RNA event type) 1009, 1224
  UDP port timeout (RNA event type) 1009, 1224
  UDP session tracking (stream) 646
timestamp
  host view 950, 1023
  service view 1043
title (vulnerability detail) 981
Token Ring 258
tos (rule keyword) 815
traffic profiles
  activating 1129
  creating 1118
  definition 884
  editing 1130
  host profile qualifications 1124
  host view 1025
  introduction 1071
  profile conditions 1122, 1130
  saving 1129
  specifying conditions 1122
  viewing 1138
traffic status monitoring 1894, 1927
Transmission Control Protocol, *see* TCP
transparent inline mode 226
transport layer 258, 303
  preprocessor 627
transport protocol (RNA event type) 1009, 1224
trend graphs 273
troubleshooting options, intrusion policies 492
TTL
  host profile 949
  rule keyword 815

Keysight, EX-1004
Page 2120

U

UDP 258
  and Back Orifice 648
  new service (RNA event type) 1009, 1224
  packet view 305
  port closed (RNA event type) 1009, 1224
  port timeout (event type) 1009, 1224
  session tracking 645
UDP service confidence update (RNA event type) 1010
UDP service information update (RNA event type) 1010, 1224
unauthorized activity 40
uninstalling software updates 1845
unique NAT ID 117, 1810
unreviewing events 278
updating
  managed 3D Sensors 1838
  unmanaged 3D Sensors 1840
updating the software 1833
  product updates widget 88
URI parsing (HTTP Inspect option) 573
urilen (rule keyword) 830
URL (host attribute type) 988
user accounts
  account management 1626
  deleting 1690
  editing 1684
  externally authenticated user accounts 1680
  LDAP users 1680
  managing 1675
  menu access per access type 1690
  password options 1681
  password settings 1681
  privileges 1629
  shell access 1654, 1668
  user authentication 1626
user authentication 1626
  external authentication 1628
  internal authentication 1628
User Datagram Protocol, *see* UDP
user history 973
user identity (host view) 1023
user preferences 43
  changing event preferences 45
  changing passwords 43
  home page 51
  time zone 50
UTF-8 (HTTP Inspect option) 571

V

variables
  and detection engines 213
  assigning values for detection engines 215
  creating 377
  creating in detection engines 217
  custom variables 391
  deleting 389
  editing 374
  predefined 372
  system default 372
  using IP addresses 398
VDB updates 1846
  scheduling downloads 1711
  scheduling installs 1714
  scheduling pushes 1712
  scheduling updates 1710
version (detected services) 953, 960, 964
violations (white lists) 1204
VLAN
  host profiles 972
  intrusion policy filtering 393
  tag information update (RNA event type) 1010, 1224
VLAN ID (host view) 1023
vulnerabilities
  activating and deactivating 982, 984
  deactivating 932, 1061
  definition 886
  field descriptions 1060, 1067
  host profiles 978
  impact qualification 982
  introduction 1057
  invalidating 932
  mapped to SIDs 1060
  misidentified 1358
  network map 932
  searching 1062
  third-party tools 1358, 1412
  understanding 1059
  viewing 1057
  vulnerability details 980
vulnerabilities, third-party
  host profiles 978
  introduction 1064
  predefined workflows 1582
  searching 1068
  understanding 1067
  viewing 1065
vulnerability database
  updating 1831
vulnerability impact 981

Keysight, EX-1004
Page 2121

vulnerability impact qualification (RNA event type) 1012
vulnerability mapping for services 1773
vulnerability mappings 966, 968, 1367, 1374, 1410, 1411
vulnerability set invalid (RNA event type) 1012, 1227
vulnerability set valid (RNA event type) 1012, 1227

W

web browser requirements 37
webroot directory traversal (HTTP Inspect option) 572
white list events 1197
  viewing 1197
white list targets 1159
  configuring network segments 1171
  creating 1171
  deleting 1174
  editing 1174
white lists 1190
  adding to compliance policies 1265
  basic information 1170
  client applications 1180
  configuring targets 1171
  creating 1167
  creating from host profiles 977
  creating targets 1171
  definition 1156
  deleting 1189
  deleting host profiles 1187
  editing 1188
  events 1197
  field descriptions 1200, 1207
  global host profiles 1161, 1175
  host profiles 1160, 1175
  host view 1026
  introduction 1156
  managing 1188
  modifying host profiles 1184
  network compliance widget 85
  operating systems 1161, 1176
  protocols 1182
  searching 1201
  services 1178
  shared host profiles 1163, 1184, 1189
  surveying your network 1168
  target evaluations 1164
  understanding 1158, 1200
  violations 976, 1165, 1204, 1208
  white list events 1165
  white list events widget 91
  white list violations 1207
whois 1027
widgets 62
  adding to a dashboard 99
  appliance status 69
  availability 63
  compliance events 69
  current sessions 71
  custom analysis 71
  deleting 101
  disk usage 82
  interface status 70
  interface traffic widget 83
  intrusion events 83
  minimizing and maximizing 101
  moving 101
  network compliance 85
  predefined 63, 67
  preferences 66
  product licensing 87
  product updates 88
  RSS feeds 89
  system load 90
  system time 90
  white list events 91
wildcards (in searches) 1555
window
  rule keyword 822
within content option (rule keyword) 781
workflows
  and custom tables 1569
  and flow data 1103, 1105
  audit workflow 1586
  common functionality 1594
  compliance workflows 1585
  components 1574
  compound constraints 1609
  constraining events 1606
  constraining time 1597
  creating 1618
  custom 1573
  custom flow data workflows 1621
  deleting 1625
  editing 1624
  flow data workflows 1582
  health event workflows 1586
  introduction 1573
  intrusion event workflows 1576
  intrusion events 280
  navigating between workflows 1614
  navigating to other pages 1613
  other workflows 1586

Keysight, EX-1004
Page 2122

predefined 1573
predefined vs.custom tables 1576
predefined vs.custom workflows 1575
RNA event 1002
RNA workflows 1580
RUA workflows 1584
saved custom workflows 1587
selecting 1590
sorting pages 1611, 1612
table view of events 1596
third-party vulnerability workflows 1582
using bookmarks 1615
using workflows 1589
viewing custom workflows 1623
workflow toolbar 1592
workflows, default workflows 50

X

xlink2state 582

Keysight, EX-1004
Page 2123