



Juniper Networks Network and Security Manager

Administration Guide

Release

2012.2



Modified: 2019-05-30

Revision 5

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Network and Security Manager Administration Guide
Copyright © 2019 Juniper Networks, Inc. All rights reserved.

Revision History
May 30 2019— Revision 6
February 02, 2017— Revision 5
November 20, 2014— Revision 4
January 21, 2014— Revision 3
January 20, 2014— Revision 2
January 15, 2013— Revision 1

The information in this document is current as of the date on the title page.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About This Guide	xli
	Objectives	xli
	Audience	xli
	Conventions	xl ii
	Documentation	xl iii
	Requesting Technical Support	xl v
	Self-Help Online Tools and Resources	xl v
	Creating a Service Request with JTAC	xl vi
Part 1	Getting Started with NSM	
Chapter 1	Introduction to Network and Security Manager	3
	About NSM	3
	Security Integration	4
	Device Configuration	4
	Network Organization	4
	Role-Based Administration	4
	Centralized Device Configuration	5
	Device Management	6
	Importing Devices	6
	Device Modeling	6
	Rapid Deployment	6
	Policy-Based Management	6
	Error Prevention, Recovery, and Auditing	7
	Device Configuration Validation	7
	Policy Validation	7
	Atomic Configuration and Updating	7
	Device Image Updates	8
	Auditing	8
	Complete System Management	8
	VPN Abstraction	8
	Integrated Logging and Reporting	9
	Monitoring Status	9
	Job Management	9
	Technical Overview	10
	Architecture	10
	User Interface	11
	Management System	11
	Managed Devices	14
	Firewall and IDP (ScreenOS/IDP) Devices	14
	Devices Running Junos OS	17

	SSL VPN Secure Access Products	20
	Juniper Networks IC Series Unified Access Control Appliances	21
	Extranet Devices	22
	Distributed Data Collection	22
	Device Schemas	22
	Security	23
	Scaling and Performance	23
	Working in the User Interface	23
	Characters Not Supported in Login Passwords	24
	Managing Blocked Login Attempts	24
	Configuring UI Preferences	24
	Disabling GUI Validation	24
	Customizing Validation	26
	UI Overview	26
	Navigation Tree	27
	Common Tasks Pane	27
	Main Display Area	27
	Menu Bar	28
	Toolbar	28
	Status Bar	28
	NSM Modules	28
	Investigate Modules	28
	Configure Modules	30
	Administer Modules	34
	Validation Icons in the User Interface	34
	Validation and Data Origination Icons	35
	Working with Other NSM Administrators	36
	Searching in the User Interface	36
	Contains String [C] Search Mode	37
	Starts With [S] Search Mode	37
	Regular Expression [R] Search Mode	38
	IP [I] Search Mode	39
	Search for an Exact Match (E)	40
	Global Search	41
	New Features in 2012.2	42
Chapter 2	Planning Your Virtual Network	43
	Configuring Devices Overview	43
	Importing Existing Devices	44
	Modeling New Devices	45
	Editing a Device Configuration	46
	Configuring IDP-Capable Devices Overview	47
	Common Criteria EAL2 Compliance	47
	Guidance for Intended Usage	47
	Guidance for Personnel	47
	Guidance for Physical Protection	47
	Supported IDP-Capable Devices	47
	Enabling Jumbo Frames (ISG1000 Only)	48

	Enabling IDP Functionality	49
	Adding an ISG2000/ISG1000 Security Device with a Security Module	49
	Updating Attack Objects	49
	Adding Objects (Optional)	50
	Configuring a Security Policy for IDP	50
	Reviewing IDP Logs	55
	Maintaining IDP	56
	Creating IDP-Only Administrators	56
	Simplifying Management	56
	Using Device Groups	57
	Using Device Templates	57
	Using Configuration Groups	58
	Merging Policies	58
	Using a Naming Convention	58
	Example: Using a Naming Convention for Devices	58
	Example: Using a Naming Convention for Address Objects	59
	Creating an Information Banner	59
	Adding an Information Banner	60
	Modifying an Information Banner	62
	Deleting an Information Banner	62
Chapter 3	Configuring Role-Based Administration	63
	Role-Based Administration	63
	Domains	63
	About Roles	64
	Using Role-Based Administration Effectively	65
	Enterprise Organizations	65
	Geographical Divisions	66
	NOC and SOC	66
	Administrator Types	66
	Service Providers	67
	Internal Network	68
	Managed Security Service Provider (MSSP)	68
	Configuring Role-Based Administration	68
	Creating Administrators	69
	Configuring General Settings	69
	Configuring Authorization	70
	RADIUS Authentication and Authorization	70
	Configuring Roles	76
	Creating Custom Roles	77
	Roles and Permissions	89
	Permissions Changes in Release 2008.1	89
	Permissions Changes in Release 2006.1	89
	Permissions Changes in Release 2005.3	90
	Assigning and Viewing Custom Roles	91
	Configuring a User Activity in a Custom Role	91
	Viewing Logged Administrators	92
	Forcing an Administrator to Log Out	92

	Creating Subdomains	93
	Viewing Current Domain Detail	93
	Example: Configuring Role-Based Administration	93
	Step 1: Create the Subdomains	94
	Step 2: Create the Subdomain Administrator	94
	Step 3: Create the Viewing and Reporting Administrator	95
	Step 4: Verify Administrator Accounts	95
Part 2	Integrating	
Chapter 4	Adding Devices	99
	About Device Creation	100
	Determine Device Status	101
	Verifying Device Configuration	102
	Managing the Device	102
	Before You Begin Adding Devices	102
	Importing Versus Modeling	103
	Importing Device Configurations	103
	Modeling Device Configurations	103
	Device Add Process	104
	Selecting the Domain	104
	Adding Single or Multiple Devices	105
	Specifying the OS and Version	106
	Determining Port Mode (ScreenOS Devices Only)	106
	Trust-Untrust Port Mode	107
	Home-Work Port Mode	107
	Dual-Untrust Port Mode	108
	Combined Port Mode	109
	Trust-Untrust-DMZ Port Mode	109
	Trust/Untrust/DMZ (Extended) Mode	110
	DMZ-Dual-Untrust Port Mode	111
	Port Mode Summary	112
	Changing the Port Mode	113
	Supported Add Device Workflows by Device Family	113
	Importing Devices	114
	Requirements	115
	Adding and Importing Devices with Static IP Addresses	115
	ScreenOS Devices	116
	IDP Sensors	117
	Junos Devices	118
	SA and IC Devices	120
	Adding Devices with Dynamic IP Addresses	121
	ScreenOS Devices	121
	IDP Sensors	122
	Adding and Importing an Infranet Controller or Secure Access Device	124
	Adding and Importing a Junos Device with a Dynamic IP Address	127

Verifying Imported Device Configurations	131
Using Device Monitor	131
Using Device Manager	131
Using Job Manager	132
Using Configuration Summaries	132
Modeling Devices	133
Requirements	134
Modeling a Device	134
Creating a Device Configuration	135
Activating a Device	135
Devices with Static IP Addresses	135
Devices with Dynamic IP Addresses	138
Using Rapid Deployment (ScreenOS Only)	142
Creating the Configlet	144
Installing the Configlet	147
Preparing the Device	147
Installing the Configlet	147
Updating the Device Configuration	149
Summarize Delta Configuration	149
Example: User Successfully Selects and Updates Two Devices with Delta Option	149
Example: User Selects Two Devices with Delta Option and One Device Fails	150
Example: User Selects Two Devices to Update Without the Delta Option	150
Example: User Selects Two Devices to Update with the Delta Option, But Has no Admin Privileges	150
Adding Vsys Devices	151
Placing the Root Device in a Global Domain or a Subdomain	151
Importing Vsys Devices	152
Modeling Vsys Devices	153
Adding L2V Root Systems	154
Adding an Extranet Device	155
Adding Clusters	155
Adding a Cluster Device Object	156
Adding Members to the Cluster	156
Adding ScreenOS or IDP Clusters	157
Adding Secure Access or Infranet Controller Clusters	157
Adding and Importing a Secure Access or Infranet Controller Cluster through Unreachable Workflow	158
Adding and Importing a Secure Access or Infranet Controller Cluster through Reachable Workflow	159
Adding Clusters of Routers Running Junos OS	160
Adding and Importing a Junos Cluster	161
Adding a Junos Cluster with Modeled Cluster Members	162
Activating and Updating a Modeled Junos Cluster	162
Example: Adding and Importing a Cluster	163
Adding the Cluster	163
Adding the Cluster Members	164

Importing the Cluster configuration	165
Example: Creating, Activating, and Updating a Cluster with Modeled Cluster Members	166
Adding the Cluster	166
Modeling the Cluster Members	166
Activating the Cluster Members	168
Updating the Cluster	170
Adding a Vsys Cluster and Vsys Cluster Members	170
Example: Adding a Vsys Cluster	170
Importing an SRX Series Cluster into NSM	172
Activating Management Access	173
Enabling Inbound Access	173
Importing an SRX Series Device in Cluster Mode	174
Importing an SRX Series Device in Virtual Chassis Mode	174
Removing Remnants from Previous Imports	175
Importing the Cluster Member	175
Adding a Blade Server	176
Automatic Grouping of Device Members Under the Blade Server	178
Manually Adding SA/IC Blade Members to the Blade Server	178
Manually Adding WXC Blade Members to the Blade Server	179
Adding Multiple Devices Using Automatic Discovery (Junos, SA, and IC Devices)	180
Adding a Device Discovery Rule	180
Running a Device Discovery Rule	181
Adding Many Devices Using CSV Files	181
Creating the CSV File	182
Devices with Static IP Addresses	183
Device with Dynamic IP Addresses	184
Undeployed Devices	187
Validating the CSV File	189
Importing Many Devices	190
Adding and Importing Many Devices with Static IP Addresses	190
Adding and Importing Many Devices with Dynamic IP Addresses	190
Modeling Many Devices	191
Using Rapid Deployment	192
Modeling and Activating Many Devices with Configlets	192
Activating Many Devices with Configlets	193
Adding Device Groups	194
Example: Creating a Device Group	194
Setting Up NSM to Work With Infranet Controller and Infranet Enforcer	195
Avoiding Naming Conflicts of the Authorization Server Object	195
Avoiding NACN Password Conflicts	198
Chapter 5 Configuring Devices	199
About Device Configuration	200
About Configuring Device Families	200
About Configuring Clusters, VPNs, Vsys Devices, Policies, and Shared Objects	200

Configuration Features	201
About the Device Editor	201
About Device Templates	202
About Configuration Groups	202
Editing Devices Using the Device Editor	202
Validation and Data Origination Icons	204
Configuring Device Features	205
Configuring ScreenOS/IDP Device Features	206
Configuring Secure Access or Infranet Controller Device Features . . .	208
Configuring Junos Device Features	209
Updating the Configuration on the Device	210
Using Device Templates	210
Modifying Values in Templates	211
Example: Creating and Applying a Device Template for DNS Settings . . .	212
Creating the Template	212
Applying the Template	213
Templates and Importing Devices	213
Promoting a Device Configuration to a Template	214
Changing Values Inherited from Templates	214
Reverting a Configuration to Default Values of a Template	215
Templates and Validation	215
Applying Multiple Templates	216
Example: Using Multiple Device Templates	216
Template Limitations	221
Maximum of 63 Templates	221
Device Groups	221
Default Values	221
Predefined Device Data	221
List Key Fields	222
Specifying the Order of List Entries	222
Combining Template Data with Device Object Data	223
Operations That Change the Sequence of Ordered Lists	224
Rules for Reordering Lists	225
Examples of Reordered Lists	225
Identifying Ordered List Entries That Do Not Match the Template or Configuration Group Order	228
Using the Template Operations Directive	229
Select OS Name Section	230
Select Devices Section	230
Select Template Section	230
Template Operation Section	231
Options Section	231
Template Operations Box Recommended Workflow	232
Removing Templates with the Template Operations Directive	233
Exporting and Importing Device Templates	234
Exporting a Device Template	234
Importing a Device Template	234

Using Configuration Groups	235
Creating and Editing Configuration Groups	236
Creating a Configuration Group	236
Editing a Configuration Group	238
Validating a Configuration Group	238
Ordered Lists and Wildcard Matching	239
Applying a Configuration Group	239
Excluding a Configuration Group	240
Editing a Device Object That Uses Configuration Groups	241
Deleting a Configuration Group	242
Adding Ordered List Entries Using Configuration Groups	242
Reordering Lists	243
Using Configuration Groups with Templates	243
Sharing Configuration Group Definitions Across Multiple Devices	244
Configuring Clusters	248
Configuring Cluster Objects Directly by Editing the Configuration	248
Configuring Cluster Objects Using Templates	248
Configuring Global Cluster Data with Configuration Groups (Junos Clusters Only)	249
Configuring Member-Level Data in a Junos Cluster	249
Configuring Junos Devices with Redundant Routing Engines	250
Configuring a Routing Engine	250
Viewing a Routing Engine Configuration	251
Overview of VRRP Support in NSM	252
Platforms on Which NSM Supports VRRP	253
Activating VRRP on a Device Interface	253
Defining a VSI as a VRRP interface	253
Managing Configuration Files	254
Viewing and Comparing Configuration File Versions	254
Updating the Device with a Configuration File Version	254
Updating the Archived Configuration File Version	254
Importing or Viewing the Current Version of the Configuration File for SRX Series Devices	255
Importing or Viewing the Current Version of the Configuration File for Devices Running ScreenOS	255
Automatic Import of Configuration Files	255
Chapter 6 Updating Devices	257
About Updating	257
How the Update Process Works	258
About Atomic Configuration—ScreenOS Devices	259
About Atomic Updating—ScreenOS Devices	260
About Atomic Configuration and Atomic Update—DMI-Compatible Devices	261

	About Implicit Updates (Secure Access and Infranet Controller Devices Only)	262
	Knowing When to Update	262
	Verifying Device Status in Device Monitor	263
	Connection Status	263
	Configuration Status	263
	Verifying Device Status in Device Manager	265
	Reviewing Logs	265
	Identifying Administrative Changes	266
	Reviewing Reports	266
	Using Preview Tools	266
	Running a Configuration Summary	267
	Using a Delta Configuration Summary	267
	Performing an Update	270
	Retrying a Failed Update	271
	Configuring Update Options	271
	Update Options for ScreenOS and IDP	271
	Update Options for DMI-Compatible Devices	272
	Tracking Device Updates	273
	Reviewing Job Information	274
	Device States During Update	276
	Understanding Updating Errors	276
Chapter 7	Managing Devices	279
	Managing Device Software Versions	280
	Upgrading the Device Software Version	280
	Upgrading a Device Software Version from NSM	282
	Upgrading a Device Software Version outside NSM	282
	Adjusting the Device OS Version	283
	Downgrading the Device OS Version	284
	Rolling back the Device OS Version	284
	Deleting the Device OS Version	285
	Upgrading Device Support	285
	Viewing and Reconciling Device Inventory	285
	Viewing the Device Inventory	286
	Comparing and Reconciling Device Inventory	287
	Managing Large Binary Data Files (Secure Access and Infranet Controller Devices Only)	289
	Uploading and Linking Large Binary Data Files	290
	Importing Custom Sign-In Pages	294
	Creating a Custom Sign-In Page	294
	Linking to a Custom Sign-In Page Shared Object	294
	Importing Antivirus Live Update Settings	295
	Uploading Live Update Settings	295
	295
	Linking to a Live Update File Shared Object	295
	Importing Endpoint Security Assessment Plug-in (ESAP) Packages	296
	Uploading ESAP Packages	296
	Linking to an ESAP Package Shared Object	297

Importing Third-Party Host Checker Policies	297
Uploading a Third-Party Host Checker Policy	297
Linking to a Third-Party Host Checker Policy Shared Object	298
Importing a Secure Virtual Workspace Wallpaper Image (Secure Access Devices Only)	298
Uploading a Secure Virtual Workspace Wallpaper Image	298
Linking to a Secure Virtual Workspace Wallpaper Image Shared Object	299
Importing Hosted Java Applets (Secure Access Devices Only)	299
Uploading a Java Applet	300
Linking to a Hosted Java Applet Shared Object	300
Importing a Custom Citrix Client .cab File (Secure Access Devices Only) . .	300
Uploading a Custom Citrix Client .cab File	300
Linking to a Custom Citrix .cab File Shared Object	300
Backing up and Restoring SA and IC Devices	301
Backing up an SA or IC Device	301
Restoring SA or IC Devices	301
Backing up multiple SA or IC Devices	302
Configuring Preferences for Backing up and Restoring SA or IC Devices . .	302
Viewing Backed up Versions for an SA or IC Device	302
Setting the RMA State on an SA/IC Device	303
Activating an SA/IC Device Set to the RMA State	303
Performing a Full Restore of an SA or IC Device	304
Managing User Sessions for SA and IC Devices	304
Activating Subscription Services	305
Managing the Attack Object Database	306
Updating the Attack Object Database	306
Updating Attack Objects for IDP-Enabled devices	307
Updating DI Attacks on ScreenOS 5.0 Devices	309
Using Updated Attack Objects	309
Verifying the Attack Object Database Version	309
Automatic Verification	310
Manual Verification	310
Managing Different Attack Database Versions	310
Example: Updating Devices with Different Attack Object Database Versions	311
Updating the IDP Detector Engine	311
Example: Confirm IDP Engine Version	312
Scheduling Security Updates	313
Example: Update Attack Objects and Push to Connected Devices . . .	315
Scheduling the Update	315
Example: Using Crontab to Schedule Attack Updates	316
Viewing Scheduled Security Updates in the Job Manager	316
Viewing Scheduled Security Updates in the Audit Log Viewer	317
Updating AV Pattern Files	317
Updating the Web Category List	317
Miscellaneous Device Operations	318
Launching a Web UI for a Device	319
Launching a Telnet CLI Window	319

	Rebooting Devices	319
	Refreshing DNS Entries	320
	Updating the Device Clock with an NTP Server	320
	Setting the Root Administrator on a Device	321
	Failing Over or Reverting Interfaces	322
	Setting the RMA State on a Device	322
	Managing Existing and Adding New Devices Using MIP IPv6 Addresses	323
	Converting Device Management IP Format from IPv4 to IPv6 Using an NSM MIP Address	323
	Adding a New Device with an IPv6 Address to the NSM Server with a MIP IPv6 Address	326
	Upgrading the OS Version During an RMA-Activate Device Workflow	330
	Troubleshooting a BGP Peer Session on a Device	330
	Reactivating Wireless Connections	331
	Finding Usages	331
	Managing ScreenOS Device Capabilities	331
	Abstract Data Model	332
	Data Model	332
	Data Model Schema	332
	Data Model Updating	333
	Data Model Importing	335
	Archiving and Restoring	337
	Archiving Logs and Configuration Data	337
	Restoring Logs and Configuration Data	338
	Managing Device Schemas Through the Juniper Update Mechanism	338
	Downloading Schemas	339
	Downloading Schemas Using the NSM UI	340
	Downloading Schemas Using the GUI Server CLI	341
	Selective Schema Loading in NSM	341
	Applying a Schema	342
Part 3	Managing	
Chapter 8	Configuring Objects	345
	About Objects	346
	Using Objects Across Domains	348
	Replacing Objects	348
	Working with Unused Shared Objects	349
	Searching for Unused Shared Objects	349
	Deleting an Unused Shared Object	349
	Working with Object Versions	350
	Searching For and Deleting Duplicate Objects	350
	Configuring Address Objects	351
	Viewing Address Objects	351
	Creating Address Objects	351
	Adding a Host Address Object	352
	Adding a Network Address Object	353
	Editing and Deleting Address Objects	353
	Replacing Address Objects	353

Adding an Address Object Group	354
Adding a Multicast Group Address Object	355
Adding Static DNS Host Addresses	355
Blocked Hosts	356
Configuring Application Objects	356
Viewing Predefined Application Objects	356
Viewing Predefined Extended Application Objects	357
Creating Custom Application Objects	358
Creating Application Group Objects	359
Validating Address Objects	360
Editing and Deleting Application Objects	361
Configuring Schedule Objects	362
Creating Schedule Objects	362
Configuring Access Profile Objects	363
Configuring Quality of Service Profiles	363
Creating a Quality of Service Profile	364
Deleting a Quality of Service Profile	364
Editing a Quality of Service Profile	365
Working with DI Attack Objects	365
Viewing Predefined DI Attack Objects	365
Viewing Attack Version Information for Attack Objects	366
Viewing Predefined DI Attack Object Groups	366
Updating Predefined DI Attack Objects and Groups	366
Creating DI Profiles	367
Working with IDP Attack Objects	369
Viewing Predefined IDP Attacks	369
Viewing Predefined IDP Attack Groups	370
Viewing Attack Version Information for Attack Objects and Groups	370
Updating Predefined IDP Attack Objects and Groups	371
Configuring Custom DI and IDP Attack Objects	371
Using the Attack Object Wizard	372
Copying and Editing Predefined Attack Objects to Create Custom Attack Objects	372
Configuring Attack Name and Description	372
Configuring Extended Information	374
Configuring External References	374
Configuring Target Platforms	375
Creating a Signature Attack Object	376
Configuring General Attack Properties	376
Configuring Attack Detection Properties	383
Configuring Header Match Properties	387
Configuring a Protocol Anomaly Attack Object	390
Configuring a Compound Attack Object	391
Configuring General Attack Properties	391
Configuring Compound Attack Members	392
Configuring the Direction Filter	395
Creating Custom DI Attack Groups	395

Creating Custom IDP Attack Groups	396
Creating Static Attack Groups	396
Creating Dynamic Attack Groups (IDP Only)	397
Updating Dynamic Groups	400
Editing a Custom Attack Group	401
Deleting a Custom Attack Group	401
Configuring Application Identification	401
Updating the NSM App-ID Database	401
Updating the Device App-ID Database	402
Viewing the Device App-ID Database Version	402
Uninstalling an App-ID Database from the Device	403
Unified Threat Management	403
Creating UTM Profiles	403
Creating an Antivirus Profile	404
Creating an Antispam Profile	405
Creating a Content Filtering Profile	406
Creating a URL Filtering Profile	406
Miscellaneous UTM Features	407
Multipurpose Internet Mail Extension (MIME) Lists	408
Extension Lists	408
Command Lists	409
URL Patterns	409
URL Categories	410
ScreenOS Threat Management Features	410
Configuring Antivirus Objects	410
Configuring External AV Profiles	411
Configuring Internal AV Profiles	412
Configuring ICAP AV Servers and Profiles	413
Configuring ICAP AV Profiles	414
Configuring Web Filtering Objects	415
Configuring Custom Policy Fields	416
Defining Metadata	417
Instantiating New Objects	417
Adding Custom Detail Object to Rules	417
Open Log Viewer	418
Configuring GTP Objects	418
Configuring Info	418
Limiting GTP Message Length	419
Limiting GTP Message Rate	419
Limiting GTP Tunnels	419
Removing Inactive GTP Tunnels	419
Validating Sequence Numbers	420
Filtering GTP-in-GTP Packets	420
Removing GTP R6 Informational Elements	420
Inspecting Tunnel Endpoint IDs	420
Configuring Traffic Logging and Counting	420
Traffic Counting	420
Traffic Logging	421

Configuring IMSI Prefix and APN Filtering	421
Creating an APN Filter	421
Creating an IMSI Prefix Filter	422
Configuring GTP Message Filtering	423
Configuring Subscriber Tracing (Lawful Interception)	423
Example: Creating a GTP Object	424
Configuring Service Objects	424
Viewing Predefined Services	425
Creating Custom Services	426
Service Object Groups	427
Example: Creating a Custom Service and Group	428
Example: Creating a Custom Sun-RPC Service	429
Example: Creating a Custom MS-RPC Service	430
Editing and Deleting Service Objects	431
Replacing Service Objects	431
Configuring SCTP Objects	432
Configuring an SCTP Object	432
Configuring Authentication Servers	432
Configuring General Authentication Server Settings	433
Configuring Authentication Server Redundancy	434
Configuring Authentication for User Types	434
Domain Name Checking	434
Domain Name Stripping	434
Configuring Authentication Server Types	435
Configuring a RADIUS Authentication Server	435
Configuring a SecurID Authentication Server	439
Configuring an LDAP Authentication Server	440
Configuring a TACACS Authentication Server	441
Configuring User Objects	441
Configuring Local Users	441
Configuring Local User Groups	442
Configuring External Users	442
Configuring External User Groups	443
Configuring VLAN Objects	445
Configuring IP Pools	445
Using Multiple IP Ranges	446
Configuring Group Expressions	447
Configuring Remote Settings	450
Configuring Routing Instance Objects	450
Viewing Routing Instance Objects	451
Creating Routing Instance Objects	451
Configuring Zone Group Objects	451
Viewing Zone Group Objects	452
Adding a Zone Group Object	452
Configuring NAT Objects	453
Configuring Legacy NAT Objects	453
Configuring DIP Objects	454
Configuring MIP Objects	454
Configuring VIP Objects	454

Configuring Destination NAT Objects	455
Configuring Junos OS NAT Objects	455
Configuring Source NAT Objects	455
Configuring Destination NAT Objects	459
Configuring Certificate Authorities	461
Using Certificate Authorities	462
Configuring Certificate Authorities	462
Configuring CRL Objects	464
Using CRLs	464
Configuring CRLs	464
Configuring Extranet Policies	464
Configuring Binary Data Objects	465
Adding Binary Data Objects	466
Viewing, Editing, and Deleting Binary Data Objects	466
Configuring Protected Resources	466
Creating Protected Resources	467
Editing Protected Resources	468
Configuring IKE Proposals	468
Creating Custom IKE Phase1 Proposals	468
Creating Custom IKE Phase 2 Proposals	469
Configuring Dial-in Objects	470
Creating a Dial-In Profile	471
Linking the Dial-In Profile with the Device	471
Setting the Time-out Period for the Modem Dial-In Authentication	471
Configuring Border Signaling Gateway Objects	471
Chapter 9	
Configuring Security Policies	473
About Security Policies	474
Viewing Rulebase Columns for a Security Policy	474
Viewing and Editing Custom Policy Fields	475
About Rulebases	476
Rule Execution Sequence	477
About Rules	478
About Firewall Rulebases	478
Firewall Rules (Zone and Global)	478
Validating Firewall Policies	479
Global Firewall Policies (Central Policy Mode)	481
Type of Global Firewall Policy	482
Mixed Rule for Zone-based Firewall Policy and Global Firewall Policies	482
Global Address Book Overview	482
Understanding Global Address Books	483
Nested Address Group Support	483
Nested Service Group Support for DMI Devices	483
Services Offload Option on High-End SRX Series Devices in NSM	484
VPN Links and Rules	484
About Rule Groups	484
About the Multicast Rulebase	485
About IDP Rulebases on ISG Family Devices	485

About IDP Rulebases on Standalone IDP Sensors	486
Enabling IPSec Null Encryption for IDP Inspection	487
Managing Security Policies	487
Creating a Security Policy	487
Configuring Objects for Rules	488
Applying the Same Object to Multiple Rules	488
Naming of Address Objects in a Security Policy That References Devices Running ScreenOS or Junos OS	489
Using the Policy Filter Tool	489
Filtering the Comment Field	489
Using a Predefined IDP Policy	490
Using the Policy Creation Wizard	490
Adding Rulebases	491
Configuring Firewall Rules	492
Defining Match for Firewall Rules	492
Configuring Source and Destination Zones for Firewall Rules	492
Configuring Source and Destination Addresses for Firewall Rules	494
Support for Any-IPv6 as a Source Address	495
Configuring Services for Firewall Rules	496
Defining Actions for Firewall Rules	496
Selecting Devices for Firewall Rules	497
Configuring Firewall Rule Options	498
Enabling NAT	498
Enabling GTP for Firewall Rules	499
Configuring Traffic Shaping in a Security Policy	499
Enabling Logging and Counting for Firewall Rules	501
Miscellaneous	502
ID	503
Configuring Web Filtering for Firewall Rules	504
Configuring Authentication for Firewall Rules	505
Configuring Antivirus for Firewall Rules	506
Configuring a DI Profile/Enable IDP for Firewall Rules	507
Limiting Sessions per Policy from Source IPs	508
Configuring the Session Close Notification Rule	509
Comments for Firewall Rules	510
Configuring Multicast Rules	510
Configuring Source and Destination Zones	510
Configuring Source and Destination Groups	510
Configuring Rule Options	511
Configuring Antivirus Rules	511
Configuring Antispam Rules	512
Configuring IDP Rules	512
Defining Match For IDP Rules	513
Configuring Source and Destination Zones for IDP Rules (Does not apply to Standalone IDP Sensor rulebases)	513
Configuring Source and Destination Address Objects for IDP Rules	514
Configuring User Roles for IDP Rules	514
Configuring Services for IDP Rules	515

Configuring Terminal IDP Rules	516
Defining Actions For IDP Rules	517
Configuring Attack Objects in IDP Rules	519
Adding IDP Attack Object Groups by Category	519
Adding IDP Attack Objects by Operating System	520
Adding IDP Attack Objects by Severity	520
Adding Custom Dynamic Attack Groups	521
Configuring IP Actions in IDP Rules	521
Choosing an IP Action	522
Choosing a Block Option	522
Setting Logging Options	522
Setting Timeout Options	523
Configuring Notification in IDP Rules	523
Setting VLAN Tags for IDP Rules	524
Setting Severity for IDP Rules	524
Setting Target Devices for IDP Rules	525
Entering Comments for IDP Rules	525
Configuring multiple IDP policies for an MX Series Router	525
Configuring Application Policy Enforcement (APE) Rules	527
Adding the APE Rulebase Using the Policy Manager	527
Adding the APE Rulebase to a Policy Using the Application Profiler	528
Defining Matches For APE Rules	529
Configuring Applications for APE Rules	529
Configuring Source and Destination Zones for APE Rules (Does not Apply to Standalone IDP Sensor Rulebases)	529
Configuring Source and Destination Address Objects for APE Rules	529
Configuring User Roles for APE Rules	530
Configuring Services for APE Rules	530
Configuring Actions For APE Rules	531
Configuring IP Actions in APE Rules	532
Choosing an IP Action	533
Choosing a Block Option	533
Setting Logging Options	533
Setting Timeout Options	534
Configuring Notification in APE Rules	534
Setting VLAN Tags for APE Rules	535
Setting Severity for APE Rules	535
Setting Target Security Devices for APE Rules	535
Entering Comments for APE Rules	535
Configuring Exempt Rules	535
Adding the Exempt Rulebase	536
Defining a Match	536
Configuring Source and Destination Zones	536
Configuring Source and Destination Address Objects	537
Setting Attack Objects	537
Specifying VLANs	537
Setting Target Devices	537
Entering Comments	537
Creating an Exempt Rule from the Log Viewer	538

Configuring Backdoor Rules	538
Adding the Backdoor Rulebase	539
Defining a Match	539
Configuring Source and Destination Zones	540
Configuring Source and Destination Address Objects	540
Configuring Services	540
Setting Operation	540
Setting Actions	540
Setting Notification	541
Setting Logging	541
Setting an Alert	541
Logging Packets	542
Setting Severity	542
Specifying VLANs	542
Setting Target Devices	542
Entering Comments	542
Configuring SYN Protector Rules	542
The TCP Handshake	543
SYN-Floods	543
Adding the SYN Protector Rulebase	544
Defining a Match	544
Configuring Source and Destination Address Objects	544
Configuring Services	544
Setting Mode	544
Setting Notification	545
Setting Logging	545
Setting an Alert	545
Logging Packets	545
Setting Severity	546
Specifying VLANs	546
Setting Target Devices	546
Entering Comments	546
Configuring Traffic Anomalies Rules	546
Detecting TCP and UDP Port Scans	547
Example: Traffic Anomalies Rule	547
Detecting Other Scans	547
Example: Traffic Anomalies Rule	547
Example: Traffic Anomalies Rule	547
Session Limiting	548
Example: Session Limiting	548
Adding the Traffic Anomalies Rulebase	548
Defining a Match	548
Configuring Source and Destination Address Objects	548
Configuring Services	548
Setting Detect Options	548
Setting Response Options	549
Setting Notification	549
Setting Logging	549
Setting an Alert	549

Logging Packets	549
Setting Severity	550
Specifying VLANs	550
Setting Target Devices	550
Entering Comments	550
Configuring Network Honeypot Rules	550
Impersonating a Port	550
Adding the Network Honeypot Rulebase	551
Defining a Match	551
Configuring the Source	551
Configuring Destination Address Objects and Services	551
Setting Operation	551
Setting Response Options	551
Setting Notification	552
Setting Logging	552
Setting an Alert	552
Logging Packets	552
Setting Severity	552
Specifying VLANs	553
Setting Target Devices	553
Entering Comments	553
Installing Security Policies	553
Assigning a Security Policy to a Device	553
Validating Security Policies	554
Rule Duplication	554
Rule Shadowing	555
Unsupported Options	555
Installing New Security Policies	556
Configuring IDP Policy Push Timeout	557
Updating Existing Security Policies	557
Updating Only the IDP Rulebases on ISG Devices	558
Managing Rules and Policies	558
Helpful Tips	559
Selecting Rules	559
Editing Rule Order	560
Using Cut, Copy, and Paste on Rules	560
Using Cut, Copy, and Paste on Rule Fields	560
Dragging and Dropping Objects	561
Deleting a Rule	562
Disabling a Rule	562
Using Rule Groups	562
Reimporting Devices and Security Policies	562
Merging Policies	563
Importing SRX Series Devices That Contain Inactive Policies	564
Exporting Policies	565
Automatic Policy Versioning	566
Setting NSM to Automatic Policy Versioning	566
Viewing Existing Policy Versions	566
Creating a New Policy Version	567

	Using a Filter to Search for a Policy Version	567
	Editing Comments for an Existing Policy Version	568
	Comparing Two Versions	568
	Restore an Older Version	569
	Viewing, Editing, Filtering, and Sorting Database Versions	569
	Displaying the Differences Between Database Versions	570
	Update Device with an Older Database Version	571
	Pre and Post Rules	572
	Rule Application Sequence	573
	ScreenOS Devices	574
	Validation of prerules and postrules	574
	Install-On Column for prerules and postrules	574
	Managing prerules and postrules	574
	Add prerules and postrules	575
	Push prerules and postrules to Regional Server	575
	Modify prerules and postrules	575
	Delete prerules and postrules	576
	Polymorphic Objects	576
	Customizing Polymorphic Objects	576
	Access Control of Polymorphic Object	577
	Validation of Polymorphic Object	577
	Supported Polymorphic Object Categories	577
	Manage Polymorphic Objects	577
	Create a Polymorphic Object	578
	Add a Polymorphic Object to a Pre/Post Rule	578
	Map a Polymorphic Object to a Real Value	579
	Mapping Polymorphic Objects Before Importing or Updating Affected Devices	579
Chapter 10	Configuring Voice Policies	581
	Adding a BSG Transaction Rulebase	581
	Adding Rules to the BSG Transaction Rulebase	582
Chapter 11	Configuring Junos NAT Policies	585
	Source NAT Policy	586
	Adding a Source NAT Rulebase	586
	Adding a Rule Set to the Source NAT Rulebase	586
	Adding a Rule to a Source NAT Rule Set	587
	Editing a Source NAT Rule or Rule Set	588
	Destination NAT Policy	590
	Adding a Destination NAT Rulebase	590
	Adding a Rule Set to a Destination NAT Rulebase	590
	Adding a Rule to a Destination NAT Rule Set	591
	Editing a Destination NAT Rule or Rule Set	592
	Static NAT Policy	593
	Adding a Static NAT Rulebase	593
	Adding a Rule Set to a Static NAT Rulebase	594
	Adding a Rule to a Static NAT Rule Set	595
	Editing a Static NAT Rule/Rule Set	595

Chapter 12	Configuring VPNs	597
	About VPNs	598
	Creating System-Level VPNs with VPN Manager	598
	Creating Device-Level VPNs in Device Manager	599
	Supported VPN Configurations	599
	Planning for Your VPN	599
	Determining Your VPN Members and Topology	600
	Using Network Address Translation (NAT)	600
	Site-to-Site	600
	Hub and Spoke	601
	Full Mesh	601
	Creating Redundancy	602
	Protecting Data in the VPN	602
	Using IPSec	602
	Using L2TP	604
	Choosing a VPN Tunnel Type	604
	About Policy-Based VPNs	605
	About Route-Based VPNs	605
	VPN Checklist	605
	Define Members and Topology	605
	Define VPN Type: Policy-Based, Route-Based, or Mixed-Mode	606
	Define Security Protocol (Encryption and Authentication)	606
	Define Method: VPN Manager or Device-Level?	606
	Preparing VPN Components	608
	Preparing Basic VPN Components	608
	Preparing Required Policy-Based VPN Components	608
	Configuring Address Objects	609
	Configuring Protected Resources	609
	Configuring Shared NAT Objects	609
	Configuring Remote Access Service (RAS) Users	610
	Configuring Required Routing-Based VPN Components	611
	Configuring Tunnel Interfaces and Tunnel Zones	612
	Configuring Static and Dynamic Routes	612
	Configuring Optional VPN Components	613
	Creating Authentication Servers	613
	Creating Certificate Objects	613
	Creating PKI Defaults	614
	Creating VPNs with VPN Manager	614
	Adding the VPN	615
	Configuring Members	616
	Adding Policy-Based Members	616
	Adding RAS Users	618
	Adding Routing-Based Members	618
	Configuring Topology	620
	Configuring Common VPN Topologies	620
	Defining Termination Points	622
	Configuring Gateways	622
	Configuring Gateway Properties	622
	Configuring Gateway Security	624

Configuring IKE IDs	625
Configuring IKE	626
IKE Properties	626
Configuring Security Level	628
Autogenerating VPN Rules	628
Configuring Overrides	629
Editing Policy Rules	629
Editing Device Configuration	630
Viewing the Device Tunnel Summary	630
Adding the VPN Link	631
Editing VPNs	631
Editing VPN Protected Resources	631
Editing Users	632
Editing the VPN Configuration	632
Editing VPN Overrides	632
VPN Manager Examples	632
Example: Configuring an Autokey IKE, Policy-Based Site-to-Site VPN	632
Example: Configuring an Autokey IKE RAS, Policy-Based VPN	637
Example: Configuring an Autokey IKE, Route-Based Site-to-Site VPN	640
Example: Configuring XAuth Authentication with External User Group	643
Creating Device-Level VPNs	647
Supported Configurations	648
Creating AutoKey IKE VPNs	648
IKEv2 and EAP Support	648
Configuring Gateways	649
Configuring Routes (Route-based only)	653
Configuring the VPN	653
Adding a VPN Rule	656
Creating Manual Key VPNs	656
Adding XAuth Users	656
Configuring Routes (Route-based only)	657
Configuring the VPN	657
Adding a VPN Rule	659
Creating L2TP VPNs	659
Adding L2TP Users	659
Configuring L2TP	660
Adding a VPN Rule	660
Creating L2TP Over Autokey IKE VPNs	660
Adding VPN Rules	661
Configuring the VPN	661
Configuring the Security Policy	662
Assign and Install the Security Policy	662
Device-Level VPN Examples	662
Example: Configuring a Route-Based Site-to-Site VPN, Manual Key	662
Example: Configuring a Policy-Based Site-to-Site VPN, Manual Key	668
Example: Configuring a Policy-Based RAS VPN, L2TP	670
Auto-Connect Virtual Private Network	671
Configuring ACVPN	672
IVE VPN Monitoring	673

Chapter 13	Central Manager	675
	Central Manager Overview	675
	Regional Server and Central Manager Self-Sufficiency	675
	Self-Sufficient Central Manager	675
	Self-Sufficient Regional Server	676
	Super Admin User	676
	Regional Server Management	676
	Management Modes for J Series and SRX Series Devices	676
	Central Management Mode	676
	Device Management Mode	677
	Using Central Manager	677
	Adding a Regional Server Object	677
	Deleting a Regional Server Object	678
	Logging into a Regional Server	678
	Installing Global Policy to a Regional Server	678
	Prerule and Postrule Updates during Global Policy Install	679
	Shared Objects Update During Global Policy Install	679
	Name Space Conflict Resolution for Shared Objects	679
	Name Space Conflict Resolution for Polymorphic Objects	680
Chapter 14	Topology Manager	681
	Overview of the NSM Topology Manager	681
	About the NSM Topology Manager	681
	Requirements for a Topology Discovery	681
	About the NSM Topology Manager Toolbar	682
	Initiating a Topology Discovery	683
	Viewing a Network Topology	684
	About the NSM Topology Map Views	684
	SubNets View	685
	Groups View	685
	Menu Options in the Topology Map View	685
	About the NSM Topology Table Views	686
	Devices View	687
	EndPoint Devices View	687
	Links View	687
	Free Ports View	687
	About Topology Manager Preferences	688
	Default Credentials Tab	688
	Refresh Interval Tab	688
	Preferred Subnets Tab	688
	Adding Discovered Devices to NSM	688
Chapter 15	Role-based Port Templates	691
	Using Role-Based Port Templates	691
	Managing Port Template Associations	692
	Apply or Edit a Port Template	692
	Detect and Resolve Configuration Conflicts	694
	Clone a Port Template	694
	Edit a Port Template	695

Chapter 16	Unified Access Control Manager	697
	Overview of the Unified Access Control (UAC) Manager Views	697
	The Infranet Controller View	697
	The Enforcement Point View	698
	Associating Enforcement Points with an Infranet Controller in the UAC Manager	698
	Disassociating Enforcement Points from an Infranet Controller in the UAC Manager	699
	Resolving Configuration Conflicts with the Infranet Controller in the UAC Manager	699
	Enabling 802.1X on Enforcement Point Ports in the UAC Manager	700
	Disabling 802.1X on Enforcement Point Ports in the UAC Manager	701
	Resolving Configuration Conflicts Between Devices and 802.1X Ports in the UAC Manager	701
 Part 4	 Monitoring	
Chapter 17	Realtime Monitoring	705
	About the Realtime Monitor	705
	Realtime Monitor Views	706
	Monitoring Managed Devices	706
	Viewing Device Status	706
	Device Polling Intervals	709
	Viewing Device Monitor Alarm Status	710
	Setting the Polling Interval For Device Alarm Status	710
	Viewing Additional Device Detail and Statistics	711
	Viewing Device Details	711
	Viewing Device Statistics	712
	Monitoring IDP Sensors	732
	Viewing IDP Device Status	732
	Viewing IDP Device Detail and Statistics	733
	Viewing IDP Device Details	734
	Viewing IDP Process Status	735
	Viewing IDP Device Statistics	735
	Monitoring VPNs	736
	Viewing the VPN Status Summary	736
	Configuring a VPN Filter	737
	Modifying a VPN Filter	738
	Deleting a VPN Filter	738
	Configuring a VPN Display Filter	739
	Viewing Active VPN Details	739
	Viewing Device-Specific VPN Information	739
	Monitoring NSRP Statistics	739
	Viewing NSRP Summary Information	739
	Viewing VSD/RTO Information	740
	Viewing VSD Counter Details	741
	Viewing RTO Counter Details	742

	Monitoring IDP Clusters	742
	Viewing IDP Cluster Summary Information	742
	Monitoring IDP Cluster Members	743
	Using the Realtime Monitor	744
	Monitoring the Management System	744
	Configuring Servers	745
	Configuring Device Servers	745
	Configuring the GUI Server	746
	Using Server Monitor	747
	Viewing Server Status	747
	Viewing Additional Server Status Details	748
	Viewing Process Status	749
	Using Management System Utilities	751
	Using Schema Information	752
	Viewing Device Schema	752
Chapter 18	Analyzing Your Network	753
	About the Dashboard	753
	About the Profiler	753
	Example of Unique Events	754
	Setting Up the Profiler	755
	Configuring the Profiler	756
	Enabling OS Fingerprinting	757
	Configuring Network Objects	757
	Configuring Context Profiles	757
	Configuring Alerts	758
	Updating Profiler Settings	758
	Starting the Profiler	759
	Stopping the Profiler	759
	Starting Profiler Operations on ISG Devices Without IDP Rules	759
	Customizing Profiler Preferences	760
	About Profiler Views	760
	About the Protocol Profiler	761
	About the Network Profiler	762
	About the Violation Viewer	763
	Configuring Permitted Objects	763
	About the Application Profiler	765
	Using Profiler Views	766
	Filtering and Sorting from the Protocol Profiler, Network Profiler, and Violation Viewer	766
	Filtering and Sorting from the Application Profiler	767
	Refreshing Profiler Data	768
	Viewing Database Information	769
	Viewing Detailed Network Information	769
	Purging the Database	770

	Recommended Profiler Options	771
	Configuring a Network Baseline	771
	Identifying a Baseline	771
	Setting a Baseline	771
	Keeping Your Network Current	772
	Proactively Updating Your Network	772
	Reacting to Vulnerability Announcements	772
	Example: Identifying Vulnerable Components	773
	Stopping Worms and Trojans	773
	Example: SQL Worm	773
	Example: Blaster Worm	774
	Accessing Data in the Profiler Database	775
	About Security Explorer	775
	Security Explorer Main Graph	776
	Graph Types	777
	Connections Detail Pane	778
	Reference Point Pane	778
	Log Viewer	778
	Reports Viewer	778
	Using Security Explorer	779
	Permissions	779
	Analyzing Relationships	780
	Viewing Data	780
	Transitioning to Other Relational Graphs	780
	Setting a Time Duration	781
	Viewing Predefined Reports	781
	Refreshing Data	781
	Adding and Removing Panels	781
	Exporting to HTML	781
Chapter 19	Logging	783
	About Logging	783
	About Log Entries	784
	About Log Events	784
	About Log Severity	785
	Viewing Logs	786
	Device Limitations for Viewing Logs	787
	Configuring the Device for Logging	787
	Configuring Severity Settings	788
	Forwarding Self Log Entries (Firewall Options)	789
	Configuring e-mail Server Settings	789
	Configuring Events Reporting Settings	790
	Screen Alarm Log Entries	790
	Event Alarm Log Entries	791
	Traffic Alarm Log Entries	791
	Deep Inspection Alarm Log Entries	792
	Configuration Log Entries	792
	Information Log Entries	793
	Self Log Entries	793

Traffic Log Entries	794
Policy Statistics	794
Attack Statistics	794
Ethernet Statistics	794
Flow Statistics	794
Protocol Distribution	794
Atomic Updating Events	795
Configuring SNMP Reporting Settings	795
Directing Logs to a Syslog Server	796
Directing Logs from DMI Devices	797
Configuring the DMI Device for Stream Mode	797
Directing Data to a WebTrends Server	798
Managing Packet Data in Logs	798
Using the Log Viewer	800
Using Log Views	801
About Predefined Log Views	801
Creating Custom Views and Folders	803
Creating Per-Session Views	804
Log Viewer Columns	804
Log Viewer Detail Panes	807
Log Viewer Status Bar	808
Navigating the Log Viewer	808
Searching Log Entries	808
Log Timeline	809
Using Flags	811
Using the Find Utility	812
Using Log ID Number	812
Filtering Log Entries by Event and Time	812
Setting a Category Filter	813
Setting an Alert Filter	813
Setting a Flag Filter	813
Setting an Address Filter	813
Setting a Protocol Filter	814
Setting a Domain Filter	814
Setting a Time-Based Filter	814
Filtering Log Entries by Range	815
Setting a Bytes In or Bytes Out Range Filter	815
Setting a Port Number Range Filter	816
Customizing Columns	816
Using Column Settings	816
Hide, Unhide, and Move Columns	816
Filtering Log Entries by Column	818
Using Log Viewer Integration	819
Jump to Policy	820
Jump to Device Configuration	820
Identifying Irrelevant Attacks	820

Using the Log Investigator	821
About the Log Investigator UI	822
Configuring Log Investigator Options	824
Configuring a Time Period	824
Configuring Axes	825
Setting a Log Entry Limit	826
Setting Log Investigator Filters	827
Example: Setting Filters in the Log Investigator	829
Investigating Log Entry Data	829
Using Rows and Columns	829
Using Cells	830
Zoom Details	831
Jumping to the Log Viewer	832
Excluding Data	832
Using the Audit Log Viewer	832
Audit Log Table	833
Managing the Audit Log Table	834
Target View and Device View	836
Setting a Start Time for Audit Log Entries	836
Managing Log Volume	837
Automatic Device Log Cleanup	837
Archiving Logs	838
Log Archival Mechanism	838
Setting Log Storage Limits	839
Date Limits	839
System-wide Retention Policy	839
Obsolete Logs	839
Required Disk Space	839
Archive Location	839
Forwarding Logs	840
Sending E-mail Notification of Downed Device	840
Using the Action Manager to Forward Logs by Domain	841
Configuring Action Parameters	841
Setting Device Log Action Criteria	843
Using the log2action Utility to Export Logs	844
Using Filters	844
Exporting to XML	847
Using XML Required and Optional Format-Specific Filters	848
Viewing XML Format Output	848
Exporting to CSV	848
Using CSV Required and Optional Format-Specific Filters	848
Viewing CSV Format Output	849
Exporting to SNMP	849
Using SNMP Required and Optional Format-Specific Filters	850
Viewing SNMP Format Output	850
Exporting to E-mail	850
Using E-mail Required and Optional Format-Specific Filters	851
Exporting to syslog	851
Using Syslog Required and Optional Format-Specific Filters	852

	Viewing Syslog Format Output	852
	Exporting to a Script	852
	Using Script Required and Optional Format-Specific Filters	853
Chapter 20	Reporting	855
	About Reporting	855
	Report Type Groupings	855
	Graphical Data Representation	856
	Integration with Logs	856
	Central Access to Management Information	856
	Report Types	857
	Predefined Reports	857
	Firewall/VPN Reports	857
	DI/IDP Reports	858
	Screen Reports	859
	Administrative Reports	860
	UAC Reports	860
	Profiler Reports	861
	AVT Reports	861
	SSL/VPN Reports	862
	EX Series Switches Report	862
	My Reports	862
	Shared Reports	862
	Working with Reports	862
	Generating a Predefined Report	863
	Creating a Custom Report	863
	Example: Creating a Custom Report	863
	Deleting Reports	864
	Organizing Reports in Folders	864
	Generating Reports Automatically	864
	Running Reports Using the guiSvrCLI.sh Utility	865
	Creating and Editing Action Scripts	866
	Using Cron with Scheduled Reports	867
	Running Xdb Audit Log Exporter Tools with High Availability	868
	Exporting Reports to HTML	869
	Setting Report Options	870
	Naming a Report	870
	Setting the Report Type	870
	Configuring Report Source Data	871
	Configuring a Report Time Period	871
	Configuring the Data Point Count	871
	Configuring the Chart Type	871
	Sharing Your Custom Report	871
	Modifying Report Filters	871
	Configuring Report Processing Warnings	872
	Saving Your Report Settings	872
	Log Viewer Integration	872
	Viewing Logs From Report Manager	872
	Generating Quick Reports	873

	Using Reports	873
	Example: Using Administrative Reports to Track Incidents	873
	Example: Using Administrative Reports to Optimize Rulebases	874
	Example: Using EX Switch Reports to Track Configuration Changes	875
	Example: Using SSL/VPN Reports to Track Authentication Failures	876
	Example: Using Screen Reports to Identify Attack Trends	876
	Example: Using DI Reports to Detect Application Attacks	876
	Using the Watch List	877
Part 5	Appendixes	
Appendix A	Glossary	881
	Network and Security Manager (NSM) Term Definitions	881
Appendix B	Unmanaged ScreenOS Commands	907
Appendix C	SurfControl Web Categories	909
Appendix D	Common Criteria EAL2 Compliance	917
	Guidance for Intended Usage	917
	Guidance for Personnel	917
	Guidance for Physical Protection	917
Appendix E	Log Entries	919
	Screen Alarm Log Entries	919
	Alarm Log Entries	921
	Deep Inspection Alarm Log Entries	922
	Configuration Log Entries	997
	Information Log Entries	999
	Self Log Entries	1001
	Traffic Log Entries	1001
	GTP Log Entries	1002

List of Figures

Part 1	Getting Started with NSM	
Chapter 1	Introduction to Network and Security Manager	3
	Figure 1: NSM Network Architecture	11
	Figure 2: NSM System Architecture	12
	Figure 3: GUI Preference Options - Disabling GUI Validation	25
	Figure 4: Disable GUI Validation Options	25
	Figure 5: GUI Preference Options - Custom Validations	26
	Figure 6: Overview of the User Interface	27
	Figure 7: UI Search Modes	37
	Figure 8: "Contains String" Search Mode Example	37
	Figure 9: "Starts With" Search Mode Example	38
	Figure 10: "Regular Expression" Search Mode Details	38
	Figure 11: "Regular Expression" Search Mode Example	39
	Figure 12: "IP Address" Search Mode Example	40
	Figure 13: Exact String Search Mode Example	41
Chapter 2	Planning Your Virtual Network	43
	Figure 14: Selecting the GUI Server in Central Manager	60
	Figure 15: Setting Up an Information Banner	61
	Figure 16: Information Banner Login into Central Manager	61
Chapter 3	Configuring Role-Based Administration	63
	Figure 17: Creating Custom Domain	72
	Figure 18: User in Domain "global" with a Predefined Role	73
	Figure 19: User in Domain "global" with Custom Role "r1"	74
	Figure 20: User in Subdomain "d1" With a Predefined Role	74
	Figure 21: User in Subdomain "d1" With a Custom Role "r1"	74
	Figure 22: Assigning Multiple Roles to a User in Global Domain	75
	Figure 23: Assigning Multiple Roles to a User in Subdomain	75
	Figure 24: Assigning Roles Defined in Domain "global"	76
	Figure 25: Assigning Roles Defined in Domain "global" to Subdomain Only	76
	Figure 26: Manage Administrators and Domains: Administrators Tab	95
Part 2	Integrating	
Chapter 4	Adding Devices	99
	Figure 27: Connecting Devices from Different Domains in VPNs	105
	Figure 28: Trust-Untrust Port Mode Bindings	107
	Figure 29: Home-Work Port Mode Bindings	108
	Figure 30: Dual-Untrust Port Mode Bindings	108
	Figure 31: Combined Port Mode Bindings	109

	Figure 32: Trust-Untrust-DMZ Port Mode Bindings	110
	Figure 33: Extended Port-Mode Interface to Zone Bindings	110
	Figure 34: DMZ Dual Untrust Port Mode	111
	Figure 35: Connecting Vsys Devices Across Domains	152
	Figure 36: Adding a Secure Access Cluster	164
	Figure 37: Adding a J Series Cluster	166
	Figure 38: Adding the First Member to a J Series Cluster	167
	Figure 39: Adding the Second Member to a J Series Cluster	167
	Figure 40: Cluster Member Icons	168
	Figure 41: Configuring Cluster Members for Paris Vsys Cluster	171
	Figure 42: Paris Cluster Members and Paris Vsys Cluster Members	172
Chapter 5	Configuring Devices	199
	Figure 43: Info and Configuration Tabs	203
	Figure 44: ScreenOS and IDP Device Configuration Information	204
	Figure 45: ScreenOS Device Object Configuration Data	206
	Figure 46: Secure Access Device Object	208
	Figure 47: Applying a Template	213
	Figure 48: Template Override Icon	214
	Figure 49: Revert to a Template or Default Value	215
	Figure 50: View Denial of Service Defense Values from DoS Template	217
	Figure 51: Configure DoS Defense Settings for the DoS2 Template	218
	Figure 52: View Template Priority (DoS Highest)	219
	Figure 53: View Values from DoS and DoS2 Templates	220
	Figure 54: View DoS2 Value for Source IP Based Session Limit	220
	Figure 55: View DoS Value for SYN-ACK-ACK Proxy Protection Setting	220
	Figure 56: View Default SYN-ACK-ACK Proxy Protection Setting	221
	Figure 57: Up and Down Arrows for Changing the Sequence of a List	223
	Figure 58: Identifying Ordered List Entries that Do Not Match the Template Order	229
	Figure 59: Template Operations Directive	230
	Figure 60: Select Template Dialog Box	231
	Figure 61: Template Operations Job Information Dialog Box	233
	Figure 62: Adding a Configuration Group	238
	Figure 63: Applying a Configuration Group	240
	Figure 64: Configuration Group Applied	240
	Figure 65: Excluding a Configuration Group	241
	Figure 66: Configuring Routing Engine Specific Parameters	251
	Figure 67: Viewing the Routing Engine Configuration	252
Chapter 6	Updating Devices	257
	Figure 68: Delta Configuration Summary Example	269
	Figure 69: Job Manager Module	274
	Figure 70: Job Information Dialog Box	275
	Figure 71: Failed Update Job Dialog Box	277
Chapter 7	Managing Devices	279
	Figure 72: Viewing the Device Inventory	286
	Figure 73: Comparing the Device Inventory with the NSM Database	288
	Figure 74: Adding a Shared Binary Data Object	292

	Figure 75: Linking to a Shared Binary Data Object	293
	Figure 76: Attack Update Summary	313
	Figure 77: Import/Update Architecture	332
	Figure 78: Data Model Update	334
	Figure 79: Data Model Importing	336
	Figure 80: GUI Server Configuration File	341
	Figure 81: Dev Server Configuration file	341
Part 3	Managing	
Chapter 8	Configuring Objects	345
	Figure 82: Shared Address Object Validation Option	360
	Figure 83: Shared Address Object Validation Warning	361
	Figure 84: New Dynamic Group	398
	Figure 85: New Dynamic Group Filters	399
	Figure 86: New Dynamic Group Members	400
	Figure 87: Configure External User Groups for Sales and Marketing	449
	Figure 88: Configure Group Expression for Sales and Marketing	449
Chapter 9	Configuring Security Policies	473
	Figure 89: Displaying the Select Visible Columns Dialog Box	475
	Figure 90: Firewall Validation Option	480
	Figure 91: Firewall Error Validation Option	480
	Figure 92: Firewall Warnings Validation Option	481
	Figure 93: Configure IP Action	522
	Figure 94: Security Policy A Rules (Before Policy Merge)	564
	Figure 95: Security Policy B Rules (Before Policy Merge)	564
	Figure 96: Security Policy Rules (Merged from Policy A and Policy B)	564
Chapter 12	Configuring VPNs	597
	Figure 97: Create Tokyo Protected Resource Object for AutoKey IKE VPN	634
	Figure 98: Create Paris Protected Resource Object for AutoKey IKE VPN	634
	Figure 99: Configure Gateway Parameters for AutoKey IKE VPN	636
	Figure 100: View Autogenerated Rules for AutoKey IKE VPN	636
	Figure 101: Add Chicago Protected Resource for AutoKey IKE RAS VPN	638
	Figure 102: Add New Local User for AutoKey IKE RAS VPN	638
	Figure 103: Configure Security for AutoKey IKE RAS VPN	640
	Figure 104: View Tunnel Summary for AutoKey IKE, RB Site-to-Site VPN	642
	Figure 105: Configure Tokyo Route for RB Site-to-Site VPN, MK	665
	Figure 106: Configure Tokyo Trust Route for RB Site-to-Site VPN, MK	665
	Figure 107: View Tokyo Routing Table for RB Site-to-Site VPN, MK	666
	Figure 108: Configure Rules for RB Site-to-Site VPN, MK	668
Part 4	Monitoring	
Chapter 17	Realtime Monitoring	705
	Figure 109: Server Monitor (Machine-wide Info)	747
	Figure 110: Process Status for the Device Server	750
	Figure 111: Process Status for the GUI Server	750
Chapter 18	Analyzing Your Network	753

	Figure 112: Security Explorer	776
Chapter 19	Logging	783
	Figure 113: View Packet Data in a Log	799
	Figure 114: Sample Packet Data	800
	Figure 115: View Category and Severity Filters	808
	Figure 116: Log Viewer Time Slider	810
	Figure 117: Log Viewer Time Display	810
	Figure 118: Filter Summary Dialog Box	819
	Figure 119: Viewing Summary Panel	821
	Figure 120: Log Investigator UI Overview	823
	Figure 121: Configure Time Period Filter	825
	Figure 122: Changing Time Period Filter	825
	Figure 123: View Log Investigator Results	829
	Figure 124: Audit Log Viewer UI Overview	833
Chapter 20	Reporting	855
	Figure 125: Generating A Quick Report	873
	Figure 126: Logs by User-Set Flag Report	874
	Figure 127: Top FW/VPN Rules Report	875
	Figure 128: Top Configuration Changes Report	876

List of Tables

	About This Guide	xli
	Table 1: Notice Icons	xlii
	Table 2: Text Conventions	xlii
	Table 3: Syntax Conventions	xlili
	Table 4: Network and Security Manager Publications	xliv
Part 1	Getting Started with NSM	
Chapter 1	Introduction to Network and Security Manager	3
	Table 5: GUI Server Processes	12
	Table 6: Device Server Processes	13
	Table 7: Supported Security Devices	14
	Table 8: J Series Services Routers and SRX Series Services Gateways NSM Supports	17
	Table 9: M Series Multiservice Edge Routers and MX Series Ethernet Services Routers NSM Supports	19
	Table 10: EX Series Ethernet Switches NSM Supports	20
	Table 11: Secure Access Products NSM Supports	21
	Table 12: IC Series UAC Appliances NSM Supports	21
	Table 13: Validation Status for Devices	34
	Table 14: Validation Icons	35
Chapter 3	Configuring Role-Based Administration	63
	Table 15: How to Authenticate Users	71
	Table 16: Predefined NSM Administrator Activities	77
	Table 17: Changes to Edit Devices, Device Groups, & Templates Activity	90
	Table 18: Changes to View Devices, Device Groups, & Templates Role	91
Part 2	Integrating	
Chapter 4	Adding Devices	99
	Table 19: Extended Bindings	111
	Table 20: Security Device Port Mode Summary (Part 1)	112
	Table 21: Security Device Port Mode Summary (Part 2)	112
	Table 22: Supported Add Device Workflows by Device Family	113
	Table 23: CSV File Information for Devices with Static IP Addresses	183
	Table 24: CSV File Information for Devices with Dynamic IP Addresses	184
	Table 25: CSV File Information for Undeployed Devices	187
Chapter 5	Configuring Devices	199
	Table 26: Validation Icons	205

Chapter 6	Updating Devices	257
	Table 27: Additional Configuration States for Devices Running ScreenOS 5.1 and Later	264
	Table 28: Delta Configuration Summary Information	268
	Table 29: Device States During Update	276
Chapter 7	Managing Devices	279
	Table 30: Adjust OS Version Directive Actions for Major and Service Releases	284
	Table 31: Scheduled Security Update (SSU) Command Line Parameters	314
Part 3	Managing	
Chapter 8	Configuring Objects	345
	Table 32: Predefined Application Table Tab Information	357
	Table 33: Predefined Extended Application Table Tab Information	357
	Table 34: Deep Inspection Profile Actions	367
	Table 35: Deep Inspection IP Actions	368
	Table 36: IP Protocol Name and Type Numbers	377
	Table 37: Supported Services for Service Bindings	378
	Table 38: Attack Pattern Syntax	383
	Table 39: Attack Pattern Syntax Example Matches	384
	Table 40: DI Attack Header Match Modifiers	387
	Table 41: Service Table Tab Information	425
	Table 42: Group Expression Operators	447
	Table 43: Source NAT Configuration Options	456
	Table 44: Destination NAT Configuration Options	460
Chapter 9	Configuring Security Policies	473
	Table 45: IDP Rule Actions	518
	Table 46: Severity Levels, Recommended Actions and Notifications	520
	Table 47: APE Rule Actions	531
	Table 48: Actions for Backdoor Rule:	541
	Table 49: Rule Shadowing Example	555
	Table 50: Polymorphic Objects	577
Part 4	Monitoring	
Chapter 17	Realtime Monitoring	705
	Table 51: Device Status Information	707
	Table 52: Device Polling Intervals	710
	Table 53: Device Detail Status Items	711
	Table 54: Device Statistics Summary	712
	Table 55: Device-Specific Views	714
	Table 56: Policy Distribution Items	715
	Table 57: Protocol Distribution Items	716
	Table 58: VPN Monitor Table	718
	Table 59: Active VPN Table	720
	Table 60: Ethernet Statistics View Data	722
	Table 61: Flow Statistics View Data	723

	Table 62: Attack Counters	724
	Table 63: Resource Statistics Items	727
	Table 64: Administrators View	728
	Table 65: Authenticated Users View	728
	Table 66: Active Sessions Items	729
	Table 67: HA Statistics View	731
	Table 68: Device Status Information	732
	Table 69: IDP Device Detail Status Items	734
	Table 70: IDP Sensor Process Status Items	735
	Table 71: Device Statistics Summary (for IDP Sensors)	736
	Table 72: VPN Tunnel Summary	736
	Table 73: NSRP Device Summary	739
	Table 74: VSD/RTO Summary	740
	Table 75: VSD Counter Details	741
	Table 76: RTO Counters Details	742
	Table 77: IDP Cluster Monitor	742
	Table 78: IDP Cluster Summary	743
	Table 79: IDP Cluster Member Monitor	744
	Table 80: Server Information	745
	Table 81: GUI Server Table	746
	Table 82: Server Monitor (Machine-wide Info) Data	748
	Table 83: Server Detail Status	749
	Table 84: Process Status	750
	Table 85: Management System Utilities	751
Chapter 18	Analyzing Your Network	753
	Table 86: General IDP Profiler Settings	756
	Table 87: Protocol Profiler Data	761
	Table 88: Network Profiler Data	762
	Table 89: Application Profiler Data	765
	Table 90: Detailed Network Information Data	769
	Table 91: Transitional Graphs	780
Chapter 19	Logging	783
	Table 92: Event-Generated Log Entries	784
	Table 93: Log Entry Severity Levels for DMI Devices	785
	Table 94: Log Entry Severity Levels for ScreenOS and IDP Devices	785
	Table 95: Destinations of Log Entry Severities	788
	Table 96: Self Log Entry Settings	789
	Table 97: Email Server Settings for Log Entries	789
	Table 98: Syslog Settings for Log Entries	797
	Table 99: WebTrends Settings for Log Entries	798
	Table 100: EX Series Switch Predefined Log Views	801
	Table 101: SSL/UAC Predefined Log Views	802
	Table 102: Predefined Log Views	802
	Table 103: Log Viewer Columns	804
	Table 104: Log Viewer Navigation Controls	808
	Table 105: Search Tools for Log Viewer	809
	Table 106: Log Viewer Flags	811
	Table 107: Irrelevant Versus Relevant Attacks	821

	Table 108: Log Investigator Filters	827
	Table 109: Log Investigator Analysis	830
	Table 110: Audit Log Information	833
	Table 111: Common Filters	845
Chapter 20	Reporting	855
	Table 112: Firewall and VPN Reports	857
	Table 113: DI/IDP Reports	858
	Table 114: Screen Reports	859
	Table 115: Administrative Reports	860
	Table 116: UAC Reports	860
	Table 117: Profiler Reports	861
	Table 118: AVT Reports	861
	Table 119: SSL/VPN Reports	862
	Table 120: EX-Switch Reports	862
Part 5	Appendixes	
Appendix A	Glossary	881
	Table 121: CIDR Translation	885
Appendix B	Unmanaged ScreenOS Commands	907
	Table 122: Unmanaged Commands for Firewall/VPN Devices	907
Appendix C	SurfControl Web Categories	909
	Table 123: SurfControl Web Categories	909
Appendix E	Log Entries	919
	Table 124: Screen Alarm Log Entries	919
	Table 125: Alarm Log Entries	921
	Table 126: Deep Inspection Alarm Log Entries	923
	Table 127: Configuration Log Entries	997
	Table 128: Information Log Entries	1000

About This Guide

- Objectives on page xli
- Audience on page xli
- Conventions on page xlii
- Documentation on page xliii
- Requesting Technical Support on page xlv

Objectives

Juniper Networks Network and Security Manager (NSM) is a software application that centralizes control and management of your Juniper Networks devices. With NSM, Juniper Networks delivers integrated, policy-based security and network management for all devices.

NSM uses the technology developed for Juniper Networks ScreenOS to enable and simplify management support for previous and current versions of ScreenOS and current versions of Junos OS. By integrating management of all Juniper Networks security devices, NSM enhances the overall security of the Internet gateway.

This guide describes NSM features and provides a technical overview of the management system architecture. It also explains how to configure basic and advanced NSM functionality, including adding new devices, deploying new device configurations, updating device firmware, managing security policies and VPNs, viewing log information, and monitoring the status of your network. Use this guide in conjunction with the NSM Online Help, which provides step-by-step instructions for many of the processes described in this document.



NOTE: If the information in the latest NSM Release Notes differs from the information in this guide, follow the NSM Release Notes.

Audience

This guide is intended for system administrators responsible for the security infrastructure of their organization. Specifically, this book discusses concepts of interest to firewall and VPN administrators, network/security operations center administrators; and system administrators responsible for user permissions on the network.

Conventions

The sample screens used throughout this guide are representations of the screens that appear when you install and configure the NSM software. The actual screens may differ.

All examples show default file paths. If you do not accept the installation defaults, your paths will vary from the examples.

Table 1 on page xlii defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xlii defines text conventions used in this guide.

Table 2: Text Conventions

Convention	Description	Examples
Bold typeface like this	<ul style="list-style-type: none"> Represents commands and keywords in text. Represents keywords Represents UI elements 	<ul style="list-style-type: none"> Issue the clock source command. Specify the keyword exp-msg. Click User Objects
Bold typeface like this	Represents text that the user must type.	user input

Table 2: Text Conventions (continued)

Convention	Description	Examples
fixed-width font	Represents information as displayed on the terminal screen.	host1# show ip ospf Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an area Border Router (ABR)
Key names linked with a plus (+) sign	Indicates that you must press two or more keys simultaneously.	Ctrl + d
<i>Italics</i>	<ul style="list-style-type: none"> Emphasizes words Identifies variables 	<ul style="list-style-type: none"> The product supports two levels of access, <i>user</i> and <i>privileged</i>. <i>clusterID</i>, <i>ipAddress</i>.
The angle bracket (>)	Indicates navigation paths through the UI by clicking menu options and links.	Object Manager > User Objects > Local Objects

Table 3 on page [xlili](#) defines syntax conventions used in this guide.

Table 3: Syntax Conventions

Convention	Description	Examples
Words in plain text	Represent keywords	terminal length
Words in italics	Represent variables	<i>mask</i> , <i>accessListName</i>
Words separated by the pipe () symbol	Represent a choice to select one keyword or variable to the left or right of this symbol. The keyword or variable can be optional or required.	diagnostic line
Words enclosed in brackets ([])	Represent optional keywords or variables.	[internal external]
Words enclosed in brackets followed by an asterisk ([]*)	Represent optional keywords or variables that can be entered more than once.	[level 1 level 2 11]*
Words enclosed in braces ({ })	Represent required keywords or variables.	{ permit deny } { in out } { clusterId ipAddress }

Documentation

Table 4 on page [xliv](#) describes documentation for the NSM.

Table 4: Network and Security Manager Publications

Book	Description
<i>Network and Security Manager Installation Guide</i>	Describes the steps to install the NSM management system on a single server or on separate servers. It also includes information on how to install and run the NSM user interface. This guide is intended for IT administrators responsible for the installation or upgrade of NSM.
<i>Network and Security Manager Administration Guide</i>	<p>Describes how to use and configure key management features in the NSM. It provides conceptual information, suggested workflows, and examples. This guide is best used in conjunction with the NSM Online Help, which provides step-by-step instructions for performing management tasks in the NSM UI.</p> <p>This guide is intended for application administrators or those individuals responsible for owning the server and security infrastructure and configuring the product for multi-user systems. It is also intended for device configuration administrators, firewall and VPN administrators, and network security operation center administrators.</p>
<i>Network and Security Manager Configuring ScreenOS Devices Guide</i>	Provides details about configuring the device features for all supported ScreenOS platforms.
<i>Network and Security Manager Configuring Intrusion Detection Prevention Devices Guide</i>	Provides details about configuring the device features for all supported Intrusion Detection Prevention (IDP) platforms.
<i>Network and Security Manager Online Help</i>	Provides procedures for basic tasks in the NSM user interface. It also includes a brief overview of the NSM system and a description of the GUI elements.
<i>Network and Security Manager API Guide</i>	Provides complete syntax and description of the SOAP messaging interface to NSM.
<i>Network and Security Manager Release Notes</i>	<p>Provides the latest information about features, changes, known problems, resolved problems, and system maximum values. If the information in the Release Notes differs from the information found in the documentation set, follow the Release Notes.</p> <p>Release notes are included on the corresponding software CD and are available on the Juniper Networks website.</p>
<i>Network and Security Manager Configuring Infranet Controllers Guide</i>	Provides details about configuring the device features for all supported Infranet Controllers.
<i>Network and Security Manager Configuring Secure Access Devices Guide</i>	Provides details about configuring the device features for all supported Secure Access Devices.

Table 4: Network and Security Manager Publications (continued)

Book	Description
<i>Network and Security Manager Configuring EX Series Switches Guide</i>	Provides details about configuring the device features for all supported EX Series platforms.
<i>Network and Security Manager Configuring J Series Services Routers and SRX Series Services Gateways Guide</i>	Provides details about configuring the device features for all supported J Series Services Routers and SRX Series Services Gateways.
<i>Network and Security Manager M Series and MX Series Devices Guide</i>	Provides details about configuring the device features for M Series and MX Series platforms.

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>

- Join and participate in the Juniper Networks Community Forum:
<https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

PART 1

Getting Started with NSM

The chapters in Part 1 of the Release 2010.3 version of the *Network and Security Manager Administration Guide* provide an overview of the management system and describe how to prepare to integrate your existing network security structure using NSM role-based administration tools.

Part 1 contains the following chapters:

- [Introduction to Network and Security Manager on page 3](#)
- [Planning Your Virtual Network on page 43](#)
- [Configuring Role-Based Administration on page 63](#)

CHAPTER 1

Introduction to Network and Security Manager

Juniper Networks Network and Security Manager (NSM) gives you complete control over your network. Using NSM, you can configure all your Juniper Networks devices from one location, at one time.

This chapter contains the following sections:

- [About NSM on page 3](#)
- [Technical Overview on page 10](#)
- [Working in the User Interface on page 23](#)
- [New Features in 2012.2 on page 42](#)

About NSM

A management system integrates your individual devices into a single security system that you control from a central location. With NSM, you can manage your network at the system level, using policy-based central management, as well as at the device level, managing all device parameters for devices.

NSM works with networks of all sizes and complexity. You can add a single device, or create device templates to help you deploy multiple devices. You can create new policies, or edit existing policies for security devices. The management system tracks and logs each administrative change in real time, providing you with a complete administrative record and helping you perform fault management.

NSM also simplifies control of your network with a straightforward user interface (UI). Making all changes to your devices from a single, easy-to-use interface can reduce deployment costs, simplify network complexity, speed configuration, and minimize troubleshooting time.

The following sections provide an overview of the key management features of NSM:

- [Security Integration on page 4](#)
- [Device Management on page 6](#)

- [Error Prevention, Recovery, and Auditing on page 7](#)
- [Complete System Management on page 8](#)

Security Integration

True security integration occurs when you can control every device on your network and see every security event in real time from one location. In NSM, this location is the NSM UI, a graphical user interface that contains a virtual representation of every device on your network. You use this console to view your network, the devices running on it, the policies controlling access to it, and the traffic that is flowing through it.

Device Configuration

You can create and manage device configurations for devices or systems in your network. NSM provides support for device configuration commands, so you can retain complete control over your devices when using system-level management features like VPNs.

Network Organization

Use domains to segment your network functionally or geographically to define specific network areas that multiple administrators can manage easily.

A domain logically groups devices, their policies, and their access privileges. Use a single domain for small networks with a few security administrators, or use multiple domains for enterprise networks to separate large, geographically distant or functionally distinct systems, or to control administrative access to individual systems.

With multiple domains, you can create objects, policies, and templates in the global domain, and then create subdomains that automatically inherit these definitions from the global domain.

Role-Based Administration

Control access to management with NSM. Define strategic roles for your administrators, delegate management tasks, and enhance existing permission structures by enabling permissions for particular tasks.

Use NSM to create a security environment that reflects your current offline administrator roles and responsibilities. You can configure multiple administrators for multiple domains. By specifying the exact tasks your NSM administrators can perform within a domain, you minimize the chance of errors and security violations, and enable a clear audit trail for every management event.

Initially, when you log in to NSM as the super administrator, you have full access to all functionality within the global domain. From the global domain, you can add NSM administrators, configure their roles, and specify the subdomains to which they have access:

- **Activities and Roles**—An activity is a predefined task performed in the NSM system. A role is a collection of activities that defines an administrative function. Use activities to create custom roles for your NSM administrators.

- **Administrators**—An administrator is a user of NSM. Each administrator has a specific level of permissions. Create multiple administrators with specific roles to control access to the devices in each domain.
- **Default Roles**—Use the predefined roles System Administrator, Read-Only System Administrator, Domain Administrator, Read-Only Domain Administrator, IDP Administrator, or Read-Only IDP Administrator to create permissions for your administrators quickly.



NOTE: In a mixed environment, an administrator with the IDP Administrator role is unable to take full command of all managed devices because of the predefined restrictions. If IDP Administrators are expected to manage other devices in a mixed environment, they need to know the restrictions and have their roles modified to include the necessary permissions.

Centralized Device Configuration

No matter how large your network, you can use several system management mechanisms to help you create or modify multiple device configurations quickly and efficiently at one time:

- **Templates**—A template is a predefined device configuration that helps you reuse specific information. Create a device template that defines specific configuration values, and then apply that template to devices to configure multiple devices at one time. For more flexibility, you can combine and apply multiple device templates to a single device configuration.
- **Configuration groups**—In Junos devices, configuration groups allow you to create a group containing configuration statements and to direct the inheritance of that group's statements in the rest of the configuration. The same group can be applied to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.
- **Shared objects**—An object is an NSM definition that is valid in the global domain and all subdomains. Any object created in the global domain is a shared object that is shared by all subdomains; the subdomain automatically inherits any shared objects defined in the global domain. You will not see global objects in the Object Manager of a subdomain; however, you can use the objects when selecting objects in a policy.

The global domain is a good location for security devices and systems that are used throughout your organization, address book entries for commonly used network components, or other frequently used objects. A subdomain, alternatively, enables you to separate firewalls, systems, and address objects from the global domain and other subdomains, creating a private area to which you can restrict access.

- **Grouping**—A group is a collection of similar devices or objects. Use device groups and object groups to update multiple devices simultaneously, simplify rule creation and deployment, and enable group-specific reporting. You can even link groups using Group Expressions to create a custom group.

Device Management

As your network grows, you might need to add existing devices, add new devices, reconfigure existing devices, update software versions on older devices, or integrate a new network to work with your existing network. NSM provides a virtual environment in which to first model, verify, and then update your managed devices with changes.

Importing Devices

If you have existing devices deployed you can use the NSM import feature to import their configurations, address books, service objects, policies, VPNs, and administrator privileges. As NSM imports your existing device configurations, it automatically creates your virtual network based on the configuration information.

You can import device configurations directly from your devices. Import all your devices at one time, or, if your network is large, import one domain at a time.

Device Modeling

Using your virtual network to change, review, and test your network configuration before deploying it to your physical network can help you discover problems like routing issues, IP conflicts, and version mismatches across your entire network before they actually occur. NSM includes configuration validation to help you identify device configuration errors and missing information, and then points you to the trouble spot so you can quickly fix the problem. When you have designed a virtual configuration that works, you can push this configuration to your devices with a single update.

You can implement a new routing protocol across your network, design and deploy a new security policy with traffic shaping, or create a new VPN tunnel that connects a branch office to your corporate network.

Rapid Deployment

Rapid Deployment (RD) enables deployment of multiple ScreenOS security devices in a large network environment with minimal user involvement. RD simplifies the staging and configuration of security devices in nontechnical environments, enabling the secure and efficient deployment of a large number of devices.

To use RD, the NSM administrator creates a small file (called a configlet) in NSM, and then sends that configlet to an on site administrator who has local access to the security device. With the help of the Rapid Deployment wizard, the onsite administrator installs the configlet on the device, which automatically contacts NSM and establishes a secure connection for device management.

RD is ideal for quickly bringing new security devices under NSM management for initial configuration. You can model and verify your device configurations for undeployed devices, and then install the completed device configuration when the device contacts NSM.

Policy-Based Management

Create simplified and efficient security policies for your managed devices. You can manage security policies either in a Central Policy Manager or through in-device policy

management, depending on the type of device. The tools at your disposal are also device-dependent, but can include:

- **Groups**—Group your devices by platform, OS version, location, or function, and then add them to your security policies.
- **Zone Exceptions**—To simplify your rules, define a common To Zone and From Zone for all devices in the rule, and then specify zone exceptions to change the To and From zones for specific devices. Zone exceptions add flexibility to your rules, enabling you to manage more devices in a single rule.
- **Filtering**—Filter on From and To Zones to see rules between zones.
- **Scheduling**—Schedule a period during which a security policy is in effect on the devices in a rule. Create schedule objects as one-time, recurring, or both.
- **Security and Protection**—Configure a rule to look for attacks, viruses, or specific URLs.
- **Traffic Shaping**—Use your firewall rules to control the amount of traffic permitted through your managed devices.

Error Prevention, Recovery, and Auditing

Using NSM's error prevention and recovery features, you can ensure that you are consistently sending stable configurations to your devices, and that your device remains connected to NSM. You can track each change made by a NSM administrator to help you identify when, how, and what changes were made to your managed devices.

Device Configuration Validation

NSM alerts you to configuration errors while you work in the UI. Each field that has incorrect or incomplete data displays an error icon:



Move your cursor over the icon to see details. For more details on validation, see [“Validation Icons in the User Interface” on page 34](#).

Policy Validation

The policy validation tool checks your security policies and alerts you to possible problems before you install them on your managed devices.

Atomic Configuration and Updating

If the configuration deployment fails for any reason, the device automatically uses the last installed stable configuration. If the configuration deployment succeeds, but the device loses its connection to the management system, the device restores the last installed configuration. This feature minimizes downtime and ensures that NSM always maintains a stable connection to the managed device.

Your security devices can be updated atomically, which enables the device to receive the entire modeled configuration (all commands) before executing those commands, instead of executing commands as they are received from the management system.

Because the device no longer needs to maintain a constant connection to the management system during updating, you can configure changes to management connection from the NSM UI.

Device Image Updates

You can update the software that runs on your devices by installing a new image on your managed devices:

- NSM updates—For ScreenOS and Junos families of devices, you can use NSM to upload the new image file to multiple devices with a single click.
- RMA updates—To replace failed devices, set the device to the RMA state, which enables NSM to retain the device configuration without a serial number or connection statistics. When you install the replacement device, activate the device with the serial number of the replacement unit.

Auditing

Use the Audit Log Viewer to track administrative actions so you'll always know exactly when and what changes were made using the management system. The Audit Log Viewer displays log entries in the order generated, and includes:

- Date and time the administrative action occurred
- NSM administrator who performed the action
- Action performed
- Domain (global or a subdomain) in which the action occurred
- Object type and name

Complete System Management

NSM provides the tools and features you need to manage your devices as a complete system, as well as individual networks and devices:

- To manage an individual device, create a single device configuration, define a security policy for that device, and monitor the device status
- To manage a network, create multiple device configurations, define and install policies for multiple devices, and view the status of all devices in the same UI.
- To manage at the system level, create templates and use them to quickly configure multiple policies and VPNs that control the flow of traffic through your network, view system-wide log information for network security events, and monitor the status of NSRP.

VPN Abstraction

Use VPN Manager to design a system-level VPN and automatically set up all connections, tunnels, and rules for all devices in the VPN. Instead of configuring each device as a VPN member and then creating the VPN, start from a system perspective: Determine which

users and networks need access to each other, and then add those components to the VPN.

Using AutoKey IKE, you can create the following VPNs with VPN Manager:

- Dynamic, route-based VPNs—Provide resilient, always-on access across your network. Add firewall rules on top of route-based VPNs to control traffic flow.
- Policy-based VPNs—Connect devices, remote access service (RAS) users, and control traffic flow (you can also create policy-based VPNs with L2TP).
- Mixed-mode VPNs—Connect route-based VPNs with policy-based VPNs, giving you flexibility.

Integrated Logging and Reporting

You can use NSM to monitor, log, and report on network activity in *real time* to help you understand what is happening on your network:

- View traffic log entries generated by network traffic events, configuration log entries generated by administrative changes, or create custom views to see specific information in the Log Viewer.
- Create detailed reports from traffic log information in the Report Manager.
- Inspect suspicious events by correlating log information in the Log Investigator.

Monitoring Status

NSM keeps you up-to-date on the health of your network.

- View critical information about your managed devices in the Device Monitor:
 - Configuration and connection status of your managed devices
 - Individual device details, such as memory usage and active sessions
 - Device statistics
- View the status of each individual VPN tunnel in the VPN Monitor.
- View NSRP status in the NSRP Monitor.
- View the status of your IDP Clusters in the IDP Cluster Monitor.
- View the health of the NSM system itself, including CPU utilization, memory usage, and swap status in the Server Monitor.

Job Management

You can view the progress of communication to and from your devices in the Job Manager. NSM sends commands to managed devices at your request, typically to import, update or reboot devices, and view configuration and delta configuration summaries. When you send a command to a device or group of devices, NSM creates a job for that command and displays information about that job in the Job Manager module.

Job Manager tracks the progress of the command as it travels to the device and back to the management system. Each job contains:

- Name of the command
- Date and time the command was sent
- Completion status for each device that received the command
- Detailed description of command progress
- Command output, such as a configuration list or command-line interface (CLI) changes on the device



NOTE: Job Manager configuration summaries and job information details do not display passwords in the list of CLI commands for administrators that do not have the assigned activity “View Device Passwords”. By default, only the super administrator has this assigned activity.

Technical Overview


NSM architecture is built on a secure foundation, with secure communication between management components and a single access point for inbound connections.



NOTE: For details on NSM architecture and help with setting up the management system, see the *Network and Security Manager Installation Guide*.

Architecture

NSM is a three-tier management system made up of a user interface (UI), management system, and managed devices. The devices process your network traffic and are the enforcement points that implement your policies. The UI and management system tiers are software, not hardware, so you can deploy them quickly and easily. Because the management system uses internal databases for storage and authentication, you do not need LDAP or an external database. See [Figure 1 on page 11](#).

Figure 1: NSM Network Architecture

The management system also provides a programmatic interface for integrating NSM into larger enterprise business systems. This NSM API provides an alternative interface to that provided by the UI. For details, see the *Network and Security Manager API Guide*.

User Interface

The user interface (UI) provides a powerful, graphical environment for centrally managing your network. It can be installed on multiple computers on your network. You use the UI to access the management system remotely.

Multiple NSM administrators can interact with managed devices using the UI and can configure unique UI preferences. The NSM GUI Server stores user preferences in the central database so that they remain consistent when you access them from different client machines. The UI also provides extensive online help.

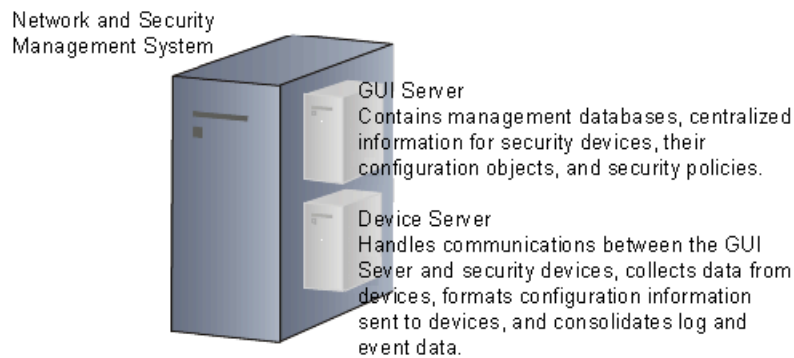
The UI communicates with the GUI Server using a secure, proprietary, TCP-based connection that encrypts and authenticates all traffic.

Management System

The management system is made up of two components:

- GUI Server
- Device Server

See [Figure 2 on page 12](#).

Figure 2: NSM System Architecture**GUI Server**

The GUI Server manages the system resources and data that drive NSM functionality. You can install the GUI Server software on a separate server or on the same server as the Device Server. The GUI Server contains the NSM databases. The GUI Server centralizes information for devices, their configurations, attack and server objects, and policies.

Specifically, the GUI Server stores all of the following information:

- Device, security policy, and VPN configuration
- NSM administrator accounts, device administrator accounts, and domains
- Objects

The GUI server also organizes and presents log entries from security devices. These log entries are actually stored on the Device Server.

The GUI Server receives logs from the Device Server on a single inbound port. When you use the UI to access NSM functionality, you connect using the same single port and access the databases stored on the GUI Server. The GUI Server communicates with the Device Server using SSP, a secure, proprietary, TCP-based connection that encrypts and authenticates all traffic.

[Table 5 on page 12](#) describes the processes that the GUI Server runs when you start it.

Table 5: GUI Server Processes

Process	Description
guiSvrManager	GUI Server Manager receives and responds to requests from the NSM UI. The GUI Server Manager forwards requests to the GUI Directive Handler or Device Directive Handler depending on the type of request for further processing.
guiSvrDirectiveHandler	GUI Directive Handler handles all directives or instructions from the NSM UI that require reading, writing, or modifying of the local data store.
guiSvrStatus Monitor	GUI Server Status Monitor monitors the status of the processes that run on the GUI Server.

Table 5: GUI Server Processes (continued)

Process	Description
guiSvrMasterController	Master Controller forwards configuration data to the NSM UI (for viewing) or to the local data store (for later retrieval).
guiSvrLicenseManager	GUI Server License Manager is responsible for license storage, retrieval, and validation.
guiSvrWebProxy	GUI Server Web Proxy responds to https requests.

Device Server

The Device Server handles communication between the GUI Server and the device, collects data from the managed devices on your network, formats configuration information sent to your managed device, and consolidates log and event data.

You can install the GUI Server and the Device Server on the same physical computer, or separate computers. Communication between a physically separate GUI Server and the Device Server is an encrypted TCP connection to a default port. The Device Server connects to the GUI Server using the default port; you can change the default port by editing the configuration files for both servers.

If the GUI Server computer and the Device Server computer have a firewall between them, you must configure a rule on that firewall to permit NSM management traffic.

[Table 6 on page 13](#) describes the processes that the Device Server runs when you start it:

Table 6: Device Server Processes

Process	Description
devSvrManager	Device Server Manager enables security devices to connect to and communicate with the NSM management system. The Device Server Manager writes log data into the local data store and routes messages and directives from the GUI Server to the Device Directive Handler for further processing.
deviceDirectiveHandler	Device Directive Handler manages directives that are issued specifically to the device (for example, a reboot, update firmware, or generate Config Summary command).
devSvrStatusMonitor	Device Server Status Monitor monitors the status of the processes that run on the Device Server.
devSvrDataCollector	Device Server Data Collector collects log data and device statistics from each device managed by NSM.
devSvrLogWalker	Device Server Log Walker performs user-specified actions on log entries (such as indexing, de-duplication, filtering).
devSvrDBServer	Device Server Database Server

Table 6: Device Server Processes (continued)

Process	Description
devSvrProfilerMgr	Device Server Profiler Manager

Managed Devices

In addition to dedicated security devices such as firewalls and IDP sensors, your managed devices can also include routers, switches, Secure Access, and Access Control devices, such as:

- [Firewall and IDP \(ScreenOS/IDP\) Devices on page 14](#)
- [Devices Running Junos OS on page 17](#)
- [SSL VPN Secure Access Products on page 20](#)
- [Juniper Networks IC Series Unified Access Control Appliances on page 21](#)
- [Extranet Devices on page 22](#)

Firewall and IDP (ScreenOS/IDP) Devices

ScreenOS/IDP devices are firewall security devices and IDP sensors and systems that you use to enable access to your network components and to protect your network against malicious traffic.

[Table 7 on page 14](#) lists the ScreenOS/IDP security devices and firmware versions supported by NSM 2011.1.

Table 7: Supported Security Devices

Security Device	Firmware Versions Supported
Juniper Networks NetScreen-5XP	ScreenOS 4.0, 5.0, 5.0 FIPS
Juniper Networks NetScreen-5XT	ScreenOS 5.0, 5.0 FIPS, 5.0 dial, 5.1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS
Juniper Networks NetScreen-5GT ADSL	ScreenOS 5.01483, 5.0 ADSL, 5.0 DSLW, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.2, 6.3
Juniper Networks Netscreen-5GT ADSLWLAN	Screen OS 5.0 DSLW, 5.0 WLAN, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.2, 6.3
Juniper Networks NetScreen-5GT WLAN	ScreenOS 5.0 WLAN, 5.0 DSLW, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.2, 6.3
Juniper Networks NetScreen-HSC	ScreenOS 5.0, 5.0 FIPS, 5.1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS
Juniper Networks NetScreen-25	ScreenOS 4.0, 5.0, 5.0 FIPS, 5.1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS

Table 7: Supported Security Devices (continued)

Security Device	Firmware Versions Supported
Juniper Networks NetScreen-50	ScreenOS 4.0, 5.0, 5.0 FIPS, 5.1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS
Juniper Networks NetScreen-204	ScreenOS 4.0, 5.0, 5.0 FIPS, 5.1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS
Juniper Networks NetScreen-208	ScreenOS 4.0, 5.0, 5.0 FIPS, 5.1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS
Juniper Networks NetScreen-500	ScreenOS 4.0, 5.0, 5.0 FIPS, 5.0 NSGP, 5.0 GPRS, 5.1, 5.1 GPRS, 5.1 shotglass, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS
Juniper Networks NetScreen-5200	ScreenOS 4.0, 5.0, 5.0 FIPS, 5.0 NSGP, 5.0L2V, 5.1, 5.1 shotglass, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.0r2, 6.1, 6.2, 6.3
Juniper Networks NetScreen-5400	ScreenOS 5.0, 5.0 L2V, 5.0 NSGP, 5.0 FIPS, 5.1, 5.1 shotglass, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.0r2, 6.1, 6.2, 6.3
Juniper Networks ISG1000	ScreenOS 5.0, 5.0 IDP1, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.0r2, 6.1, 6.2, 6.3
Juniper Networks ISG2000	ScreenOS 5.0, 5.0 FIPS, 5.0 IDP1, 5.2, 5.3, 5.3 TMAV, 5.4, 5.4 FIPS, 6.0r2, 6.1, 6.2, 6.3
Juniper Networks SSG5-ISDN	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG5-ISDN-WLAN	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG5-Serial	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG5-SB	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3.
Juniper Networks SSG5-Serial-WLAN	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG5-V92	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG5-V92-WLAN	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG-20	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG-20-WLAN	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG-140	ScreenOS 5.4, 5.4 FIPS, 6.0r2 or later, 6.1, 6.2, 6.3
Juniper Networks SSG-320	ScreenOS 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks SSG-320M	ScreenOS 6.0r2 and later, 6.1, 6.2, 6.3

Table 7: Supported Security Devices (continued)

Security Device	Firmware Versions Supported
Juniper Networks SSG-350	ScreenOS 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks SSG-350M	ScreenOS 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks SSG-520	ScreenOS 5.1 SSG, 5.4, 5.4 FIPS, 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks SSG-520M	ScreenOS 5.4, 5.4 FIPS, 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks SSG-550	ScreenOS 5.1 SSG, 5.4, 5.4 FIPS, 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks SSG-550M	ScreenOS 5.4, 5.4 FIPS, 6.0r2 and later, 6.1, 6.2, 6.3
Juniper Networks IDP10	IDP 4.0, 4.1
Juniper Networks IDP50	IDP 4.0, 4.1
Juniper Networks IDP 75	IDP 4.1, 5.0, 5.1
Juniper Networks IDP100	IDP 4.0, 4.1
Juniper Networks IDP200	IDP 4.0, 4.1, 5.0, 5.1
Juniper Networks IDP 250	IDP 4.1, 5.0, 5.1
Juniper Networks IDP500	IDP 4.0, 4.1
Juniper Networks IDP 600C	IDP 4.0, 4.1, 5.0, 5.1
Juniper Networks IDP 600F	IDP 4.0, 4.1, 5.0, 5.1
Juniper Networks IDP800	IDP 4.1, 5.0, 5.1
Juniper Networks IDP 1000	IDP 4.0, 4.1
Juniper Networks IDP 1100C	IDP 4.0, 4.1, 5.0, 5.1
Juniper Networks IDP1100F	IDP 4.0, 4.1, 5.0, 5.1
Juniper Networks IDP8200	IDP 4.2, 5.0, 5.1



NOTE: NSM supports the following ScreenOS releases: 5.0r11, 5.1r4, 5.2r3, 5.3r10, 5.4r11, 6.0r2, 6.1r4, 6.2, and 6.3.

NSM supports the IDP Release 4.0, 4.1, 4.2, and 5.0.



NOTE: Customers of NetScreen 5GT must upgrade directly from 5.4 to 6.2 as there are no intermediate releases for it. If you need to go through an intermediate non-certified release when upgrading from one certified release to the next, you must plan for a service outage and a longer upgrade time.

SSG-5-SB replaces NetScreen 5GT. SSG-5-SB is a 10-user variant of SSG-5, similar to the existing 10-user variant of NS-5GT.

Devices Running Junos OS

Devices running Junos OS and managed by NSM are listed in the following sections:

- [Juniper Networks J Series Services Routers and SRX Series Services Gateways on page 17](#)
- [Juniper Networks M Series Multiservice Edge Routers and MX Series Ethernet Services Routers on page 18](#)
- [Juniper Networks EX Series Ethernet Switches on page 19](#)



NOTE: NSM only supports the domestic version of the Junos OS and not the export version.

Juniper Networks J Series Services Routers and SRX Series Services Gateways

These routers and gateways offer not only a rich set of routing protocols and interfaces, but also firewall and IPsec virtual private network capabilities, providing high levels of security.

[Table 8 on page 17](#) lists the J Series Services Gateways, and SRX Series Services Routers, and the operating system versions supported by NSM.

Table 8: J Series Services Routers and SRX Series Services Gateways NSM Supports

Device	Versions of Junos OS NSM Supports
Juniper Networks J2320 Services Router	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks J2320 Services Router with IDP	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks J2350 Services Router	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks J2350 Services Router with IDP	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks J4350 Services Router	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4

Table 8: J Series Services Routers and SRX Series Services Gateways NSM Supports (continued)

Device	Versions of Junos OS NSM Supports
Juniper Networks J4350 Services Router with IDP	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks J6350 Services Router	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks J6350 Services Router with IDP	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX100 Fixed Platform	Junos OS Release 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX 210 Modular Platform	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX220 Modular Platform	Junos OS Release 10.3, 10.4
Juniper Networks SRX240 Modular Platform	Junos OS Release 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX650 Modular Platform	Junos OS Release 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX1400	Junos OS Release 10.4
Juniper Networks SRX3400	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX3600	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX5600	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX5600—Modular DPC	Junos OS Release 9.5, 9.6, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX5800	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.1, 10.2, 10.3, 10.4
Juniper Networks SRX5800—Modular DPC	Junos OS Release 9.5, 9.6, 10.1, 10.2, 10.3, 10.4

Juniper Networks M Series Multiservice Edge Routers and MX Series Ethernet Services Routers

Table 9 on page 19 lists the M Series and MX Series Routers, and the versions of Junos OS that NSM supports.

Table 9: M Series Multiservice Edge Routers and MX Series Ethernet Services Routers NSM Supports

Device	Versions of Junos OS NSM Supports
Juniper Networks M7i	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks M10i	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks M40e	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks M120	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks M320	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX80	Junos OS Release 10.2, 10.3, 10.4
Juniper Networks MX240	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX240 with MS-DPC PIC	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX240 with IDP services	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX480	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX480 with MS-DPC PIC	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX480 with IDP services	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX960	Junos OS Release 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX960 with MS-DPC PIC	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks MX960 with IDP services	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4

Juniper Networks EX Series Ethernet Switches

Table 10 on page 20 lists the Ethernet Switches and the versions of Junos OS that NSM supports.

Table 10: EX Series Ethernet Switches NSM Supports

Device	Versions of Junos OS NSM Supports
Juniper Networks EX2200–24P	Junos OS Release 10.1, 10.2, 10.3, 10.4
Juniper Networks EX2200–24T	Junos OS Release 10.1, 10.2, 10.3, 10.4
Juniper Networks EX2200–48P	Junos OS Release 10.1, 10.2, 10.3, 10.4
Juniper Networks EX2200–48T	Junos OS Release 10.1, 10.2, 10.3, 10.4
Juniper Networks EX3200–24P	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX3200–24T	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX3200–48P	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX3200–48T	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX4200–24F	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX4200–24P	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX4200–24T	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX4200–48P	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX4200–48T	Junos OS Release 9.2, 9.3, 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX4500–40F	Junos OS Release 10.2, 10.3, 10.4
Juniper Networks EX8208	Junos OS Release 9.4, 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4
Juniper Networks EX8216	Junos OS Release 9.5, 9.6, 10.0, 10.1, 10.2, 10.3, 10.4

SSL VPN Secure Access Products

Typical deployments of Secure Access products and clusters tend to scale rapidly as deployments grow and adapt to a wider range of applications. NSM provides a convenient way to centralize logging, monitoring, and reporting for your growing network.

[Table 11 on page 21](#) lists the Secure Access products and operating system versions supported by NSM 2011.1.

Table 11: Secure Access Products NSM Supports

Security Device	Versions of SA Software NSM Supports
Juniper Networks Secure Access 2000	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 2500	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 4000	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 4000 (FIPS)	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 4500	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 4500 (FIPS)	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 6000	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 6000 (FIPS)	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 6500	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks Secure Access 6500 (FIPS)	SA Release 6.3, 6.4, 6.5, 7.0
Juniper Networks VA-SPE	SA Release 7.0
Juniper Networks VA-DTE	SA Release 7.0

Juniper Networks IC Series Unified Access Control Appliances

In a Unified Access Control (UAC) solution, Infranet Controller (IC) products provide policy management. ScreenOS firewalls can provide the enforcement points.

[Table 12 on page 21](#) lists the Infranet Controller products and firmware versions supported by NSM 2011.1.

Table 12: IC Series UAC Appliances NSM Supports

Security Device	Versions of Firmware NSM Supports
Juniper Networks Infranet Controller 4000	IC Release 2.2, 3.0, 3.1, 4.0
Juniper Networks Infranet Controller 4500	IC Release 2.2, 3.0, 3.1, 4.0
Juniper Networks Infranet Controller 6000	IC Release 2.2, 3.0, 3.1, 4.0

Table 12: IC Series UAC Appliances NSM Supports (continued)

Security Device	Versions of Firmware NSM Supports
Juniper Networks Infranet Controller 6500	IC Release 2.2, 3.0, 3.1, 4.0
Juniper Networks Infranet Controller 6500 (FIPS)	IC Release 3.0, 3.1, 4.0

Extranet Devices

Your managed network can also include extranet devices, which are firewalls or VPN devices that are not Juniper Networks security devices.

Distributed Data Collection

The distributed data collection system provides a robust method for managing multiple objects. Each device is described by a unique Data Model (DM) that contains all the configuration data for that individual device. The Abstract Data Model (ADM) contains configuration data for all objects in a specific domain. When you use the UI to interface with your managed devices, the ADM and DMs work together:

- When you update a device configuration, the GUI Server translates the objects and object attributes in the ADM domain into device configuration information in a DM. The Device Server then translates the device configuration information in the DM into CLI commands and sends the commands to the device for ScreenOS devices. For DMI based devices, Device Server converts the DM into XML configlet and sends the configlet through NetConf protocol to the device.
- When you import a device configuration, for ScreenOS devices, the device sends CLI commands to the Device Server, which translates the CLI commands into a DM with device configuration information. For DMI devices, the device sends the configuration through NetConf protocol as an XML document to the Device Server, which translates it into a DM with device configuration information. The GUI Server then translates the device configuration in the DM into objects and object attributes in the ADM, and uses the ADM to display current information in the UI.

For more details on the ADM and DMs, see [“Managing Devices” on page 279](#).

Device Schemas

The structure of the ADM and the DMs is defined by a DM schema, which lists all the possible fields and attributes for a type of object or device. The DM schema reads from a capability file, which lists the fields and attributes that a specific operating system version supports, to determine the supported features for the operating system version that is running on the managed devices. NSM uses capability files to enable Juniper Networks software upgrades without changing the device configuration in NSM.

The device schemas for each of the firmware versions supported for ScreenOS and IDP devices are built into Network-Security Manager.

Device families introduced in Release 2008.1 and later are described by schemas that are maintained on a schema repository owned by Juniper Networks. These schemas can be added dynamically to NSM. These devices include:

- Devices running Junos OS:
 - J Series Services Routers and SRX Series Services Gateways
 - M Series Multiservice Edge Routers and MX Series Ethernet Services Routers
 - EX Series Ethernet Switches
- Secure Access products
- Infranet Controller products

See “[Managed Devices](#)” on [page 14](#) for lists of specific models of these products that support management through NSM.

Unlike schemas for ScreenOS and IDP devices, schemas for these devices can be updated asynchronously with releases of NSM. You decide when to check for new schemas, which schemas to download, and when to activate them.

Security

NSM integrates application-level encryption and authentication and uses high-grade encryption and public-key algorithms to eliminate the need for separate IPsec tunnels between each device and the management station.

For communication between the UI and the GUI Server, NSM uses Transport Layer Security (TLS), a cryptographic protocol that provides secure communication.

For communication between the GUI Server, and the Device Server, NSM uses Secure Server Protocol (SSP), a modified version of TCP that is more reliable than ordinary TCP, requires less CPU and memory resources from servers, and reduces the number of acknowledgement packets on the network. SSP uses AES encryption and SHA1 authentication for all connections.

Scaling and Performance

As you add devices or network components to your physical network, you also add them to your virtual NSM network, where you can manage all future configurations. An NSM Device Server can support up to 1000 devices; the management system supports up to 30,000 log entries per second.

Working in the User Interface

Using the NSM UI, you can configure NSM administrators, add devices, edit policies, view reports, and access the full functionality of the NSM system.



NOTE: This manual provides an overview of the UI. For step-by-step instructions on using the User Interface, click **Help** in the menu bar of the UI to access the *Network and Security Manager Online Help*.

Characters Not Supported in Login Passwords

The following characters are not supported for NSM administrator names and passwords:

- Period (.)
- Number sign (#)
- Dollar sign (\$)
- Asterisk (*)
- Ampersand (&)
- Circumflex (^)



NOTE: Passwords in the NSM UI are case-sensitive.

Managing Blocked Login Attempts

The NSM UI blocks hosts that fail to login after 10 attempts by default. Use the **Tools > Preferences > System Properties** option to change the number of attempts. Use the **Tools > Manage Blocked Hosts** option to unblock hosts that have been locked out of the UI because of excessive failed login attempts.

Configuring UI Preferences

You can configure additional preferences for UI behavior, such as appearance, external tool use, polling statistics, and UI timeout. For details on configuring these settings, see the topics under “Network and Security Manager User Interface” in the *Network and Security Manager Online Help*.

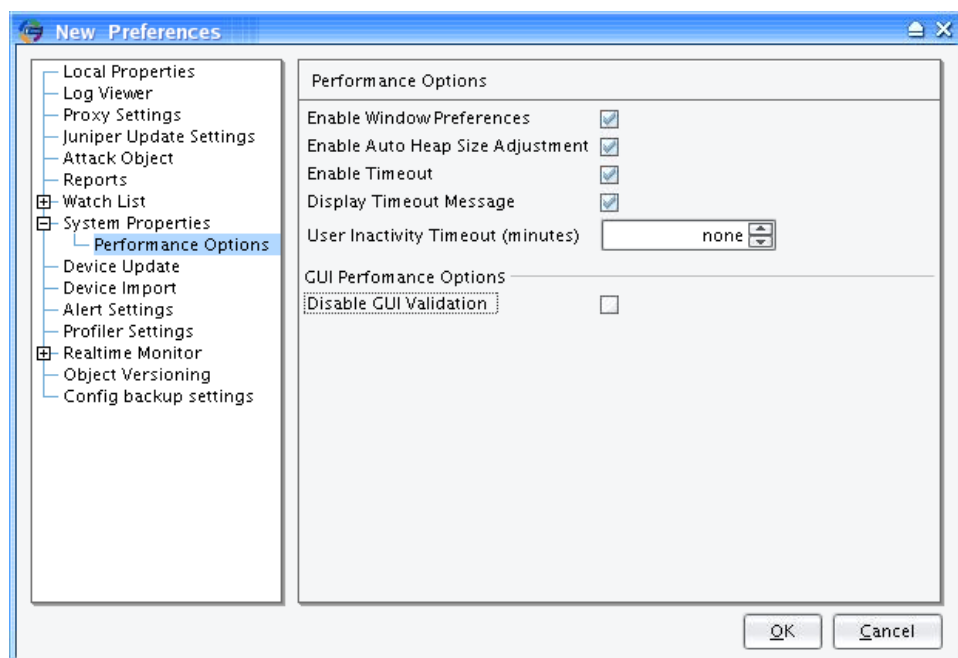
Disabling GUI Validation

Validation processing time depends on data size; large data bytes increase validation time. GUI performance issues resulting from delays in validation can be improved by enabling the **Disable GUI Validation** flag. Disabling GUI validation is client-specific and each client has its own validation disable flag.

To configure preferences to disable validation:

1. Select **Tools > Preference** from the menu bar.
The New Preferences dialog box appears.
2. In the preference navigation tree, select **System Properties > Performance Options**.
The Performance Options appear. [Figure 3 on page 25](#) shows the preference options.

Figure 3: GUI Preference Options - Disabling GUI Validation



3. Select the **Disable GUI Validation** check box.



NOTE: The validation option is disabled by default and performs normal validation. You can enable the option to increase the performance and the time consumed for validation.

4. Select the GUI category where the validation need to be turn off. [Figure 4 on page 25](#) shows the GUI category options.

Figure 4: Disable GUI Validation Options



NOTE:

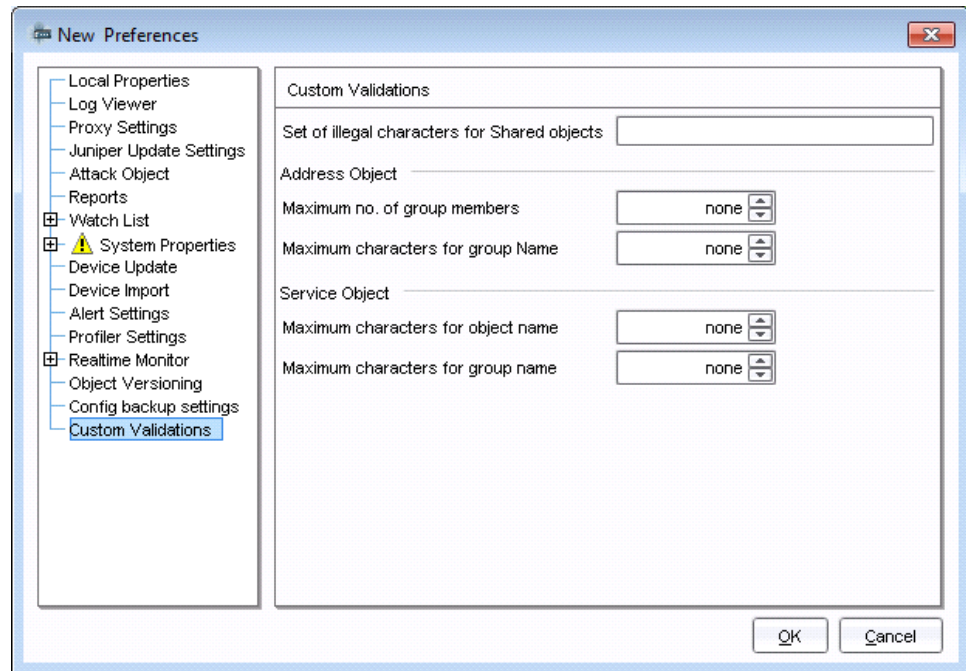
- For each selection, NSM will display a warning about how and where the selection will affect.
- When the shared object option is deselected validation for Address, services and attacks is stopped completely.
- When the Firewall policy option is deselected the warning validate is stopped. But, the error message validation is still active.

5. Click **OK**.

Customizing Validation

Beginning in NSM 2012.2 release, you can customize the validation options as per your requirement. This option supports various invalid character set and name length in SRX Series devices. [Figure 5 on page 26](#) shows the custom validation options.

Figure 5: GUI Preference Options - Custom Validations



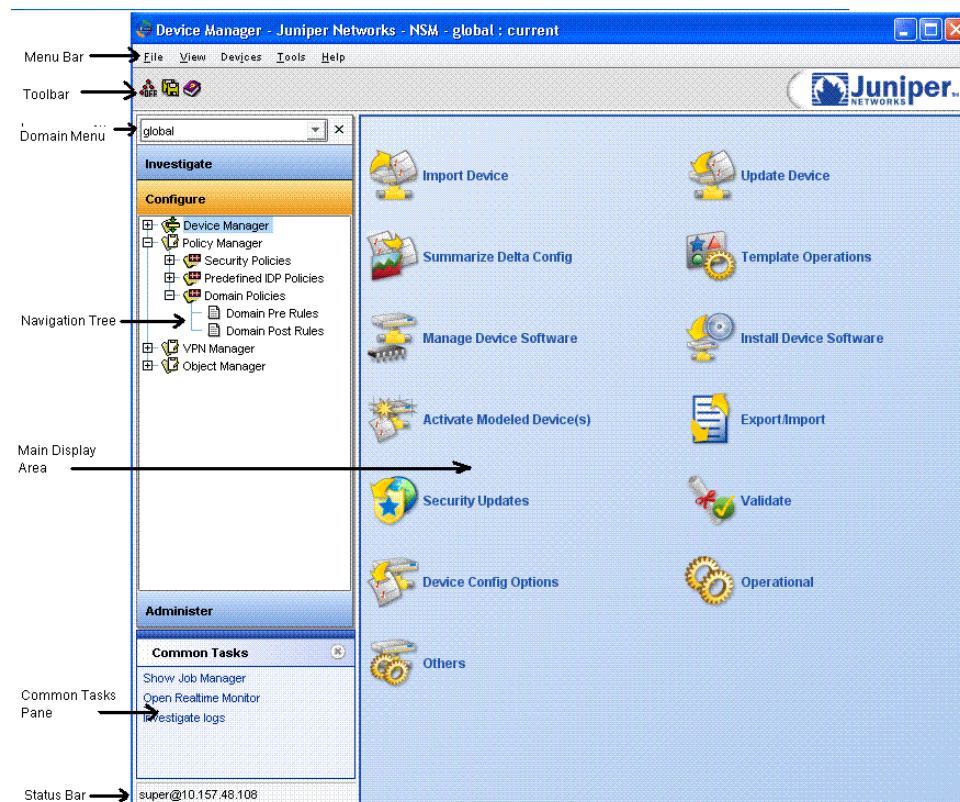
The options available are:

- **Set of illegal characters for shared objects**—Specifies the illegal characters for shared objects.
- **Address Object > Maximum number of group objects**—Specifies the maximum number of group objects to be allowed in a group object for an address object.
- **Address Object > Maximum characters for group name**—Specifies the maximum characters to be allowed in a group name for an address object.
- **Service Object > Maximum characters for object name**—Specifies the maximum characters to be allowed in an object name for a service object.
- **Service Object > Maximum characters for group name**—Specifies the maximum characters to be allowed in a group name for a service object.

UI Overview

The NSM UI appears after you log in, and displays a set of menus and toolbar icons at the top of the UI window. For some components, right-click menus are available to perform tasks. [Figure 6 on page 27](#) shows a sample UI screen.

Figure 6: Overview of the User Interface



Navigation Tree

The navigation tree provides three panels:

- Investigate panel—Provides NSM modules with tree structures for monitoring your network.
- Configure panel—Provides NSM modules with tree structures for configuring devices, policies, VPNs, and other objects.
- Administer panel—Provides NSM modules with tree structures for managing the NSM servers, ongoing jobs, and other actions.

For details about each module, see [“NSM Modules” on page 28](#).

Common Tasks Pane

The Common Tasks pane provides links to commonly accessed tasks throughout the UI. These common tasks change depending on what tasks are often selected in the UI.

Main Display Area

The main display area displays content for the selected module, module contents, or launch pad views. Launch pad views make use of otherwise blank panes to provide access to commonly used functionality within the module.

Menu Bar

The menu bar contains clickable commands. You can access many menu bar commands using keyboard shortcuts. For a complete list of keyboards shortcuts, see the *Network and Security Manager Online Help*.

Toolbar

The toolbar contains buttons for common tasks. The buttons displayed in the toolbar are determined by the selected module.

Status Bar

The status bar displays additional information for a selected module.

NSM Modules

The navigation tree splits top-level modules into three panels:

- [Investigate Modules on page 28](#)
- [Configure Modules on page 30](#)
- [Administer Modules on page 34](#)

Investigate Modules

The Investigate panel includes the following top-level modules:

- [Log Viewer on page 28](#)
- [Report Manager on page 29](#)
- [Log Investigator on page 29](#)
- [Realtime Monitor on page 29](#)
- [Security Monitor on page 29](#)
- [Audit Log Viewer on page 30](#)

Log Viewer

The Log Viewer displays log entries that your security devices generate based on criteria that you defined in your security policies, on the GUI Server, and in the device configuration. Log entries appear in table format; each row contains a single log entry, and each column defines specific information for a log entry.

You can select which log entries and what log information is shown using log filters or by changing the column settings.

Use the Log Viewer to:

- View summarized information about security events and alarms.
- View information about a specific log entry.
- Show, hide, or move columns to customize the Log Viewer.
- Filter log entries by column headings.

- Create and save custom views that display your filters and column settings.
- Set flags on Log Viewer entries to indicate a specific priority or action.

NSM supports log management for ScreenOS devices, IDP sensors, J Series devices, and EX Series devices. NSM does not support log management for SRX Series devices, M Series devices, and MX Series devices.

For more details on using the Log Viewer, see [“Logging” on page 783](#).

Report Manager

The Report Manager contains summary, graphs, and charts that describe specific security events that occur on your network. NSM generates reports to show the information contained in your log entries. You can use reports to summarize security threats to your network, analyze traffic behavior, and determine the efficiency of NSM. To share reports or to use report information in other applications, you can print or export report data.

NSM supports report management for ScreenOS devices, IDP sensors, J Series devices, and EX Series devices. NSM does not support report management for SRX Series devices, M Series devices, and MX Series devices.

Log Investigator

The Log Investigator contains tools for analyzing your log entries in depth. Use the Log Investigator to:

- Manipulate and change constraints on log information.
- Correlate log entries visually and rapidly.
- Filter log entries while maintaining the broader picture.

Realtime Monitor

Realtime Monitor provides a graphical view of the current status of all devices managed by NSM:

- Device Monitor—Tracks the connection state and configuration state of your managed devices. You can also view device details to see CPU utilization and memory usage for each device, or check device statistics.
- VPN Monitor—Tracks the status of all VPN tunnels.
- NSRP Monitor—Tracks the status of security devices in clusters.
- IDP Cluster Monitor—Tracks the status of IDP clusters.

You can customize Realtime Monitor to display only the information you want to see, as well as to update information at specified intervals. You can also set alarm criteria for a device or process. For more details on Realtime Monitor, see [“Realtime Monitoring” on page 705](#).

Security Monitor

Security Monitor provides access to the Dashboard, Profiler, and Security Explorer. These tools enable you to track, correlate, and visualize aspects about your internal network,

enabling you to create more effective security policies and minimize unnecessary log records. For more details, see [“Analyzing Your Network” on page 753](#).

The Security Monitor applies to ScreenOS devices and IDP sensors. It does not apply to J Series, SRX Series, Secure Access, Infranet Controller, M Series, MX Series, or EX Series devices.

Audit Log Viewer

The Audit Log Viewer contains a log entry for every change made by an NSM administrator. For more details on Audit Log Viewer, see [“Using the Audit Log Viewer” on page 832](#).

Configure Modules

The Configure panel include the following top-level modules:

- [Device Manager on page 30](#)
- [Policy Manager on page 31](#)
- [VPN Manager on page 31](#)
- [UAC Manager on page 32](#)
- [Object Manager on page 32](#)

Device Manager

The Device Manager contains the device objects that represent your managed devices. You can create or modify:

- ScreenOS security devices and IDP sensors—The devices you use to enable access to your network and to protect your network against malicious traffic.
- Devices running Junos OS:
 - EX Series Ethernet Switches—Enterprise-class switches managed by NSM.
 - J Series Services Routers—Routers managed by NSM.
 - SRX Series Services Gateways.
 - M Series Multiservice Edge Routers and MX Series Ethernet Services Routers.
- Secure Access products—SSL VPN systems managed by NSM.
- Infranet Controller products—Unified Access control systems managed by NSM.
- Vsys devices—Virtual devices that exists within a physical security device.
- Clusters—Two managed devices joined together in a high availability configuration to ensure continued network uptime.
- Vsys cluster—A vsys device that has a cluster as its root device.
- Extranet devices—Firewalls or VPN devices that are not Juniper Networks security devices.
- Templates—A partial device configuration that you can define once, and then use for multiple devices.
- Device Groups—A user-defined collection of devices.

- **Device Discovery Rules**—Sets of rules that define subnets or ranges of IP addresses to scan for EX Series devices in your network.
- **Topology Views**—Graphic and tabular views of the network topology generated by the topology discovery engine, providing device information as well.

Policy Manager

The Policy Manager manages security policies that contain the firewall, multicast, and VPN rules that control traffic on your network for devices that support centralized policy management. Using a graphical, easy-to-use rule building platform, you can quickly create and deploy new policies to your security devices.

Use the Policy Manager to:

- Add or modify existing security policies.
- Add or modify existing VPN rules.
- Add or modify existing IDP rules.
- Create new policies based on existing policies.
- Install policies on one or multiple devices.
- Delete policies.

If the device configurations that you import from your security devices contain policies, the Policy Manager displays those imported policies. For details on editing those imported policies or creating new policies, see [“Configuring Security Policies” on page 473](#), or [“Configuring VPNs” on page 597](#).

You can configure policies for ScreenOS and IDP devices using Policy Manager. For Secure Access, Infranet Controller, and EX Series devices, you must configure policies in the device. For J Series routers, SRX Series gateways, and MX Series routers, you can configure policies either in the Central Policy Manager or in the device, but not both.

VPN Manager

The VPN Manager contains the VPN objects that control the VPN tunnels between your managed devices and remote users. Using VPN objects, such as Protected Resources and IKE Proposals, you can create multiple VPNs for use in your security policies.

Use the VPN Manager to:

- Define the protected resources on your network—the network resources you want to protect in a VPN.
- Create custom IKE Phase 1 and 2 Proposals.
- Configure AutoKey IKE, L2TP, and L2TP-over-AutoKey IKE VPNs in policy-based or route-based modes. You can also create an AutoKey IKE mixed mode VPN to connect policy-based VPN members with route-based VPNs members.
- Configure AutoKey IKE and L2TP policy-based VPNs for remote access services (RAS) and include multiple users.

NSM supports VPN management for ScreenOS devices, IDP sensors, J Series devices, and SRX Series devices.

UAC Manager

The UAC Manager enables you to create and view associations between Infranet Controllers (IC) and Enforcement Points (EP) in a network. You can choose between IC views and EP views. The IC view provides a list of EPs associated with the IC and their location groups. You can associate or disassociate EPs from a particular IC. The EP view provides a list of associated ICs and their port details. You can use this feature to resolve configuration conflicts, and enable or disable 802.1X ports on enforcement points.

Object Manager

The Object Manager contains objects, which are reusable, basic NSM building blocks that contain specific information. You use objects to create device configurations, policies, and VPNs. Objects are shared by all devices and policies in a domain.

You can create the following objects in NSM:

- Access Profiles—An access profile consists of a set of attributes that defines access to a device. You can create access profile objects and share them across security policies that are assigned to J Series Services Routers and SRX Series Services Gateways managed by NSM.
- Address objects—Represent components of your network—hosts, networks, servers.
- Attack objects—Define DI profiles and IDP attack objects.
 - DI Profiles—Define the attack signature patterns, protocol anomalies, and the action you want a security device to take against matching traffic.
 - IDP attack objects—Define attack patterns that detect known and unknown attacks. You use IDP attack objects within IDP rules.
- Custom Policy Fields objects—Represent metadata information that you can store and use in a structured manner. Users can add custom objects to the policy table, such as ticket number, vendor contact, and so on, for each rule in the rulebase. NSM provides a shared object to store these custom details while the table contains a column that corresponds to these custom details.
- AV objects—Represent the AV servers, software, and profiles available to devices managed by NSM.
- ICAP objects—Represents the Internet Content Adaptation Protocol (ICAP) servers and server groups used in ICAP AV objects.
- GTP objects—Represent GTP client connections.
- Authentication Servers—Represent external authentication servers, such as RADIUS and SecurID servers. You can use an authentication server object to authenticate NSM administrators (RADIUS only), XAuth users, IKE RAS users, and L2TP users.
- Certificate Authority objects—Represent the certificate authority's certificate.
- CRL objects—Represent the certificate authority's certificate revocation list.

- Group Expressions—These logical expressions include OR, AND, and NOT statements that set conditions for authentication requirements.
- IP Pools—Represent a range of IP addresses. You use IP pools when you configure a DHCP Server for your managed devices.
- NAT objects—Represent MIPs, VIPs, and DIPs.
- Remote Settings—Represent DNS and WINS servers. You use a remote settings object when configuring XAuth or L2TP authentication in a VPN.
- Routing instance objects— A routing instance is a collection of routing tables, interfaces contained in these routing tables, and routing option configurations. A routing instance object configured in Object Manager can be included in the RADIUS server and LDAP server configurations within the access profile object. A routing instance object is a polymorphic object (similar to zone objects) that maintains the mapping between the actual routing instance and the device in which it is created.
- Regional Servers—Represent NSM servers managed by a Central Manager.
- Zone objects—Represent zones in a Central Manager or Regional Server.
- Schedule objects—Represent specific dates and times. You can use schedule objects in firewall rules to specify a time or time period that the rule is in effect.
- Web filtering objects (Web Profiles)—Define the URLs, the Web categories, and the action you want a security device to take against matching traffic.
- Service objects—Represent services running on your network, such as FTP, HTTP, and Telnet. NSM contains a database of service objects for well-known services; you can also create new service objects to represent the custom services you run on your network.
- User objects—Represent the remote users that access the network protected by the security device. To provide remote users with access, create a user object for each user, and then create a VPN that includes those user objects.
- VLAN objects—Limit rule matching to packets within a particular VLAN.
- VSYS Profile object—Represent profiles of resource limits for vsys devices.
- Unified Threat Management objects — Create threat management profiles for common objects.
- Extranet policies objects—Enable you to configure and manage extranet devices, such as routers from other vendors.
- Binary Data—Enables efficient management of large binary data files used in the configuration of Secure Access and Infranet Controller devices.

You can use the Object Manager to:

- View and edit the object properties.
- Create, edit, or delete objects.
- Create custom groups of objects.

For more details on objects, see [“Configuring Objects” on page 345](#).

Administer Modules

The Administer panel includes the following top-level modules:

- [Server Manager on page 34](#)
- [Job Manager on page 34](#)
- [Action Manager on page 34](#)

Server Manager

Server Manager contains server objects that represent your management system components.

- Servers—Manage the individual server processes that make up your NSM system.
- Server Monitor—Monitors the status of your NSM servers.
- Schema Information—Allows you to manage the update and activation of schemas.

Job Manager

Job Manager contains the status of commands (also called directives) that NSM sends to your managed devices. You can view summaries or details for active jobs and completed jobs. For more details on Job Manager, see [“Tracking Device Updates” on page 273](#).

Action Manager

The Action Manager enables you to forward logs on a per-domain basis. For more details on using the Action Manager, see [“Using the Action Manager to Forward Logs by Domain” on page 841](#).

Validation Icons in the User Interface

NSM uses automatic validation to help you identify the integrity of a configuration or specific parameter at a glance. The icons shown in [Table 13 on page 34](#) might appear as you work in the UI:

Table 13: Validation Status for Devices





Icon	Meaning
	Error. Indicates that a configuration or parameter is not configured correctly in the NSM UI. Updating a device with this modeled configuration will cause problems on the device.
	Warning. Indicates that a configuration or parameter is not configured correctly in the NSM UI. Updating a device with this modeled configuration might cause problems on the device.
	Needs Validation. Indicates that a configuration or parameter has not been validated. Although NSM automatically validates all parameters when entered, this icon might appear for a template-driven value after you have changed a template. We highly recommend that you validate all parameters before updating a device.







Table 13: Validation Status for Devices (continued)

Icon	Meaning
	Valid. Indicates that a configuration or parameter is configured correctly in the NSM UI.

Validation and Data Origination Icons

Data origin tooltips show the user where field data originates. These tooltips are implemented as additional types of validation messages (beyond the current Error and Warning messages), adding Template Value, Override, and From Object messages. Each has its own icon and text color in the tool tips, as shown in [Table 14 on page 35](#).

Table 14: Validation Icons

Icon	Message Type	Meaning	Priority
	Error	Indicates that a configuration or parameter is not configured correctly in the NSM UI. Updating a device with this modeled configuration will cause problems on the device.	Highest
	Warning	Indicates that a configuration or parameter is not configured correctly in the NSM UI. Updating a device with this modeled configuration might cause problems on the device.	
	Override	Indicates that the displayed value was set manually and that the value overrides whatever value might come from a template. The icon can also indicate an override of a VPN-provided value or a cluster-provide value. Changes to a template will not change this value unless "Remove conflicted device values" is selected in the template Operations dialog box.	
	Template Value	Indicates the value was inherited from a template. Changes to the template are also shown in the device edit dialog box.	
	Configuration Group	Indicates the value was inherited from a configuration group. Changes to the configuration group are also shown in the device edit dialog box.	
	From Object	Indicates that a value is set for a field in a template or configuration group definition. This icon is shown only in a template or configuration group definition. From Object messages appear only when you view template objects to help find fields set in the template.	Lowest

When more than one type of icon appears within a panel, the highest-priority icon appears next to the icon in the tree and the panel title bar.

Working with Other NSM Administrators

When multiple NSM administrators access the NSM system at the same time, NSM ensures that all edits are synchronized by locking an active object. Only one administrator at a time can edit existing values for an object, but multiple administrators can still view the existing values for that object.

- When an NSM administrator begins editing an object, the UI locks that object to prevent other administrators from editing the object's value.
- During lockout, NSM makes "lazy" saves of all edits made and stores them in an in-memory database. If NSM crashes during a lazy save, edits made since the last lazy save are lost, and NSM prompts the NSM administrator to roll back to the last lazy save.
- When the administrator completes and saves the edit, the object is unlocked, enabling other administrators to edit it. However, because the UI does not immediately refresh the object values, you must manually refresh the UI to view the most recent versions.

When you attempt to open a locked object, a warning message indicates that the object is locked and can be opened only as a read-only object. The warning message also contains the name of the NSM administrator who is editing the object. Depending on your administrator privileges, you can locate contact information for the administrator in the Manage Administrators and Domains area of the UI (From the file menu, select **Tools > Manage Administrators and Domains**). For details on working with administrators and domains, see ["Configuring Role-Based Administration" on page 68](#).

For example, let's say Bob and Carol are both NSM administrators with the same roles. If both administrators view the same object, but Bob also edits and saves the object, NSM does not notify Carol that a newer version of the object exists. To see the newest version, Carol must first close, then open the object again or refresh the console.


Searching in the User Interface

You can use the integrated search feature in NSM to quickly locate a specific setting within a UI screen or dialog box.

To locate a word, begin typing the word. The search window appears in the top left of the selected screen or dialog box. The UI attempts to match your entry to an existing value; as you enter more characters, the UI continues to search for a match. Use the arrow keys to move between matching values. If your entry appears in red, no matching value was found within the selected screen or dialog box.

To locate a different data type, such as an IP address, change the search mode. To display all available search modes, press the backslash key (\). The search mode window appears, as shown in [Figure 7 on page 37](#).

Figure 7: UI Search Modes

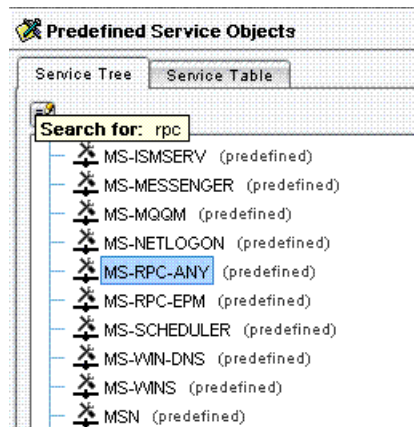
	[C]-find pattern inside entire search string.
	[S]-search for value starts with pattern.
 Select Mode:	[R]-use a Regular Expression to find a value.
	[I]-search for an IP. "*" notation is allowed.
	[E]-search for Exact Match.

Press the key that represents the search mode you want to use, and then begin typing the search criteria. You can search within a category or across different categories in the Object Manager. Press the ESC key to end the search operation and close the window. The following sections provide examples of each search mode.

Contains String [C] Search Mode

Use to locate a pattern anywhere in a string. For example, to locate the pattern “RPC” in service objects:

1. In the main navigation tree, select **Object Manager > Service Objects > Predefined Service Objects**, and then select the Service Object icon at the top of the Service Tree tab.
2. Press the backslash key (\) to display the search mode window.
3. Enter **C**, and then enter **RPC**. The UI automatically highlights the first match, MS-RPC-ANY, as shown in [Figure 8 on page 37](#).

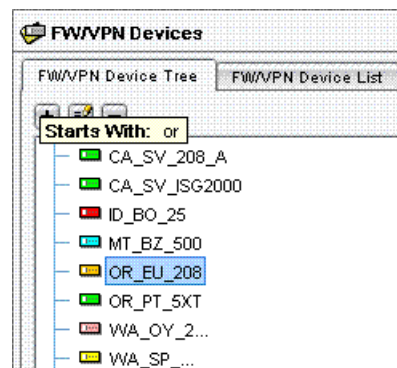
Figure 8: “Contains String” Search Mode Example

Starts With [S] Search Mode

Use to locate a pattern at the beginning of a string. For example, to locate the pattern “OR” in devices:

1. In the main navigation tree, select **Device Manager > Devices**, then select the security devices icon at the top of the Device Tree window.
2. Press the backslash key (\) to display the search mode window.
3. Enter **S**, then enter **OR**. The UI automatically highlights the first match, OR_EU_208, as shown in [Figure 9 on page 38](#).

Figure 9: “Starts With” Search Mode Example

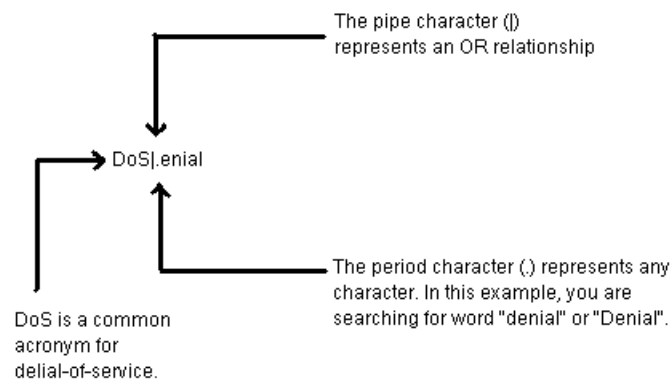


Regular Expression [R] Search Mode

Use to locate a value using a regular expression. For example, to locate all DI attack objects that detect denial-of-service attacks:













1. In the main navigation tree, select **Object Manager > Attack Objects > DI Objects**, and then select the Predefined Attacks tab.
2. Select the first entry in the column Name, and then press the backslash key (\) to display the search mode window.
3. Enter **R**, and then enter the following characters: **DoS|.enial**. [Figure 10 on page 38](#) details this expression:

Figure 10: “Regular Expression” Search Mode Details



The UI automatically highlights the first match; click the down arrow key to highlight the next match. Both matches are shown in [Figure 11 on page 39](#).

Figure 11: “Regular Expression” Search Mode Example

Match 1	Predefined Attacks	Predefined Attack Groups
	VR[el]: "DoS].enial	Severity /
	 SMB Error: Invalid Message Length	 Major
	 SMB Error: Malformed Message	 Major
	 DOS Network Device: 3Com OfficeConnect HTTP Router Denial of Service	 Major
	 CISCO IOS httpd DoS	 Major
	 FTP:Line Too Long	 Major
	 FTP:Password Too Long	 Major



NOTE: The regular expression search mode supports all common regular expressions. For more information about regular expressions, refer to a dedicated resource, such as *Mastering Regular Expressions*, 2nd Edition, by Jeffrey E. F. Friedl.

IP [I] Search Mode

Use to locate an IP address. For example, to locate the IP address 5.5.5.50 and 5.5.5.51 in address objects:

1. In the main navigation tree, select **Object Manager > Address Objects**, then select the Address Table tab.
2. Select the first entry in the column IP/Domain Name, and then press the backslash key (\) to display the search mode window.
3. Enter I, and then enter **5.5.5.***. The UI automatically highlights the first match, **5.5.5.50**. Click the down arrow key to highlight the next match, **5.5.5.51**.

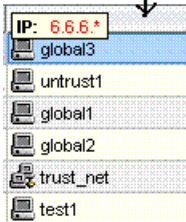
When searching in a table, your search criteria are applied only to the selected column. If you select a different column, such as Name, and perform the same search, the results differ. [Figure 12 on page 40](#) shows both search results.



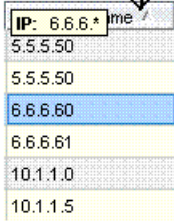
NOTE: NSM Release 2009.1 allows you to search for an IP address with its specific netmask.

Figure 12: “IP Address” Search Mode Example

Name	Type	IP/Domain Name	Netmask
global3	Host	5.5.5.50	32
untrust1	Host	5.5.5.50	32
global1	Host	6.6.6.60	32
global2	Host	6.6.6.61	32
trust_net	Network	10.1.1.0	24
test1	Host	10.1.1.5	32
Kayak	Host	10.1.1.76	32
Internal Network	Network	10.10.1.1	24
Web Server	Host	10.10.1.255	32
FTP Server	Host	10.10.10.254	32
trust_jan	Network	10.100.2.0	24



Unsuccessful search



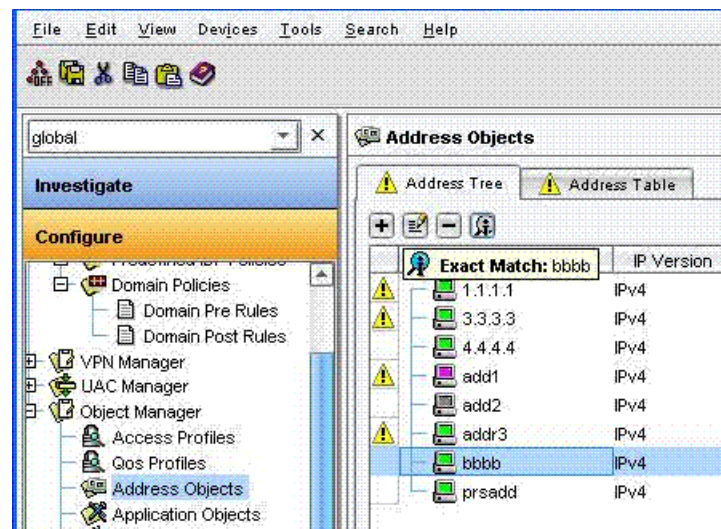
Successful search

Search for an Exact Match (E)

You can search for an exact string in all the address, attack, service and security policy categories. For example, to locate the string **bbbb**:

1. In the main navigation tree, select **Object Manager > Address Objects**, then select the **Address Table** tab.
2. Select any entry in the **Name** column, and then press the backslash key (\) to display the search mode window.
3. Enter **E** and then type **bbbb**. NSM highlights the matching object as depicted in Figure 10.

Figure 13: Exact String Search Mode Example



Global Search

To perform a global search:

1. Select **Global Search** from the **Search** menu. The **New Search** dialog box appears. It allows you to specify search criteria across different views and to switch between different views.
2. Select a category from the drop-down list in the **Search Category** field. You can select **All**, **Address**, **Service**, **Attacks**, **IDP Attack Group**, **DI Profile**, **Attack Group**, or **Policy**.
3. Select a filter to be applied to the chosen category from the drop-down list in the **Filter** field. You can select **Name**, **IP**, **Comment**, **Service**, or **Install On**.

Depending on the filter you selected, NSM prompts you for related information.

- If you select **Name**, you must enter the name of the object in the **Name** field. You can then specify whether you want the search to be a **Case Sensitive** or **Regular Expression** type of search.
- If you select **IP** as a filter, NSM asks you for the **IP Version**, the **IP address** and the **Netmask**. If you enter both the IP address and the netmask, NSM uses the combined criteria for the search.
- If you select **Comment**, enter the comment you are searching for in the **Comment** field. You can then specify whether you want the search to be a **Case Sensitive** or **Regular Expression** type of search.
- If you select **Service**, you can click the **Select Service** button to view a list of services. Check the desired services and click **OK** to select multiple services. Your selection

of services appears in the **Services** box. If you select multiple services, NSM uses the combined criteria for the search.

- If you select **Install On**, you can click the **Select Target** button to view a list of target devices. Check desired devices and click **OK** to select multiple target devices. Your selection of devices appears in the **Install On** box.
4. Click the **Search** button to execute the search. The **Search Results** appear at the bottom of the dialog box. The applicable search category is listed to the left and the matching search objects are listed to the right.
 5. Use the buttons above the list of search results to add or search for more results, edit a result, and delete a result.
 6. Click **Close** to exit the search.

New Features in 2012.2

NSM release 2012.2 supports:

- NSM supports automatic purging of database versions at configured intervals. When saved database files exceed the maximum threshold, only the configured minimum database version is retained.
- NSM supports loading of device schema based on the startup options provided in the installer script. After NSM is installed you can also enable this option in the GUI and dev server configuration files.
- NSM caching has been enhanced to clean up all objects in the cache to make conservative use of available memory in the Java process.
- On ISG Series and SRX Series devices, NSM provides the option to update only the firewall configuration without updating the IDP configuration and reducing the total update time required for the devices.
- NSM supports software upgrade for eight Junos OS devices concurrently.
- NSM provides nested service group support for J Series and SRX Series devices.
- Performance of shared object loading such as address, service and custom attack object in NSM has been enhanced.
- Performance of xdb query responses in NSM database version-restriction commands has been enhanced.
- Performance of delta and update workflows in NSM has been enhanced.
- Performance of firewall policy loading in NSM GUI has been enhanced.
- Find usage workflow has been enhanced, to cleanup stale references. You will be prompted to cleanup the stale references, if they are found as a part of the find usage process.

CHAPTER 2

Planning Your Virtual Network

When you use Network and Security Manager (NSM) to manage your devices, you are creating a virtual network that represents your physical network. Using this virtual network, you can create, control, and maintain the security of your physical network at a system level.

This chapter provides information to help you decide how best to create your virtual network and simplify management tasks.



NOTE: Not all devices support all features described in this guide. For device-specific datasheets that include an updated feature list for each device, go to <http://www.juniper.net>. In the Products and Services section, select the product family.

This chapter contains the following sections:

- [Configuring Devices Overview on page 43](#)
- [Configuring IDP-Capable Devices Overview on page 47](#)
- [Simplifying Management on page 56](#)
- [Creating an Information Banner on page 59](#)

Configuring Devices Overview

To manage Juniper Networks devices that already exist on your network, you can import their device configurations into NSM. Each imported device appears in the NSM UI, where you can view or make changes to the device, such as change settings in the device configuration, edit the security policy for the device, and upgrade device firmware.

For new devices that do not yet exist on your network, you can create their device configuration in NSM. When you physically deploy your device, you can install the modeled device configuration on that device to instantly get it up and running. After you install the modeled configuration on the device, you can manage the device just as you would an imported device.



NOTE: You cannot install a modeled device configuration on a Secure Access or Infranet Controller device. These devices must be added to NSM by importing.



NOTE: Juniper Networks also offers security devices with Intrusion Detection and Prevention (IDP) capability. For details on how to enable IDP functionality on these devices, see [“Configuring IDP-Capable Devices Overview” on page 47](#).

Importing Existing Devices

For networks with deployed devices, if you have already designed, staged, and set up a working physical device, you don't need to repeat that process; you can import that device so it exists (virtually) inside the management station. Importing includes the routing, IP configuration, access and security policies, access privileges, and other device-specific information defined on the device.

To import existing devices:

1. Add the security device and import your device configuration.
 - a. In the NSM main navigation tree, select **Device Manager > Devices**.
 - b. In the main display area, click the Add icon and select **Device**. Follow the instructions in the Add Device Wizard to import an existing device.

As NSM imports the existing device configuration, it automatically creates all objects and policies in the configuration.



NOTE: NSM does not import IDP rulebases in a security policy when importing the device configuration.

For details on adding and importing existing devices, see [“Importing Devices” on page 114](#).

2. Verify the imported device configuration and related information:
 - Run a Delta Config Summary and view the results to check for differences between the physical device configuration and the device object configuration imported into NSM.
 - Check device configuration information.
 - Check Address, Service, Schedule, and NAT objects.
 - Check security policies.
 - Check protected resources.
 - Check VPNs.

3. Correct any validation errors, if found, and check for duplicate objects (such as address objects, custom service objects). Be sure to consolidate any duplicate objects before importing another device.

You can also delete devices from NSM, and reimport them if necessary. Deleting a device removes all device configuration information from the management system, but might be the best solution if you need to perform extensive troubleshooting or reconfigure the device locally. After you have made the necessary changes locally, you can then reimport that device into the NSM system.

For details on adding devices, see [“Adding Devices” on page 99](#).

Modeling New Devices

For new networks or networks that do not use a previously deployed Juniper Networks device, you should review your network topology thoroughly and design a security system that works for your organization.

When creating a new security network using NSM:

1. Create the domain structure that best suits your network topology and access requirements.
2. Create NSM administrators and set their permission level by creating and assigning roles. See [“Configuring Role-Based Administration” on page 63](#) for details.
3. Add your devices and model their device configurations in NSM.
 - Use templates to configure multiple devices. Templates help you reuse common information to quickly create configurations for similar devices.
 - For ScreenOS 5.x and later devices, you can use Rapid Deployment (RD) to deploy multiple devices in nontechnical locations. Use RD to stage and configure devices quickly, and then simultaneously update all devices with policies to control traffic as desired in multiple locations.



NOTE: Secure Access and Infranet Controller devices must be imported into NSM.

4. Create the objects used in your security policies. These objects might include:
 - NAT objects for policy-based network address translation
 - Address objects for your network components
 - Service objects for your custom network services (NSM includes an object database of common transport and application-level services)
 - AV objects for detecting viruses in your network traffic
 - GTP objects for inspecting GTP packets

For details about creating objects, see [“Configuring Objects” on page 345](#).

5. Create security policies.

NSM integrates policy management, linking multiple devices to one security policy that defines the type of traffic permitted on the network and how that traffic is treated inside the network.



NOTE: You can use the NSM Policy Manager to centralize policy management for certain devices. Devices that support central policy management include ScreenOS and IDP devices and any J Series routers or SRX Series gateways configured for central policy management.

- Add a policy, and then create firewall rules that specify source, destination, service, and action. You can also create multicast rules to handle multicast control traffic.
- Verify each policy using the Policy Validation tool.

For details of configuring policies, see [“Configuring Security Policies” on page 473](#).

6. Update devices after they are deployed. This action pushes the modeled configuration to the deployed device.

- Resolve any validation issues with the device configuration.
- View a summary of the device configuration to ensure that all device parameters are correct.
- Check progress in Job Manager.

For details about pushing a configuration to a device, see [“Updating Devices” on page 257](#).

7. Create VPN rules.

- Create Protected Resources.
- Create user objects and User Groups for RAS VPNs.
- Use VPN Manager to select VPN members, and then automatically generate the rules for each member.

For details about configuring VPNs, see [“Configuring VPNs” on page 597](#)

For details on adding devices, see [“Adding Devices” on page 99](#). For details on configuring devices, see [“Configuring Devices” on page 199](#).

Editing a Device Configuration

After importing or modeling a device configuration in NSM, you can edit that configuration in NSM. For configuration changes to become effective, however, you must update the device by pushing the edited configuration to it. For details about pushing a configuration to a device, see [“Updating Devices” on page 257](#).

Conversely, the device configuration can be edited by the device administrator using the device's native GUI or CLI. To synchronize the device object configuration in NSM with the actual device, you must then reimport the device.

Configuring IDP-Capable Devices Overview

Although firewalls provide basic protection, they are not designed to detect all attacks. Advanced attack methods often elude firewall detection by embedding an attack within permitted traffic or by using attack vectors that are outside the firewall's detection capability.

When deployed inline in your network, Juniper Networks Intrusion Detection and Prevention (IDP) technology can detect—and stop—attacks. Unlike IDS, IDP uses multiple methods to detect attacks against your network and prevent attackers from gaining access and doing damage. IDP can drop malicious packets or connections before the attacks can enter your network. IDP is designed to reduce false positives and ensure that only actual malicious traffic is detected and stopped. You can also deploy IDP as a passive sniffer, similar to a traditional IDS, but with greater accuracy and manageability.

Common Criteria EAL2 Compliance

All Juniper Networks IDP Sensors meet the Common Criteria requirements for Common Criteria EAL2. This section describes actions that are required for a security administrator to properly secure the NSM system and NSM User Interface to be in compliance with the Common Criteria EAL2 security target for Juniper Networks NetScreen-IDP 4.x.

The NSM system consists of the Device Server and the GUI Server; the NSM User Interface is a client application used to access information stored in the NSM system.

Guidance for Intended Usage

The NSM system must be installed on dedicated systems. These dedicated systems must not contain user processes that are not required to operate the NSM software.

Guidance for Personnel

The following items are also required for Common Criteria EAL2 compliance:

- There must be one or more competent individuals assigned to manage the NSM system and User Interface, and the security of the information that they contain.
- The authorized administrators must not be careless, willfully negligent, or hostile and must follow and abide by the instructions provided by the NSM documentation.
- The NSM system and User Interface must be accessed only by authorized users.

Guidance for Physical Protection

The processing resources of the NSM system and User Interface must be located within facilities with controlled access which prevents unauthorized physical access.

Supported IDP-Capable Devices

NSM supports IDP on standalone IDP Series Intrusion Detection and Prevention Appliances (IDP 10, 50, 100, 200, 500, 600C, 600F, 1000, 1100C, and 1100F); as part of ISG2000 and ISG1000 security systems running ScreenOS 5.0.0-IDP1 or ScreenOS 5.4 and later; as well as J Series, SRX Series, and MX Series devices.

Standalone IDP Sensors

The ISG2000 and ISG1000 security module is an optional component that provides IDP functionality. If you have an ISG2000 or ISG1000 device that does not have IDP capability, you can upgrade the device to be an IDP-capable system by replacing the memory chip in the CPU, installing up to three security modules, and installing the Advanced and IDP license keys for IDP. See the *ISG2000 Field Upgrade Guide* or the *ISG1000 IDP Field Upgrade Guide* for instructions on how to upgrade your device to include IDP capabilities.

You can use the ISG2000 or ISG1000 device with IDP capability as a fully integrated firewall/VPN/IDP security system that not only screens traffic between the Internet and your private network, but also provides application-level security. Alternatively, you can use the ISG2000 or ISG1000 device as a dedicated IDP system to protect critical segments of your private network, such as Web servers or corporate accounting servers.



NOTE: IDP Series Appliances are standalone appliances that provide IDP functionality without integrated firewall/VPN capabilities.

NSM is the sole means for configuring and managing IDP on the ISG2000 and ISG1000 devices. Although you can use the ScreenOS CLI or Web UI to configure the firewall/VPN capabilities of the security device, you must use the NSM UI to enable and configure IDP capabilities on the security module.

Enabling Jumbo Frames (ISG1000 Only)

NSM supports jumbo frames on ISG1000 devices running ScreenOS 6.0r2 and later. When the jumbo frame feature is enabled, the four predefined ports on the ISG1000 are disabled.

If, however, you use a template to apply a predefined Ethernet or Fast Ethernet port to an ISG1000 device with jumbo frames enabled, the port will be visible and can be edited, even though it is not applicable to the device type.

To enable jumbo frames for the ISG1000:

1. In the Configure panel of the NSM main navigation tree, select **Device Manager > Devices**.
2. Click the Add Device icon, and then select **Device** from the list. The New Security Device wizard appears.
3. Enter a device name, and then select the **Model Device** option button.
4. Click **Next** to continue.
5. Select **ScreenOS/IDP** from the OS Name list.
6. Select **nsISG1000** from the Platform list.
7. Select **6.0** or greater from the Managed OS Version list.
8. Select the **Enable Jumbo Frame** check box, and then click **Finish**.

Enabling IDP Functionality

To enable IDP functionality on a managed device and deploy that functionality to protect your network, you must perform the steps described in the following sections:



NOTE: Juniper Networks devices require a license to activate the feature. To understand more about Network and Security Manager (NSM) Licenses, see, [Licenses for Network Management](#). Please refer to the Licensing Guide for general information about License Management.

- [Adding an ISG2000/ISG1000 Security Device with a Security Module on page 49](#)
- [Updating Attack Objects on page 49](#)
- [Adding Objects \(Optional\) on page 50](#)
- [Configuring a Security Policy for IDP on page 50](#)
- [Reviewing IDP Logs on page 55](#)

Adding an ISG2000/ISG1000 Security Device with a Security Module

You must add an ISG2000 or ISG1000 security device with at least one security module to the NSM UI before you can enable the IDP functionality in the security module.

NSM automatically detects the security module when you:

- Import an ISG2000 or ISG1000 device running ScreenOS 5.0.0-IDP1 and the security module is already installed.
- Install a security module in an existing ISG2000 or ISG1000 device that is currently managed by NSM, then upgrade the device firmware to ScreenOS 5.0.0-IDP1.



NOTE: After you have upgraded the firmware, you must reimport the device configuration.

To view the security module in the UI, open the device configuration and select **Network > Chassis**.

Updating Attack Objects

You must update the attack object database before you can use IDP functionality. To update the IDP and DI databases and the IDP detector engine, download new attack objects from the attack object database server to the GUI Server.



NOTE: You must have DNS enabled on the NSM GUI server before you can update your attack objects.

To update the IDP and DI attack object databases on the NSM GUI Server:

1. Select **Tools > View/Update NSM Attack Database** to open the Attack Update Manager wizard,
2. Follow the instructions in the Attack Update Manager wizard to download the new Signature and Protocol Anomaly attack objects to the NSM GUI Server. The management system contacts the server and downloads the latest database version to the GUI Server.

After you have updated the attack object database on the GUI Server, you can use that database to update the attack object database on your managed devices.

IDP attack objects are loaded onto IDP-capable devices with the IDP rulebase.

To load a new detector engine onto an IDP-capable device:

1. From the Device Manager launch pad, select **Security Updates > Update ScreenOS Device Detector** or **Update Junos Device Detector**.
2. Click **Next**, then select the devices on which you want to load the detector engine.
3. Click **Finish**.

To download the DI attack object database update to your DI-capable devices:

1. From the Device Manager launch pad, select **Update Device Attack Database** to open the Change Device Sigpack wizard.
2. Follow the directions in the Change Device Sigpack wizard to update the attack object database on the selected managed devices.

Adding Objects (Optional)

Create address objects for the network components you want to protect with IDP. These components can be routers, servers, workstations, subnetworks, or any other object connected to your network. You can also create address object groups, which represent multiple address objects. (If you have previously created network objects for use with your devices, you do not need to create them again.)

For more information about creating address objects, see [“Configuring Address Objects” on page 351](#).

For more information about adding address object for standalone IDP sensors, see the *IDP Concepts & Examples Guide*.

Configuring a Security Policy for IDP

Because the security module on the device processes traffic *after* the firewall/VPN management module, you must configure a firewall rule to pass permitted traffic to the IDP rulebases. Enabling IDP functionality in a security policy is a two-step process: first enable a firewall rule to pass permitted traffic to the IDP rulebases, then create the IDP rules that detect and prevent malicious traffic from entering your network.

When creating a new security policy for your IDP deployment, we highly recommend you use a security policy template. Each security policy template contains the IDP rulebase

and IDP rules that use the default actions associated with the attack object severity and protocol groups. You can customize these rules to work on your network as needed, such as selecting your own address objects as the Destination IP and choosing IDP actions and notifications that reflect your security needs.

If you do not use a security policy template, you must add the IDP rulebase manually, as detailed in [“Adding the IDP Rulebases” on page 52](#).

Configure Firewall Rules (ISG Only)

You can enable IDP within an existing rule, or create a new rule. Configure the firewall rule as you would normally, setting the source and destination zones, address objects, services, and so on to define the type of network traffic you want to permit.

When configuring the firewall rule, consider the following:

- Traffic that is denied by a firewall rule cannot be passed to IDP rules. To enable IDP in a firewall rule, the action must be permitted.
- When deploying the ISG2000 or ISG1000 device as a dedicated IDP system, configure a single firewall rule that directs all traffic to the IDP rules. (By default, the firewall denies all traffic.)



NOTE: When operating the security device in a nontransparent mode, you must have configured basic security device settings, such as assigning interfaces to zones, setting the administrative password, and configuring default routes. For details about configuring these settings, see the user guide that shipped with the device.

When operating the security device in transparent mode and using it as a dedicated IDP system, you do not need to configure additional firewall settings.

- For firewall rules that pass traffic to the IDP rulebases, the Install On column must include IDP-capable devices only.

Setting the IDP Mode (ISG Only)

Because the security module is part of the inline security device, IDP protects your network while directly in the path of traffic coming and going on your network.

To set the IDP mode:

1. In the Configure panel of the main navigation tree, select **Policy Manager > Security Policies**, and then double-click the policy name in the Security Policies window to open the firewall rulebase.
2. In the Rule Options column of a firewall rule, select **IDP**.
3. Select one of the following modes:
 - **Inline**—In the inline mode, IDP is directly in the path of traffic on your network and can detect and block attacks. For example, you can deploy the ISG2000 or ISG1000

with integrated firewall/VPN/IDP capabilities between the Internet and the enterprise LAN, WAN, or special zones such as DMZ.

- **Inline Tap**—In the inline tap mode, IDP can detect attacks and provide notification. IDP receives a copy of a packet while the original packet is forwarded on the network. IDP examines the copy of the packet and flags any potential problems. IDP's inspection of packets does not affect the forwarding of the packet on the network.



NOTE: You must deploy the ISG2000 or ISG1000 device inline. You cannot connect a device that is in the inline tap mode to an external TAP or SPAN port on a switch.

Selecting either mode enables IDP for the firewall rule, and configures the security device to forward all permitted traffic to the IDP rulebases for further processing.

Adding the IDP Rulebases

After you have enabled one or more firewall rules to pass traffic to the IDP rulebases, you must add one or more of the following IDP rulebases to the security policy:

- **The IDP Rulebase**—This is the main rulebase for IDP rules. Add this rulebase when you want to configure rules that use attack objects to detect specific malicious or anomalous activity in your network traffic.

For an overview of creating rules in the IDP rulebase, see [“Configuring a Security Policy for IDP” on page 50](#). For details, see [“Configuring IDP Rules” on page 512](#).

- **The Exempt Rulebase**—This rulebase works in conjunction with the IDP rulebase. When traffic matches a rule in the IDP rulebase, the security module attempts to match the traffic against the Exempt rulebase before performing the specified action or creating a log record for the event.

Add the Exempt rulebase:

- When an IDP rule uses attack object groups containing one or more attack objects that produce false positives or irrelevant log records.
- To exclude a specific source, destination, or source and destination pair from matching an IDP rule (prevents unnecessary alarms).
- When the IDP rulebase uses static or dynamic attack object groups that contain one or more attack objects that produce false positives or irrelevant log records.

For details on creating rules in the Exempt Rulebase, see [“Configuring Exempt Rules” on page 535](#).

- **The Backdoor Detection Rulebase**—This rulebase detects backdoor traffic from components on your internal network. A *backdoor* is a mechanism installed on a host computer that facilitates unauthorized access to the system. Attackers who have already compromised a system often install a backdoor to make future attacks easier. However, when attackers enter commands to control a backdoor, they generate interactive traffic that your security device can detect.

Add this rulebase to your security policy when you want to configure rules that detect backdoor activity on your internal network. For details on configuring rules in the Backdoor Detection Rulebase, see [“Configuring Backdoor Rules” on page 538](#).



NOTE: NSM does not import IDP rulebases in a security policy when importing the device configuration from an existing IDP-capable security device.

If you are using a security policy template, the IDP rulebases are automatically added to the policy. However, if you are not using a template, you must manually add the IDP rulebases to your policy.

To add the IDP, Exempt, or Backdoor Detection rulebase:

1. In the Configure panel of the main navigation tree, select **Policy Manager > Security Policies**, then double-click the policy name in the Security Policies window.
2. Click the **Add** icon in the upper right corner of the Security Policy window and select **Add Backdoor Rulebase** to open the selected rulebase tab.

Configure IDP Rules

IDP detection and prevention capabilities work against attacks by dropping connections during the attack detection process, preventing attacks from reaching the target system.

To add a rule to a rulebase:

1. Click the rulebase tab for the rulebase in which you want to add a rule.
2. On the left side of the Security Policy window, click the **Add** icon to open a default rule.

For rules in the IDP rulebase, you define the type of network traffic to monitor, the attacks to detect, the action to be taken against matching traffic, and the notification you want to receive. Specifically, you must configure the following:

- **Configure Match Criteria**—Define the type of network traffic you want the IDP security module to monitor for attacks, such as source-destination zones, source-destination address objects, and the application layer protocols (services) supported by the destination address object. You can also negate zones, address objects, or services.

You configure the match criteria in the following IDP rulebase columns:

- From Zone
- Source
- To Zone
- Destination
- Service

For details on configuring match criteria within the IDP rulebase, see [“Defining Match For IDP Rules” on page 513](#).

- Add attack objects—Add the attacks you want the IDP security module to match in the monitored network traffic. Each attack is defined as an attack object, which represents a known pattern of attack. Whenever this known pattern of attack is encountered in the monitored network traffic, the attack object is matched. You can add attack objects by groups (category, operating system, severity, and so on) or individually.

You configure attack objects within the Attack column of the IDP rule:

- For details on selecting attacks within IDP rules, see [“Configuring Attack Objects in IDP Rules” on page 519](#)
- For details on IDP attack objects, see [“Working with IDP Attack Objects” on page 369](#).
- For details on creating your own custom IDP attack objects, see [“Configuring Custom DI and IDP Attack Objects” on page 371](#).
- Configure Action— Define the action the IDP security module takes when a particular attack is detected. You can define an IDP action (action that the security module takes against the current connection) and an IP action (action that the security module takes against the current and future connections to or from the same IP address).

You configure IDP actions in the Action column of an IDP rule. For details, see [“Defining Actions For IDP Rules” on page 517](#).

You configure IP actions in the IP Action column of an IDP rule. For details, see [“Configuring IP Actions in IDP Rules” on page 521](#). (The IP Action column appears only when viewing the security policy in Expanded Mode. To change the view mode of a policy, from the menu bar, select **View > Show Expanded Mode**, **View > Show Compact Mode**, or **View > Show Custom Mode**).

- Configure Notification—Define the logging and notification activities you want the IDP security module to take when the IDP rule is matched. You can configure the module to generate log entries, trigger alarms, and log captured packets.



NOTE: J Series routers and SRX Series gateways do not send packet data to NSM. If your policy rules attempt to do so, then no data is logged.

Configure Notification settings in the Notification column of an IDP rule. For details, see [“Configuring Notification in IDP Rules” on page 523](#).

Assign, Validate, and Install the Security Policy

After you have created the necessary firewall and IDP rules within the security policy, you must perform the following steps to apply the policy to your network traffic:

1. Assign the policy to a device.

Assigning a policy to a device links the device to that policy.

To assign an existing policy to the ISG2000 or ISG1000 device:

- a. In Device Manager, right-click the ISG2000 or ISG1000 device and select **Policy > Assign Policy**.
 - b. From the Security Policy Name list, select the security policy you just created.
2. Validate the security policy (optional).

Validating a security policy can identify potential problems before you install it.

 - a. In the navigation tree, select **Device Manager**.
 - b. From the Device Manager launchpad, select **Validate > Validate IDP Policy** and select the device. A Job Manager window displays job information and progress.

If NSM identifies a problem in the policy during policy validation, it displays information about the problem at the bottom of the selected rulebase. For example, if you included a non-IDP capable security device in the Install On column of an IDP rule, policy validation displays a error message.
3. Install the security policy.

During policy installation, NSM installs the entire security policy, including the firewall and IDP rules, on the security devices you selected in the Install On column of each rule.

To install a policy:

 - a. In the navigation tree, select **Device Manager**.
 - b. From the Device Manager launchpad, select **Update Device**.
 - c. Select the ISG2000 or ISG1000 security device.
 - d. Click **OK**. A Job Manager window displays job information and progress.

Reviewing IDP Logs

After you have enabled IDP on the device and installed a security policy that uses the IDP detection and prevention functionality, IDP logs begin to appear in the NSM Log Viewer (assuming you enabled IDP logging for each IDP rule). Depending on the attack objects you included in the IDP rule, the IDP log entries you receive might provide details of events such as attacks against your network, protocol anomalies, or even simple login attempts.

To view IDP log entries:

1. Go to the main navigation tree and expand the **Investigate** panel.
2. Select **Log Viewer > Predefined > 3-IDP/DI**. The Log Viewer displays all IDP logs generated by the security device.



NOTE: The DI/IDP Logs view is a predefined custom view applied to all log entries received by NSM. To view all log entries for all devices in the selected domain without filters, select the Log Viewer module in the main navigation tree.

We recommend you review and analyze these log entries to determine the effectiveness of your current security policy and IDP rules. Log entries are often a valuable insight into your network traffic. You can see where traffic is coming from, where traffic is going to, and what malicious content (if any) the traffic contains.

Maintaining IDP

Attackers are constantly devising new and better ways to infiltrate your network. Juniper Networks actively discovers these new attacks and creates new attack objects to detect them—so you can prevent the attacks from entering your network. To ensure that the IDP security module and security policies remain highly effective against all emerging and evolving threats, we highly recommend that you perform frequent updates to the attack object database and to the IDP detection engine, described in [“Managing the Attack Object Database” on page 306](#).

Creating IDP-Only Administrators

You can use NSM's role-based administration (RBA) to create a custom role for administrators working with IDP functionality on a device. For example, if your organization's IDS or IDP administrators do not configure firewall/VPN security devices, you can restrict administrative privileges for those administrators within the NSM system to IDP tasks only.



NOTE: The NSM “super” administrator automatically has all IDP-related permissions.

A custom role for IDP administrators might include the following permissions:

- Attack Update
- Create/View/Edit/Delete Policies
- Create/View/Edit/Delete Backdoor and IDP Rulebases
- View Firewall Rulebases
- Create/Edit/Delete Shared Objects and Groups

For details on RBA in NSM, see [“Configuring Role-Based Administration” on page 68](#); for an example that shows how to create an IDP-only administrator, see [“Creating Administrators” on page 69](#).

Simplifying Management

When you add devices to NSM, you are creating the network organization that you use to manage your security system. Before you begin the device creation or device import process however, first review your network topology and decide how you want it to appear in NSM. This is particularly important when you are creating a new network, but is also helpful when you are importing networks, because you might want to edit your network design to take advantage of key NSM management features.

These features include:

- [Using Device Groups on page 57](#)
- [Using Device Templates on page 57](#)
- [Using Configuration Groups on page 58](#)
- [Merging Policies on page 58](#)
- [Using a Naming Convention on page 58](#)

Using Device Groups

You can create groups of devices to manage multiple devices at one time. Group your device by region, device type, or even OS version, and then use the groups to:

- Deploy new or updated device configurations to the entire device group.
- Deploy new or updated policies to the entire device group.
- Create reports using the log information from the entire device group.

Using Device Templates

A template is a predefined device configuration that helps you reuse common information. A domain can contain multiple templates, and you can use templates to quickly configure and deploy multiple devices. A device template looks much like a device configuration—the template page displays boxes for interfaces, zones, and virtual routers in which you can enter values. When you add a new device that uses similar information as a previously added device, you can use a device template to fill in specific configuration values so you do not have to reenter information.

For example, you might create a generic NetScreen-5GT device template that you can use each time you add a device of that type. Or you can apply multiple templates to the same device. You can map a maximum of 63 templates to the same device; you set the priority of the template to determine the order in which they applied.

For example, you might create the following templates:

- DNS setting template
- Default PKI Settings template
- Authentication template

Apply these templates to a single device to instantly configure the DNS, PKI, and Authentication settings for the device.



NOTE: You cannot create VPNs between devices in different domains.

For details about device templates, see [“Using Device Templates” on page 210](#).

Using Configuration Groups

Configuration groups are similar to device templates in that you define configuration data to be used multiple times. Configuration groups, which are used only in Junos devices, are different in that the configuration data is used within the same device but at several levels in the configuration. A special use of configuration group is to apply configuration data in different members of a cluster.

For details about configuration groups, see [“Using Configuration Groups” on page 235](#).

Merging Policies

You can create new policies for all your managed devices from the central NSM UI and deploy them with a single click. Alternatively, NSM can import all existing policies from your device. You can import all security and access policies from your devices, and import all VPN tunnels (route-based and policy-based) from your devices.

Each time you import a policy from a managed device, that policy appears in NSM as a separate, individual policy in the Security Policies list. To simplify policy management and maintenance, you can merge two policies into a single policy. For details on merging policies, see [“Configuring Security Policies” on page 473](#).

Using a Naming Convention

A naming convention is a method for assigning names to your network devices (firewalls, servers, workstations, and so on) that enables you to quickly identify where the device is and what its purpose is.

If your network is small, you might choose a simple naming convention, such as planet names, car models, or mountain names. When using this type of informal method to name your network components, be sure to choose a theme that is easily understood by your users and administrators, and that still has room to grow. For example, you might use the naming convention. <city><name>, with a naming theme of Greek mythology figures; some sample device names might be la_ns5gt_Athena, sf_ns5XT_Zeus, or oak_ns204_Hermes.

If your network is larger, however, you need a more formal naming schema that is more descriptive of the network component’s location and purpose. Having a logical and standardized naming convention can help you quickly identify the appropriate administrator for the component, as well as quickly identify the component location without having to review subnet tables.

A typical naming convention for large, distributed networks consists of a standardized location identification code, followed by the department code, a description of function, and a numerical sequence.

Example: Using a Naming Convention for Devices

You use the naming convention: nation_state_platform_name for your security devices. Your devices use names similar to the following:

- us_ca_ns5gt_01

- us_co_ns204_05
- us_tx_ns5200_10

Example: Using a Naming Convention for Address Objects

For address objects that represent networks or hosts, use the following naming convention. state_function_service_00:

- State—A two-character postal abbreviation for the state where the server resides.
- Function—Some common functional abbreviations:
 - SV (Server)
 - WS (Workstation)
 - IIS (Web Server)
 - MSX (Mail Server)
 - SQL (SQL Server)
 - SMS (SMS Server)
 - APP (Application Server)
- Service—Abbreviated name of the main service on that machine
- Number—A sequential number starting with 01

For example, the first Apache Web server installed in the state of California would be: ca_ws_apache_01.

For address objects that represent client hosts, use the naming convention: state_firstname_(m or w)_os

- State—A two-character postal abbreviation for the state where the user is located
- FLastname—The first initial and last name of the main user (or general account name if it is a multiuser machine)
- (M or W)—A single letter to designate Mobile computer or Workstation
- OS—A two-character abbreviation for the operating system

For example, Wendy Parker, working in Texas on a Windows 2000 Pro laptop, would see her machine name as: tx_wparker_m_2kpro.

Creating an Information Banner

Central Manager administrators and regional server “super” administrators have the ability to display an informational banner when users log into NSM. This banner is created using custom text that is stored on the server. Once it has been created, the banner is displayed as a splash screen after users enter their login credentials. This text is used server-wide, which does not depend on user, role, domain, and so on. Users are unable

to proceed into the NSM UI until they accept the message to continue. If this banner is used, users are required to accept the message each time they log in.

You can add an information banner from Central Manager or from a regional server.

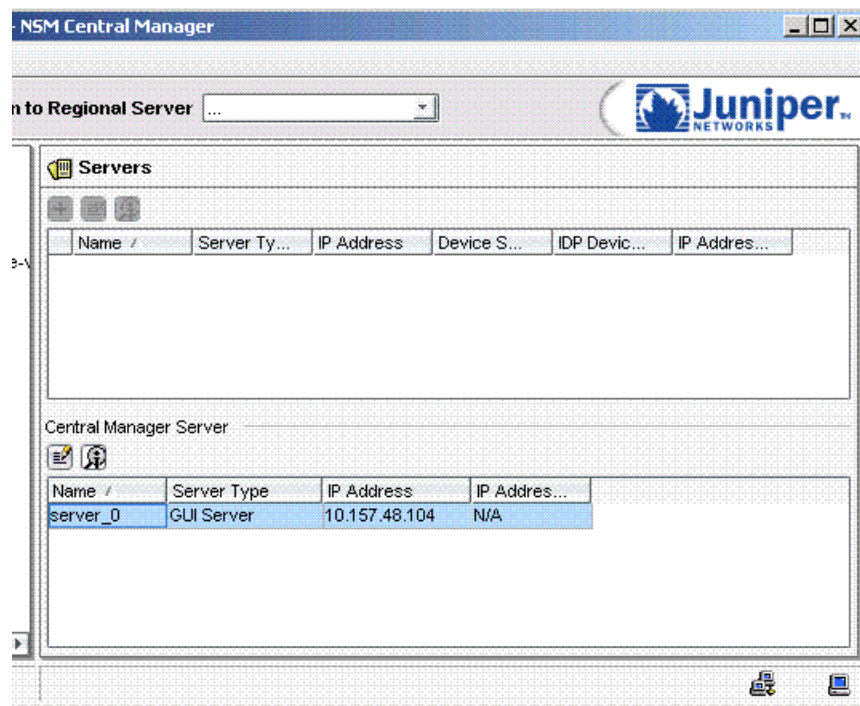
Adding an Information Banner

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client or a super user is logged into a regional server.

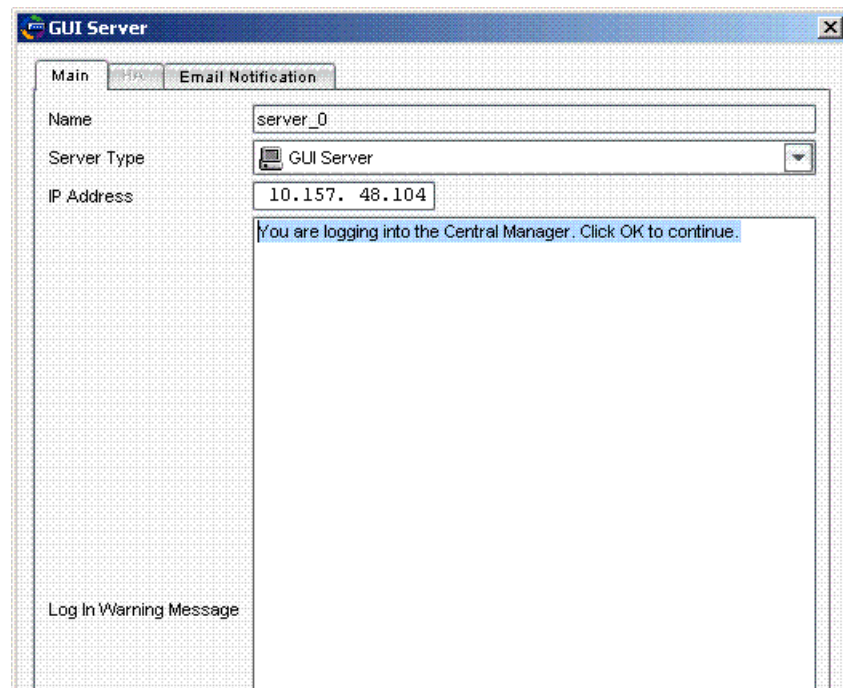
To add an informational banner:

1. In the Administer panel of the main navigation tree, select **Server Manager > Servers**.
2. Select the GUI server to which you want to add the banner server-wide, and click the Edit icon, as shown in [Figure 14 on page 60](#).

Figure 14: Selecting the GUI Server in Central Manager



3. Enter the customized text in the Log In Warning Message text box, and then click **OK**, as shown in [Figure 15 on page 61](#).

Figure 15: Setting Up an Information Banner

The message is immediately available to NSM users connected to the server, as shown in Figure 16 on page 61.

Figure 16: Information Banner Login into Central Manager

The NSM user must click **Yes** to access the GUI server.

Modifying an Information Banner

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client or a super user is logged into a regional server and an information banner has been created.

To modify an informational banner:

1. In the Administer panel of the main navigation tree, select **Server Manager > Servers**.
2. Double-click the GUI server to which you want to change the banner server-wide.
3. Change the customized text in the Log In Warning Message text box, and then click **OK**.

The message is immediately available to all NSM users server-wide.

Deleting an Information Banner

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client or a super user is logged into a regional server and an information banner has been created.

To delete an informational banner:

1. In the Administer panel of the main navigation tree, select **Server Manager > Servers**.
2. Double-click the GUI server for which you want to delete the banner server-wide.
3. Delete the customized text in the Log In Warning Message text box, and then click **OK**.

The message is immediately removed from the login screen to all NSM users server-wide.

CHAPTER 3

Configuring Role-Based Administration

This chapter details how to use the Juniper Networks Network and Security Manager (NSM) role-based administration (RBA) feature to configure domains, administrators, and roles to manage your network. Your organization probably already has an existing permission structure that is defined by job titles, responsibilities, and geographical access to your security devices. Using role-based administration, you can re-create the same permission structure in NSM.

RBA is particularly useful for Enterprise and Service Provider organizations that have different administrative roles associated with managing a large network and security infrastructure. You can define custom roles with specific permissions to create the exact administration structure your organization requires.

After you have created an RBA-based structure for your network, you can begin thinking about your central management strategy and how to prepare your network for NSM. NSM includes many features specifically designed for managing multiple Juniper Networks devices, such as device groups and templates.

This chapter contains the following sections:

- [Role-Based Administration on page 63](#)
- [Using Role-Based Administration Effectively on page 65](#)
- [Configuring Role-Based Administration on page 68](#)

Role-Based Administration

The NSM role-based administration (RBA) feature enables you to define strategic roles for your administrators, delegate management tasks, and enhance existing permission structures using task-based functions.

Use NSM to create a secure environment that reflects your current administrator roles and responsibilities. By specifying the exact tasks your NSM administrators can perform within a domain, you minimize the probability of errors and security violations, and enable a clear audit trail for every management event.

Domains

A domain is a logical grouping of devices, their security policies, and their access privileges. A domain can contain devices, templates, objects, policies, VPNs, administrators, activities,

authentication servers, groups—a representation of all or a subset of the physical devices and functionality on your network.

NSM contains a default top-level domain, called the global domain, which can contain additional domains, called subdomains. Use subdomains to manage multiple domains in a single hierarchical structure. You can create all your devices and their configurations in the global domain, or you can configure additional subdomains within the global domain.



NOTE: You can create only one level of subdomains in NSM.

Typically, multiple domains are used for two main reasons: to define network structure and to control administrator access. Multiple domains help to separate large, geographically distant systems into smaller, more manageable sections, and also to control administrative access to individual systems.

For example, a small organization might only have one domain (the global domain) for their entire network, while a large, international organization might have dozens of subdomains that exist within the global domain to represent each of its regional office networks across the world. A service provider might use domains to build a virtual network for each client network, and then assign access permissions for each client domain.

Domain selection is important if you plan to use VPNs in your network. Because you can create VPNs only between devices in the same domain, be sure to add the devices you want to connect with a VPN to the same domain.

About Roles

Roles define who can perform which task and view which information. NSM uses a powerful, role-based access control system that enables you to create custom roles for individual administrators. Use role-based management to control administrative access to NSM functionality.

All NSM users are some type of administrator. During NSM installation, you are prompted for a password for the (default) administrator account for NSM; this administrator account is the first administrator, and is therefore the super administrator. The super administrator automatically has all permissions, and can create other domains, administrators, and roles. As super administrator, you specify who has what permissions for NSM functionality for the entire NSM system, a single domain, or specific functionality within a domain.



NOTE: All passwords handled by NSM are case-sensitive.

System administrators can be active or read-only. All system administrators, including those assigned a Read-Only role, can create and run their own reports.

You can define multiple NSM administrators and assign dedicated roles to each administrator:

- A role is a set of activities that specify the functions the administrator can perform.

- Activities are predefined tasks within NSM. The NSM administrator can combine multiple activities into a custom role.



NOTE: You cannot define a custom activity.

With role-based administration, you can specify who has what permissions for NSM functionality for the entire NSM system, a single domain, or even specific functionality within a domain. You can even delegate NSM administrator management, enabling existing NSM administrators to create other NSM administrators, assign domains, and define or create roles.



NOTE: A device administrator is the person responsible for managing a device directly, using the command line or GUI for the local OS (ScreenOS, Junos, IC, or SA). If a device administrator uses only the local OS command line or GUI to manage devices, do not create an NSM administrator account for the device administrator; however, if a device administrator uses both the local OS and NSM to manage devices, you must create an NSM administrator account for the device administrator.

Using Role-Based Administration Effectively

The structure of your NSM domains should reflect both your existing network structure and your desired permission structure.

- Network Structure—Use multiple domains to segregate large, geographically distant networks into locally managed sections.
- Permission Structure—Use multiple domains to segregate critical devices and systems from less important network areas, and then restrict administrator access to devices in the critical domain.

Your organization probably already has an existing permission structure that is defined by job titles, responsibilities, and geographical access to your security devices. You can re-create this same permission structure in NSM.

Role-based administration is particularly useful for Enterprise and Service Provider organizations that have different administrative roles associated with managing a large network and security infrastructure. RBA is also helpful for any size of organization that wants to provide access to other device statistics to non-administrators within the organization, such as creating a role for the CIO to access reports.

Enterprise Organizations

Each enterprise defines administrative roles differently. With NSM, you have the flexibility to create the appropriate permission level.

Geographical Divisions

To manage large, geographically diverse networks, you can create domains for each separate geographical location. Typically, the larger the Enterprise, the deeper and more complex your geographical divisions. Two common geographical divisions are defined below:

- **Corporate**—The corporate domain is the global domain. In the global domain, the super administrator creates the devices, objects, and policies that exist in the corporate network, and creates subdomains for each region.
- **Region**—Each region is a subdomain. Within each subdomain, the super administrator creates a regional administrator to manage the subdomain. The super administrator also specifies the roles the regional administrator has to view and manipulate devices, remote users, configuration actions, and report information within that subdomain.

NOC and SOC

To ensure continual network uptime and provide prompt response to network attacks, each geographical division is often monitored by a dedicated network operations center (NOC), a security operations center (SOC), or both. The NOC and SOC are typically the same location for small organizations, but might be physically separate for larger, more complex organizations. Whether combined or separate, NOC and SOC administrators perform distinct roles:

- NOC administrators focus on network connectivity and status.
- SOC administrators focus on network attacks and events associated with security policies.

Administrator Types

Many organizations have different types of administrators for different roles within the company. Each organization has a unique vision for the granularity of their permission structure.

Tiered NOC/SOC

Typically, a NOC/SOC uses a three-tier permission structure. The administrators in each tier have a specific level of skill and understanding of the underlying network and technology, as well as access permissions to view or change configurations. An example NOC/SOC center might use the following role structure:

- Tier 1 administrators view events and audit configurations.
- Tier 2 administrators view events and audit configurations, but also change network configurations during troubleshooting.
- Tier 3 administrators have full access to all functionality on the device, and make configuration and policy changes.

Configuration Responsibilities

Some enterprise organizations use different administrator groups to manage specific aspects of device configuration. Configuration responsibilities might use the following role structure:

- **IT group**—Integrates new devices into the existing network infrastructure. This group has roles with activities for setting up Layer 2 and Layer 3 aspects of the device (IP addressing, Routing, VLANs, Syslog, and so on). Within the IT group, the network administrator might also have a role with an activity for managing the management system.
- **Security group**—Creates and manages security policies. This group has roles with activities for defining custom services, address objects, and firewall rules on devices for which they have responsibility.
- **Remote Connectivity group**—Creates and manages VPNs and RAS user configuration. This group has roles with activities for configuring VPNs and remote users.

Specific Tasks

- **Configuration Validation**—An audit administrator approves all configuration changes before those changes are made on the network. Only the auditor has a role with activities for updating devices on the network.
- **Reporting**—A reporting administrator views reports for one or more domains. A regional reporting administrator has a role with activities for viewing reports for their regional subdomain; a corporate reporting administrator has a role with activities for viewing reports for the global domain and all subdomains.
- **Configuration Update**—An update administrator updates firmware for devices. The update administrator has a role with activities for updating firmware on the devices in their assigned domain.
- **Administrative Management**—A management administrator creates administrators and manages their permissions. The super administrator creates a management administrator to delegate administrator management. For example, a NOC Tier 2 administrator has a role that includes the activity to create new administrators, but cannot assign them an activity that is not included in their own role. Typically, a subdomain has only one management administrator to control the creation of administrators.
- **Device Installation**—A device install administrator creates new devices. The device install administrator has a role with activities for adding, updating, and viewing device configurations.

Service Providers

Service Providers can use NSM domain, subdomains, and roles to manage their internal infrastructure and their customers' infrastructures.

Internal Network

Internally, a Service Provider network is similar to an enterprise network; both view their networks as regions with dedicated NOC/SOC, and both use the same types of administrators.

Managed Security Service Provider (MSSP)

Telcos and Service Providers use their networks to generate revenue. Customers pay the MSSP to deploy devices and to manage the VPN or firewall infrastructure. MSSPs use different role structures that best match their organizational structure:

- MSSP owns devices; customer manages infrastructure.
- Customer owns devices; MSSP manages infrastructure.
- Customer leases devices; MSSP manages the infrastructure.
- MSSP owns devices and manages infrastructure (Customer Network Management (CNM)).

CNM Service Providers vary widely in how they control access to their customer networks. Some CNMs assign one or more customers to a network administrator that has control over the device and policies used by those customers. Other CNMs assign one network administrator to view reports for all customers. CNMs might use the following role structure:

- Super administrator. At the global domain, the super administrator creates
 - The internal network of the CNM.
 - A subdomain for each customer. The customer subdomain contains the devices and objects that belong to the customer network. Because the customer network is completely contained within a subdomain, it is isolated from other subdomains for other customers.
 - Customer administrators to manage one or more subdomains. The super administrator assigns roles to the customer administrator in one or more customer subdomains, enabling the customer administrator to handle multiple customer networks without access to the CNM internal network.

Additionally, the super administrator can create a role structure that maps to the specific tasks performed by each customer administrator, as described in [“Specific Tasks” on page 67](#).

MSSPs can also use virtual systems (available on NetScreen-500 and NetScreen-5000 series) to share a single device between multiple customers. For each customer, the MSSP creates a customer subdomain and a virtual system within that subdomain.

Configuring Role-Based Administration

When you have analyzed your network and permission structure and designed your domain strategy, you are ready to create subdomains and new NSM administrators for

those subdomains. When you create NSM administrators for your subdomains, you can set their permissions so that they can see only the domains to which they have access.

From the menu bar, click **Tools > Manage Administrators and Domains** to display the RBA settings for NSM:

- Administrators—Configure administrators for NSM or IDP.
- Roles—View or edit default roles, or create your own custom roles for your NSM or IDP administrators.
- Subdomains—Create subdomains to segregate networks.
- Current Domain Detail—View the information about the current domain, such as assigned administrators, authentication method, and default authentication servers.

The following sections explain how to configure these RBA settings.

Creating Administrators

The super administrator automatically has full permissions for all subdomains, so you do not need to assign new subdomains to the super administrator. However, to assign a subdomain to another administrator, you must first create the administrator and specify their permissions within a selected subdomain.

You can create NSM administrators at the global domain or subdomain level:

- To assign the new administrator permissions in the global domain or multiple subdomains, create the administrator in the global domain.
- To assign the new administrator permissions in only one subdomain, create the administrator in that subdomain.

Configuring General Settings

To create an NSM administrator account, click the Add icon in the Administrator tab to display the New Admin dialog box. In the General tab, enter a name and contact information (e-mail, telephone, and other basic information) for the new administrator.



NOTE: The following characters are not supported for NSM administrator names:

- Period (.)
- Number sign (#)
- Dollar sign (\$)
- Asterisk (*)
- Ampersand (&)
- Circumflex (^)

Configuring Authorization

To configure the authorization method for the new administrator, click the Authorization tab and select local or remote authentication:

- For locally authenticated administrators, the NSM management server handles authentication. You must specify the password that NSM uses to authenticate the administrator; the administrator must enter this password at the NSM UI login screen.



NOTE: All NSM passwords are case-sensitive.

- For remotely authenticated administrators, a RADIUS authentication server handles authentication. Because the administrator password is stored on the RADIUS server, you do not need to enter the password again, however, the administrator must enter the password at the NSM UI login screen.

To configure the RADIUS authentication server for NSM administrators, see the *Network and Security Manager Online Help* topic “Editing the Domain Contact.”



NOTE: The super administrator has full permissions. You cannot change or delete permissions for the super administrator; you can only change the password. Because the super administrator has complete control over NSM functionality, we recommend that you consider the security of the super administrator password appropriately. If you forget or lose the super administrator password, please contact the Juniper Technical Assistance Center (JTAC).

RADIUS Authentication and Authorization

NSM supports both local and RADIUS user authentication. It manages access control both through the local database and through the RADIUS server.

You are not required to define RADIUS users in the local NSM database. The AUTH Handler looks at the local database to find the user, and then, if no match is found, to the RADIUS server. You can also define the role assignment for each user directly from the RADIUS server.



NOTE: You must configure your RADIUS server individually for each domain.

NSM also supports a secondary RADIUS server for administrator authentication and authorization when the primary RADIUS server cannot be contacted.

There are two kinds of users: local users and RADIUS users. The local user is created locally and authentication data is stored in the local database. The default authentication mode is local mode. The RADIUS user is created only on a RADIUS server and can only be authenticated using a remote RADIUS server.

There are also two kinds of authentication modes for NSM users: local mode and RADIUS mode. Both User and Domain can define these modes and Domain's authentication mode is applied to all the users within it. User's Authentication mode has a higher priority and can override Domain's mode.

The NSM user is authenticated based on the rules listed in [Table 15 on page 71](#).

Table 15: How to Authenticate Users

Rule	User in Local Database	User Auth Mode	Domain Auth Mode	Authentication Results	Authorization
1	Defined	Local	Local	Authenticates user locally.	Local
2	Defined	Local	Remote	Authenticates user locally first. If fails, RADIUS authentication is used.	Local
3	Defined	Remote	Local	Authenticates user remotely.	Local
4	Defined	Remote	Remote	Authenticates user remotely.	Local
5	Not Defined	—	Local	Authenticates user remotely.	Remote
6	Not Defined	—	Remote	Authenticates user remotely.	Remote

Dictionary File

To authenticate local or remote users from a RADIUS server, you must first define role mapping assignments and domain names in NSM. If you use Steel Belted RADIUS, you can copy the NSM RADIUS dictionary to your RADIUS server.

This file (**netscreen.dct**) is available in the NSM. If you installed NSM using the default options, you can find the dictionary in the following location:

```
/usr/netscreen/GuiSvr/utls/netscreen.dct
```

RADIUS VSA Definition

RADIUS vendor specific attribute (VSA) is available to allow vendors to support their own extended attributes. If you use a RADIUS server other than Steel-Belted RADIUS, you must enter the following NSM attributes in your RADIUS dictionary file.

These attributes are case sensitive and must be entered exactly as they appear below:

```
ATTRIBUTE NS-NSM-User-Domain-Name 26 [vid=3224 type1=220 len1=+2 data=string]
ATTRIBUTE NS-NSM-User-Role-Mapping 26 [vid=3224 type1=221 len1=+2 data=string]
```

When a user is defined only in RADIUS, you must define NS-NSM-User-Domain- Name and role mapping assignment. Auth Handler checks if the domain name matches the

user's login domain name when NSM authenticates the user. Role mapping lists are used for NSM access control purposes.

Custom Roles

Figure 17 on page 72 shows the format for a custom role. The format for the custom role of NS-NSM-User-Role-Mapping is:

domainName1:domainName2.roleName

- **domainName1** is the domain that the current user can access.
- **domainName2** is the domain that the current role (**roleName**) belongs to.

If you create a custom domain, NS-NSM-User-Domain-Name should include the domain's full path. Do not omit the word "global" and include the full path for **domainName**, for example, global.d1, global, or global.d2. Figure 17 on page 72 shows an example.

Figure 17: Creating Custom Domain

Attributes

☐ Use profile: View

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global.d1	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:global.d1.r1	<input type="checkbox"/>

Add... Edit... Delete

In Figure 17 on page 72, users belong to domain d1 and role r1 is defined in domain1. Therefore, the domain name is global.d1 and the role is global.d1:global.d1.r1.

Predefined Roles

The current predefined role names, which users can use, are listed below:

- Domain Administrator
- IDP Administrator
- Read-Only Domain Administrator
- Read-Only IDP Administrator
- Read-Only System Administrator
- System Administrator

Predefined roles do not belong to any domain. The format for predefined roles is:

DomainName1:(predefined-role-name)

- *DomainName1* is the domain that the current user can access.
- *predefined-role-name* is one of the options listed above.

For example, if a user is in domain d1 with a role of IDP Administrator, the domain name is global.d1 and the role is global.d1:IDP Administrator.

Creating Roles

If a user is defined in the local database or defined in a RADIUS server, NSM uses a role mapping list from the local database. The custom roles must be created in NSM. If the custom role belongs to a subdomain, it must be created in that subdomain. If the role is created in the global domain, it is automatically inherited into the subdomain and can be assigned to a subdomain user.



NOTE: A role defined in a subdomain belongs only to that subdomain.

Assigning Roles

If a user is defined in the local database, NSM uses a role mapping list from the local database. Otherwise, the RADIUS administrator must configure the role mapping list for each user on the RADIUS server.

Figure 18 on page 73 through Figure 24 on page 76 show examples of assigning predefined and custom roles through RADIUS. All examples assume that the user will be authenticated and authorized using a RADIUS server.

Figure 18: User in Domain "global" with a Predefined Role

User-specific:

Check list

Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global:Domain Administrator	<input type="checkbox"/>

<

>

Add...

Edit...

Delete

Figure 19: User in Domain "global" with Custom Role "r1"

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global:global.r1	<input type="checkbox"/>

< ||| >

Add... Edit... Delete

The "r1" role was created in the NSM in "global" domain.

Figure 20: User in Subdomain "d1" With a Predefined Role

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global.d1	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:IDP Administrator	<input type="checkbox"/>

< ||| >

Add... Edit... Delete

Figure 21: User in Subdomain "d1" With a Custom Role "r1"

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global.d1	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:global.d1.r1	<input type="checkbox"/>

< ||| >

Add... Edit... Delete

Create the custom role "r1" in the subdomain "d1."

Figure 22: Assigning Multiple Roles to a User in Global Domain

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global:global.r2	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global:global.r1	<input type="checkbox"/>

< ||| >

Add... Edit... Delete

Roles "r1" and "r2" are the custom roles assigned to the user.

Figure 23: Assigning Multiple Roles to a User in Subdomain

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global.d1	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:global.d1.r1	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:global.d1.r2	<input type="checkbox"/>

< ||| >

Add... Edit... Delete

Both "r1" and "r2" are the custom roles assigned to the user.

Figure 24: Assigning Roles Defined in Domain "global"

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global.d1	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:global.r1	<input type="checkbox"/>

Add... Edit... Delete

The user role "r1" is defined in global domain, but the user has access to only a subdomain d1 and therefore gets a the global role "r1."

Figure 25: Assigning Roles Defined in Domain "global" to Subdomain Only

User-specific:

Check list Return list

Attribute	Value	Echo
NS-NSM-User-Domain-Name	global	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d1:Domain Administrator	<input type="checkbox"/>
NS-NSM-User-Role-Mapping	global.d2:global.d2.r1	<input type="checkbox"/>

Add... Edit... Delete

The user is defined in domain "global" but has access to subdomains only. The user is a "Domain Administrator" in subdomain "d1," but has a custom role r1 for subdomain "d2."

Configuring Roles

To assign a role to the new administrator, select the **Permissions** tab and choose a role for the new administrator. When you assign a role to an NSM administrator, the administrator can perform the predefined system activities specified in that role.

You can select a default or custom role for that administrator. NSM includes default roles for common job responsibilities:

- Domain Administrator—Can perform all activities in the domain.
- Read-Only Domain Administrator—Can perform all read-only activities in the domain.
- IDP Administrator—Can perform all IDP activities. All other activities are excluded.
- Read-Only IDP Administrator—Can perform all read-only IDP activities.

- **System Administrator**—Can perform all system-wide activities, Domain Administrator activities, and IDP Administrator activities.
- **Read-Only System Administrator**—Can perform all read-only system-wide activities and Domain Administrator activities.

Each default role contains activities that relate to the traditional responsibilities for a specific job title. Use a default role to create quickly an NSM administrator or to create administrators when your organization's existing permission structure maps closely to the permissions defined in the default role.

All roles, default and custom, are created from activities. In a default role, the activities are chosen for you; in a custom role, you choose the activities that make up the desired functionality. See [“Creating Custom Roles” on page 77](#) for details.



NOTE: Role assignment is additive. When you assign multiple roles to a single administrator, the permissions specified by the activities in the role are added.

You must also select a domain. You can assign administrators to the global domain, or to one or more subdomains (the subdomain must already exist). Administrators must log in to the domain they were created in. For example, the super administrator has access to all domains, but must log in to the global domain first, and then switch to a subdomain using the domain menu. For details on creating a subdomain, see [“Creating Subdomains” on page 93](#).

Creating Custom Roles

For more complex and diverse permissions requirements, create custom roles to specify the exact level of permission you want to give an administrator. An *activity* is a predefined task that defines access to a function in NSM. To assign one or more activities to an NSM administrator, create a role that includes those activities and assign the role to the administrator.

Some activities are dependant on other activities. If you select a dependant activity, NSM automatically selects the prerequisite activities. You can clear prerequisite activities from a custom role, but doing so affects permissions granted in the dependant activity. For example, if you create a role that includes the activity “Create VPNs”, the activities “Edit VPNs” and “View VPNs” are automatically selected for you.

Click the Add icon to display the New Role dialog box and all available activities. NSM includes many predefined activities, grouped by similar functionality. See [Table 16 on page 77](#).

Table 16: Predefined NSM Administrator Activities

Function	Task	Description
Action Attributes	View	The Action Manager is a node on the main navigation tree that enables you to configure the management system to forward logs generated within a specific domain or subdomain.
	Modify	

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Access Profile Objects	View	Use access profile objects to share access profiles across security policies that are assigned to J Series Services Routers and SRX Series Services Gateways managed by NSM.
	Create	
	Edit	
	Delete	
Admin Roles	View	An admin role defines the access privileges for an NSM administrator.
	Create	
	Edit	
	Delete	
Admins	View	An admin is a user of the NSM management system.
	Create	
	Edit	
	Delete	
Address Objects	Create	An address object is a representation of a component of your network, such as a workstation, router, switch, subnetwork, or any other object that is connected to your network. You use address objects in NSM to specify the network components you want to protect.
	Delete	
	Edit	
	View	
Allow Installation of Pre/Post Rules from Central Manager	N/A	Pre-rules and post-rules are ordered lists of rules that are defined from the Central Manager at the global domain and subdomain levels as well as on regional servers in standalone NSM installations. This activity allows Central Manager to install Central Manager pre/post rules into a Regional server.
Antivirus Profiles	Create	An antivirus profile defines the parameters for performing virus scans.
	Delete	
	Edit	
	View	
Attack Update	N/A	This activity enables an administrator to update the attack object database on the NSM system and on each managed device that supports Deep Inspection.
Auditable Activities	Edit	Allows the administrator to select read/write or read only actions to determine what actions get reported to the Audit Log Viewer.
	View	

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Audit Logs	View	An audit log records an administrative change (such as login, update, or policy change) to the managed devices or management system.
Authentication Server	Create	An authentication server provides authentication services for NSM administrators and remote access services (RAS) users on your network. The information stored in an authentication server determines the privileges of each administrator.
	Delete	
	Edit	
	View	
AV Pattern	Update	Updates the pattern file on the devices with the latest AV signatures.
Backdoor Rulebase	Create	This activity enables an administrator to manage the Backdoor Rulebase within a security policy. Rules configured in this rulebase are supported only on IDP-capable security devices, such as an ISG2000 or ISG1000 gateway running 5.0–IDP1.
	Edit	
	View	
Blocked IP	View	Allows an administrator to view a list of IP addresses blocked because of repeated failed attempts to log in to the server.
CA	Create	A CA object represents a Certificate Authority which is a trusted third party that verifies an electronic signature.
	Delete	
	Edit	
	View	
Catalog Objects	Create	Catalog objects enable the management of report folders.
	Delete	
	Edit	
	View	
Channel	View	
CLI-based Reports	N/A	This activity enables an administrator to generate predefined and shared historical log reports using the guiSvrCli command utility.
CLI-based Security Update	N/A	This activity enables an administrator to update the attack object database on the NSM system using guiSvrCli command utility.
Config Sync Status	Check	Verifies the configuration status of the device.
Configlet	View	A configlet is a small, static configuration file that contains information on how a security device can connect to NSM.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
CRLs	Create	A Certificate Revocation List (CRL) identifies invalid certificates. You can obtain a CRL file (.crl) from the CA that issued the local certification and CA certificate for the device, and then use this file to create a Certificate Revocation List object.
	Delete	
	Edit	
	View	
Custom Troubleshoot Commands	Edit	Enables add, delete, or viewing of custom troubleshooting commands available through the Device Manager for troubleshooting devices. Custom troubleshooting commands are in addition to the standard device troubleshooting commands debug , exec , and get .
	View	
Dashboard	View	Known targets and sources of attacks or suspected targets and sources of attacks can be added to source or destination watch lists. The Dashboard is a near-real time monitor of watch lists of known targets and sources of attacks and the top 10 attacks within the previous hour.
Database Versions	Create	These activities allow an administrator to manage the database snapshot feature that provides versioning of NSM objects, policies, and managed devices.
	Delete	
	Edit	
	View	
Deep Inspection Pack Selection	Edit	Deep Inspection is a mechanism for filtering traffic that a security device permits. You can enable Deep Inspection in firewall rules to examine permitted traffic and take specific actions if the DI module finds matching attack signatures or protocol anomalies. This activity enables a system administrator to view or set the deep inspection package that is loaded on the device.
	View	
Device Admins	Import	Device administrators have permissions to administer the devices through the CLI or UI for the device itself. Importing a device administrator allows that administrator to use the NSM UI.
Device BGP Operations	N/A	Allows you to perform Border Gateway routing Protocol (BGP) operations on a device, such as connect or disconnect with a neighbor, or test the TCP connection to a specific neighbor.
Device Certificates	Generate and Upload	A device certificate authenticates packets passing through a device.
	Get	
	Delete	
Device Config To/From File	Export	Allows a system administrator to import a device configuration from a file or export a device configuration to a file.
	Import	

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Device Configuration	View Update	A device configuration is the modeled configuration that exists for a managed device within NSM.
Device CPU Limit	N/A	
Device Delta Config	View	A device delta config is a report that lists the differences between the device configuration running on the physical device and the modeled device configuration in NSM.
Device Firmware	Update	The device firmware is the software image used on the managed device.
Device Log Comments	Update	A device log comment is a user-defined description of a security event that is recorded in a device log.
Device Log Flags	Update	A device log flag is a visual icon that can be assigned to a device log. Administrators can assign flags to indicate severity, status, and other options to a device log.
Device Logs	View Hide and Unhide Purge Archive Retrieve	A device log records a security event that occurred on a security device.
Device Passwords	View	This activity enables an administrator to view device passwords in configuration summaries and Job Manager information details. Note: All passwords handled by NSM are case-sensitive.
Device Reboot	Reboot	A device reboot is a reboot command sent to a managed device to power down, and then power up.
Device Running Config	View	A device running config is a report that details the device configuration running on the physical device.
Device Site Survey	N/A	When setting up the NetScreen-5GT Wireless (ADSL) device as a wireless access point (WAP), this activity enables an administrator to scan the broadcast vicinity to see if there are any other WAPs broadcasting nearby.
Device Software Keys	Install	A device software key provides, enhances, or adds functionality for a managed device.
Device Status Monitor	View	The device status monitor tracks the status of devices, VPN tunnels, and NSRP.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Device UrlCategory	Update	The device URL category list contains predefined Web categories used in Web filtering profiles. You can update the device Web category list from the system Web category list.
Devices, Device Groups, and Templates	View	A device is a managed device.
	Create	A device group is a collection of managed devices.
	Edit	A template is a device configuration that contains predefined, static configuration information, such as networking settings, interface settings, or DNS settings.
	Delete	
DI Objects	Create	Deep Inspection (DI) attack objects contain attack patterns and protocol anomalies for known attacks and unknown attacks that attackers can use to compromise your network. However, DI attack objects don't work on their own—they need to be part of an attack object group and a DI Profile object before you can use them in a firewall rule to detect known attacks and prevent malicious traffic from entering your network.
	Delete	
	Edit	
	View	
Extranet Policy Objects	Create	Extranet policies enable you to configure and manage third-party routers.
	Delete	
	Edit	
	View	
Dial-in Objects	Create	Use Dial-in objects to dial-in and manage a device as a console. You can create and edit lists of allowed numbers and lists of blocked numbers and set a policy for unknown CNIDs.
	Delete	
	Edit	
	View	
Failover Device	N/A	Use failover to enable the device to switch traffic from the primary interface to the backup interface, and from the backup to the primary when both primary and backup interfaces are bound to the Untrust zone.
Firewall Rulebases	Create/Edit	The firewall rulebases (Zone and Global) in a security policy contain rules that handle traffic passing through the firewall. These activities enable an administrator to control or view rules in the firewall rulebases.
	View	
Force Logout Administrators	N/A	The Force Logout Administrators activity enables a system administrator to forcibly log out another administrator when there is a resource contention.
Get Entitlement from Entitlement Server	N/A	Entitles the administrator to configure devices to receive services that require subscriptions, such as internal AV or Deep Inspection Signature Service.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Group Expressions	Create	Group expressions allow the system administrator to set conditions for authentication requirements for external users stored on a RADIUS server.
	Delete	
	Edit	
	View	
GTP Objects	Create	GPRS Tunneling Protocol (GTP) objects applied to a security policy rule enable a security device to manage GTP traffic. If a GTP packet matches the rule, the device attempts to further match the packet data with the parameters set in the GTP object.
	Delete	
	Edit	
	View	
HA for guiSvrClusterMgr	N/A	Enables the system administrator to configure HA parameters for the GUI server.
Historical Log Reports	View	A historical log report is a report generated using historical log entries. If an administrator can view historical log reports, then that administrator can also view shared historical log reports and their definitions.
ICAP Objects	Create	An ICAP object defines a server or server group to act as an ICAP AV server.
	Delete	
	Edit	
	View	
IDP	Create	Enables a system manager to create, delete, edit, or view IDP devices.
	Delete	
	Edit	
	View	
IDP Cluster Monitor	View	Enables a system administrator to run the IDP Cluster Monitor and monitor IDP clusters.
IDP Migration	N/A	
IDP Profiler Operations		This activity enables an administrator to view the Profiler.
IDP Rulebase	Create	This activity enables an administrator to manage the IDP Rulebase within a security policy. Rules configured in this rulebase are supported only on IDP capable devices, such as the ISG2000, ISG1000 running 5.0-IDPI, SRX Series, J Series, and MX Series.
	Edit	
	View	
Import Device		Allows an administrator to import device information to NSM from the device, including Inventory information.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Infranet Controller Operations	N/A	This activity enables an administrator to configure a security device to interoperate with an Infranet Controller.
Inventory	View	Allows an administrator to view the hardware, software, and license inventories.
Investigative Log Reports	View	The Log Investigator generates investigative log reports based on selected criteria. This activity enables an administrator to view those log reports.
IP Pools	Create Delete Edit View	This activity allows an administrator to manage IP pools. An IP pool object contains IP ranges (a range of IP addresses within the same subnet). You use IP Pool objects to assign IP addresses to L2TP users in L2TP VPNs or local users on a specific device.
Job Status Logs	Purge	Each time NSM performs a task for which a job is created, Job Manager creates a job status log. This activity enables an administrator to purge those logs from the management system.
Jobs	View Cancel active	A job is a task that NSM performs, such as updating a device, generating a device certificate request, or importing a device.
Large Binary Objects	Create Delete Edit View	Enables an administrator to manage large binary data files. These files permit the efficient handling of large amounts of configuration data often associated with Secure Access and Infranet Controlled devices.
License	Install View	This activity allows an administrator to install or view a new NSM license.
Logged in Admins	View	This activity allows an administrator to view the administrators logged into NSM.
Multicast Rulebases	Create/Edit View	The multicast rulebase in a security policy contains multicast rules, which can handle IGMP and PIM-SM traffic. These activities enable an administrator to control or view rules in the multicast rulebase.
NAT Objects and subtree	Create Delete Edit View	Allows an administrator to manage NAT objects, which allow multiple devices to share a single object.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Network Honeypot Rulebase	Create	Allows an administrator to configure network honeypot rules that mimic seemingly vulnerable counterfeit ports to entrap would be attackers.
	Edit	
	View	
NSRP Monitor	View	NSRP Monitor tracks NSRP statistics. To enable NSM to track these statistics, you must enable "NSRP Monitor" in the NSRP properties for each cluster device.
Permitted Objects	Create	Allows an administrator to manage permitted objects. You configure permitted objects in Profiler consisting of source ID, destination, and service. These objects define what you should see on the network, rather than using attack objects which define what you do not want to see.
	Delete	
	Edit	
	View	
Phase1/Phase2 Proposal	Create	Allows an administrator to manage custom IKE phase 1 and phase 2 proposals,
	Delete	
	Edit	
	View	
Policy Custom Field Data Objects	Create	Allows an administrator to manage custom objects added to a Policy table.
	Delete	
	Edit	
	View	
Policy Custom Field Metadata Objects	Create	Allows an administrator to manage metadata used in defining custom Policy objects.
	Delete	
	Edit	
	View	
Policy Lookup Table	Modify	
Polymorphic Address Objects	Create	Allows an administrator to create polymorphic address objects. Polymorphic objects can be used as place holders for values that will be defined in a different context (in a regional server domain or subdomain, for instance).
Polymorphic Service Objects	Create	Allows an administrator to create polymorphic service objects. Polymorphic objects can be used as place holders for values that will be defined in a different context (in a regional server domain or subdomain, for instance).

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Post Domain Policies	Edit	Allows an administrator to edit post domain policy.
Pre Domain Policies	Edit	Allows an administrator to edit pre domain policy.
Profiler	View	Enables an administrator to view the Profiler.
Profiler Data	Purge	Enables an administrator to purge or clear data in the ProfilerDB.
	Clear	
Remote Settings	Create	Allows an administrator to manage a remote settings object that defines the DNS and WINS servers that are assigned to L2TP RAS users after they have connected to the L2TP tunnel. You can use remote settings objects in an L2TP VPN, and when configuring a local user on a specific device.
	Delete	
	Edit	
	View	
Routing Instance Object	Create	Use a routing instance object to share routing instances in the Radius server & LDAP server configurations within the access profile object. A routing instance object is a polymorphic object (similar to zone objects) that maintains the mapping between the actual routing instance and the device in which it is created.
	Delete	
	Edit	
	View	
Rulebases	Delete	A rulebase in a security policy contains rules that manage specific types of traffic passing through the managed device. These activities enable an administrator to delete a rulebase.
Run Online Retention Command	N/A	
Security Explorer	View	Allows an administrator to run the Security Explorer—a graphical tool that enables you to visualize and correlate network behavior based on data collected in the Profiler, Log Viewer, and Report Manager.
Schedule Objects	Create	Allows an administrator to manage Schedule objects. A schedule object defines a time interval for which a firewall rule is in effect.
	Delete	
	Edit	
	View	
Schema	Apply	Allows an administrator to download or apply schemas for managing devices.
	Download	
Schema Details	View	Allows an administrator to view the details of schemas for managing devices.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Security Policies	View	A policy is a set of rules that determines how a device handles traffic passing through the firewall.
	Create	
	Edit	
	Delete	
Servers	View	The Device Server and GUI Server form the NSM System.
	Create	
	Edit	
Service Objects	Create	Allows an administrator to manage service objects. Service objects represent the IP traffic types for existing protocol standards.
	Delete	
	Edit	
	View	
Shared Historical Log Report	Create	A shared historical log report is a user-defined historical log report that is available to users with the appropriate permissions in a domain. These reports appear under “Shared Reports” in the UI and can be generated offline with the guiSvrCli utility. If an administrator can create shared historical log reports, then that administrator can also move a report from “My Reports” to “Shared Reports”. An administrator requires permission to delete shared historical log reports in order to move a report from “Shared Reports” to “My Reports”.
	Edit	
	Delete	
Subdomain and Groups	View	A subdomain is a separate, unique representation of other networks that exist within your larger network.
	Create	
	Edit	
	Delete	
Supplemental CLIs in Devices & Templates	Edit	The Supplemental CLI option enables you to configure features on security devices not yet formally supported in NSM. This applies to security devices running a future release of ScreenOS.
SYNProtector Rulebase	Create	Allows an administrator to manage the SYN-Protector rulebase, which protects your network from SYN floods by ensuring that the three-way handshake is performed successfully for specified TCP traffic.
	Edit	
	View	
System Status Monitor	View	The system status monitor displays the status of NSM servers (GUI Server and Device Server) and the processes running on each server.
System UrlCategory	Update	The system URL category list contains predefined Web categories used in Web filtering profiles. You can update the system Web category list from the master Web category list maintained by SurfControl.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
Template Operations	N/A	Allows an administrator to perform template operations.
Unified Threat Management (UTM) objects	Create Delete Edit View	Allows an administrator to manage threats to the network through the creation of objects and profiles.
Network Topology Manager	View Manage	Allows an administrator to view and manage the network topology with the help of table and map views.
Traffic Signature Rulebase	Create Edit View	Allows an administrator to manage the traffic signature rulebase.
Troubleshoot Devices		The Troubleshoot option in the Device Manager enables you to query the status of a security device using a debug , exec , or get command.
Unblock IP	N/A	Allows an administrator to remove IP addresses from a list of IP addresses blocked because of repeated failed attempts to log in to the server.
URL Filtering	Create Delete Edit View	Allows an administrator to manage a web filtering profile for all devices by binding the profile to a firewall rule.
User Objects	Create Delete Edit View	Allows an administrator to manage local and external user objects that represent the users of your managed devices.
VLAN Objects	Create Edit Delete View	Allows an administrator to manage VLAN objects, which limit rule matching to packets within a specific LAN.
VPN Monitor	View	VPN Monitor tracks VPN tunnel statistics. To enable NSM to track these statistics, you must enable "VPN Monitor" in the Gateway properties for each VPN.

Table 16: Predefined NSM Administrator Activities (continued)

Function	Task	Description
VPNs	View	Allows an administrator to manage Virtual Private Networks (VPNs). A VPN can exist between two or more security devices, extranet devices, or RAS users.
	Create	
	Edit	
	Delete	
VSYS Profile Objects	Create	This activity enables a system administrator to configure the resource limits for a vsys device by creating or editing a vsys profile and assigning it to the vsys device.
	Delete	
	Edit	
	View	
Zone Objects	Create	Allows an administrator to create zone objects. Zone objects are available in both Central Manager and Regional Server. Similar to other polymorphic objects, the resolution in rulebases of zones is based on the mapping table that the user entered in the zone table.
	Delete	
	Edit	
	View	

Roles and Permissions

The permissions granted to some activities have changed across releases, which can cause behaviors to change following migration.

Permissions Changes in Release 2008.1

In Release 2008.1, the Create Device, Device Groups & Templates role does not allow Import Device, that is, the importing of the configuration from the device into NSM. Specifically, the Create Device, Device Groups & Templates role no longer allows the following directives:

- importConfig
- importConfigOffline
- importConfigFile
- importVsysDevice

A new Import Device role has been created to allow device import. However, for backward compatibility, any roles created prior to 2008.1 with permission to "Create Devices, Device Groups & Templates" will contain the Import Device permission as well after upgrade.

Permissions Changes in Release 2006.1

In Release 2006.1, the Create Devices, Device Groups, & Templates activity does not allow permission to run the following directives:

- Import Admin
- Export Device Config To File
- Import Device Config From File

The Import Device Admins role allows permission to run the Import Admin directive.

A new role (Export/Import Device Config to File) has been created to allow permission to run the Export Device Config To File and Import Device Config From File directives.

Permissions Changes in Release 2005.3

In Release 2005.3, the Edit Devices, Device Groups, & Templates activity no longer allows permission to run the directives listed in [Table 17 on page 90](#). Use the activities listed in the table instead.

Table 17: Changes to Edit Devices, Device Groups, & Templates Activity

Activity	Directives
View Configlet	View Configlet
Edit Device Admin	Set Root Admin (4.x device only) Set Admin Ports (4.x device only) Set Admin SSH Enable Disable (4.x device only)
Failover Device	Failover
Device BGP Operations	Modify BGP Peer Session BGP Refresh Route BGP Update Route on Peer
Update AV Pattern	Update AV Pattern
Get Entitlement from Entitlement Server	Get Entitlement from Entitlement Server
Check Config Sync Status	Check Config Sync
Intranet Controller Operations	Connect To Intranet Controller Disconnect From Intranet Controller

In Release 2005.3, the View Devices, Device Groups, & Templates activity no longer allows permission to run the directive listed in [Table 18 on page 91](#). Use the Device Site Survey activity instead.

Table 18: Changes to View Devices, Device Groups, & Templates Role

Activity	Directives
Device Site Survey	Site Survey

Assigning and Viewing Custom Roles

When you create an administrator, you can assign a custom role just as you would a default role. However, you cannot assign an activity or role that you do not possess to another administrator (the activity or role is not visible in the list of available activities or roles).

Within a domain, you can view only the custom roles that you have created or that have been assigned to you. You cannot view custom roles created by other administrators, even if the role is in the same domain and includes the same activities already assigned to you.

Configuring a User Activity in a Custom Role

Role-based administration enables you to use filters to fine-tune the permissions on the "Edit Devices, Device Groups, & Templates" and "View Devices, Device Groups, & Templates" activities. These filters include:

- Routing Configuration (for ScreenOS/IDP devices)—Allows editing of the routing configuration from ScreenOS/IDP devices, which includes the virtual router, routing configuration on the interface, and policy-based routing (PBR).
- IDP Policy Configuration (for EX Series switches)—Allows editing of policy configuration of EX Series switches in the device itself.
- Firewall Rulebase Configuration (for Junos devices that support central policy management)—Allows editing of the policy configuration of J Series routers or SRX Series gateways in the Central Policy Manager of NSM.
- Remaining Configuration—Allows editing of all device configurations, except the routing configuration for ScreenOS/IDP devices and policy configuration for EX Series switches.

To edit the filter configuration:

1. From the menu bar, click **Tools > Manage Administrators and Domains**.
2. In the RBA settings of NSM, select the **Roles** tab.
3. In the Roles dialog box, click the Edit icon to edit an existing custom role or click the Add icon to create a new role. You can also edit the filter configuration while creating a new role.
4. In the activities listed, click the **Edit Devices, Device Groups, & Templates** link. The Filter Configuration dialog box appears with a list of filters. By default, all filters of the activity are enabled.
5. Disable the filters that are not required and click **OK**.

Viewing Logged Administrators

NSM lets you view information associated with all the administrators currently logged into the system. This information includes the following columns:

- Home Domain—The name of the domain in which the administrator was created.
- Admin Name—The name of the administrator who is logged in.
- Status—Whether a user has been active in the last 5 minutes). When the administrator's status becomes inactive, NSM sends an update to the server at 1-minute intervals. This update automatically refreshes the screen with the new information. When the administrator become active again, NSM sends another update to the server that the status has changed.
- IP Address—The administrator's IP address.
- Locked Object—A detailed list about the objects locked by each administrator listed. These locked objects include object identifiers, length of time the object is locked, the lock type, and so on.

Using this information, system administrators can monitor and manage users more effectively.

Access to this feature is granted only to system administrators and read-only administrators. You can access this information from the Tools menu by selecting the Logged In Administrators menu item. By default, this activity is assigned to the predefined system administrator role.

Forcing an Administrator to Log Out

As of Release 2007.3, the system administrator can forcibly log out an administrator.

To log out an administrator forcibly:

1. From the menu bar, click **Tools > Logged In Administrators**. A list of logged-in administrators appears.
2. Right-click the name of the administrator to be logged out under the Admin Name column. The Logout button appears.
3. Click **Logout** to forcibly log out the administrator.

The logged out administrator sees the Session Logged Out dialog box. After the administrator clicks **OK**, the dialog box disappears.

By default, this activity is assigned to the system administrator. Custom roles can be created in the NSM installation with this activity enabled. The Audit log shows which system administrator logged out which administrator.

When the administrator is forcibly logged out, the following rules apply:

- Any objects that were locked by the administrator during the login session are unlocked.
- The server operations triggered by the logged-out GUI (Jobs, Reports, Log-Viewer, and so on) run to completion in the server.
- An administrator cannot forcibly log out from his own session.
- Server resources such as the GUI Server connection to a client and a port are freed.
- In a central or a regional server setup, forced logout applies only to a server. The administrator is not logged out from other servers.

Creating Subdomains

To create a subdomain, in the Subdomains tab, add a new subdomain and click **OK**. The new subdomain appears in the subdomain list.



NOTE: You cannot create VPNs between devices in different domains,

You can add unlimited subdomains in the global domain. However, you cannot create subdomains within a subdomain. When you view the Manage Administrators and Domains dialog box from within a subdomain, the Subdomains tab does not appear. To view a subdomain in the main display area, select it from the list at the top of the navigation tree.



NOTE: Objects and groups defined in the global domain are not visible in subdomains.

Viewing Current Domain Detail

The domain detail displays the subdomains, administrators, their roles, and authentication server for the currently selected domain (subdomains appear only when you view the global domain).

You can designate a default RADIUS authentication server for the global domain and for each subdomain. The default authentication server is used:

- To authenticate administrators when they log into the NSM system
- To authenticate RAS users in VPNs

For step-by-step instructions on configuring a RADIUS authentication server to authenticate administrators and users, see the *Network and Security Manager Online Help* topic “Editing the Domain Contact.”

Example: Configuring Role-Based Administration

In this example, you configure a domain structure for an Internet service provider (ISP) with a co-location facility in New York that handles customers across four states. The company uses a two-letter state postal code combined with the customer name. That

ISP's goal is to manage all devices and policies from the co-location facility and provide read-only permission for customers to view log entries and generate reports. No VPNs are used.

To configure this domain structure, use the following process:

- Create the subdomains.
- Create the subdomain administrators.
- Create the read-only customer administrator.
- Log in as each administrator (for verification).

Step 1: Create the Subdomains

In this step, you create a subdomain for each company that uses the ISP.

1. Log in to the global domain as the super administrator.
2. From the Menu bar, select **Tools > Manage Administrators and Domains**.
3. Click the **Subdomains** tab, then click the Add icon to create a subdomain for the first customer. Configure the following four subdomains:
 - MA_company1
 - NH_company2
 - RI_company3
 - VT_company4
4. Click **OK** to save your changes.

Step 2: Create the Subdomain Administrator

In this step, you create a subdomain administrator with full permissions for the domain.

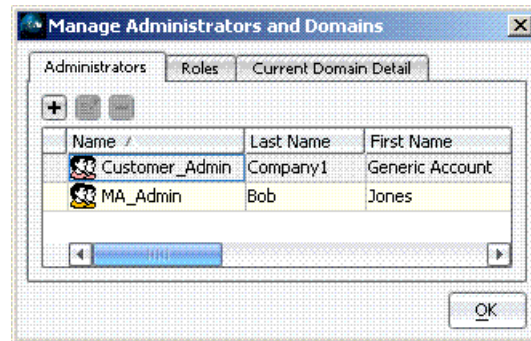
1. Using the domain menu (at the top of the navigation tree), select the first subdomain (MA_company1). NSM loads the subdomain.
2. From the Menu bar, click **Tools > Manage Administrators and Domains**.
3. In the Administrators tab, click the Add icon to create the primary administrator for this domain:
 - In the General Properties tab, enter a name, color, and contact information for the administrator.
 - In the Authorization tab, leave the default authentication as Local and configure a password for the administrator.
 - In the Permissions tab, click the Add icon, then configure the role as Domain Administrator (predefined) and the Domain as MA_company1.
4. Click **OK** to save your changes.
5. Repeat step 1 through step 4 for each subdomain.

Step 3: Create the Viewing and Reporting Administrator

In this step, you create a custom role and administrator account that permits the ISP customers to view log entries and generate reports for devices in their subdomain.

1. Using the domain menu (at the top of the navigation tree), select the first subdomain (MA_company1). NSM loads the subdomain.
2. From the Menu bar, click **Tools > Manage Administrators and Domains**.
3. In the Roles tab, click the Add icon. Name the new role, add an optional comment, choose a color, and then select activities and reporting permissions for this domain.
4. Click **OK** to save your changes.
5. In the Administrators tab, click the **Add** icon to create the customer administrator for this domain:
 - In the General Properties tab, enter a name, color, and contact information for the administrator.
 - In the Authorization tab, leave the default authentication as **Local** and configure a password for the administrator.
 - In the Permissions tab, click the **Add** icon, then configure the role as Viewing & Reporting and the Domain as MA_company1.
6. Click **OK** to save your changes and return to the Administrators tab, which now displays the following administrators:

Figure 26: Manage Administrators and Domains: Administrators Tab



7. Click **OK** to save your changes.
8. Repeat step 1 through step 7 for each subdomain.

Step 4: Verify Administrator Accounts

In this step, you log in as each administrator to verify their permissions (administrators must log in to the domain in which they were created). Start a new instance of the NSM UI, and then log in as the following administrators to test permissions:

- Logging in as the Domain Administrator—To log in as the domain administrator, in the login screen, enter the subdomain/domain administrator name (**MA_company1/MA_Admin**), the password, and the GUI Server IP address.

Click **OK** to log in. The NSM navigation tree and main display area appear. Because the domain administrator account has full permissions for the domain, the UI displays all modules and enables all functionality for the domain. However, the domain menu (at the top of the navigation tree) displays only the current domain, restricting the domain administrator to that domain.

Repeat for each subdomain and domain administrator.

- Logging in as the Customer Administrator—To log in as the customer administrator, in the login screen, enter the subdomain/domain administrator name (**MA_company1/Customer_Admin**), the password, and the GUI Server IP address.

Click **OK** to log in. The NSM navigation tree and main display area appear. Because the customer administrator account has permission only for viewing and reports, the UI displays only the modules that are used for those permissions (note that Server Manager, Job Manager, and the Audit Log Viewer do not appear). Additionally, all Add, Edit, and Delete icons appear in gray, indicating that the administrator cannot perform these tasks.

Repeat for each subdomain and customer administrator.

PART 2

Integrating

- [Adding Devices on page 99](#)
- [Configuring Devices on page 199](#)
- [Updating Devices on page 257](#)
- [Managing Devices on page 279](#)

CHAPTER 4

Adding Devices

This chapter provides information about adding Juniper Networks devices to your network. These devices can include routers and switches, as well as the security devices that protect your network against malicious traffic.

Juniper Networks Network and Security Manager (NSM) can manage all Juniper Networks devices running ScreenOS 5.x and later, IDP 4.0 and later, Junos 9.0 and later, IC 2.2 or later, or SA 6.3 or later. NSM can also manage vsys configurations, NetScreen Redundancy Protocol (NSRP) clusters, Junos Redundancy Protocol (JSRP) clusters, vsys clusters, and extranet devices.



NOTE: NSM does not support NetScreen-5 devices, NetScreen-100 devices, or NetScreen-1000 devices. NSM supports the following ScreenOS releases: 5.0r11, 5.1r4, 5.2r3, 5.3r10, 5.4r11, 6.0r2, 6.1r4, 6.2 and 6.3.

Before you can manage a device with NSM, you must add the device to the management system. NSM supports adding individual devices or many devices at a time.



NOTE: If you have been managing your IDP Sensors using the IDP Management Server and UI, you should migrate your Sensors instead of adding them manually. Refer to the *IDP-NetScreen-Security Manager Migration Guide* for procedures and additional information.

Use Rapid Deployment (RD) to quickly add ScreenOS devices in nontechnical environments with no staging requirements.

This chapter contains the following sections:

- [About Device Creation on page 100](#)
- [Before You Begin Adding Devices on page 102](#)
- [Supported Add Device Workflows by Device Family on page 113](#)
- [Importing Devices on page 114](#)
- [Modeling Devices on page 133](#)
- [Using Rapid Deployment \(ScreenOS Only\) on page 142](#)

- [Adding Vsys Devices on page 151](#)
- [Adding an Extranet Device on page 155](#)
- [Adding Clusters on page 155](#)
- [Importing an SRX Series Cluster into NSM on page 172](#)
- [Adding a Blade Server on page 176](#)
- [Adding Multiple Devices Using Automatic Discovery \(Junos, SA, and IC Devices \) on page 180](#)
- [Adding Many Devices Using CSV Files on page 181](#)
- [Adding Device Groups on page 194](#)
- [Setting Up NSM to Work With Infranet Controller and Infranet Enforcer on page 195](#)

About Device Creation

Before NSM can manage devices, you must first add those devices to the management system using the UI. To add a device, you create an object in the UI that represents the physical device, and then create a connection between the UI object and the physical device so that their information is linked. When you make a change to the UI device object, you can push that information to the real device so the two remain synchronized. You can add a single device at a time or add multiple devices all at once.



NOTE: The connection between a managed device and the NSM Device Server must be at least 28.8 Kbps.

You can add the following types of devices:

- Physical devices—“[Importing Devices](#)” on page 114 and “[Modeling Devices](#)” on page 133 later in this chapter provide details on how to add an existing or new device into NSM. These devices include:

- ScreenOS/IDP Security devices—A security device or system (such as a NetScreen-5GT device, an ISG device, or an IDP Sensor) manages firewall, VPN, or IDP activities on your network.

Adding an IDP Sensor to NSM does not migrate existing settings. If you have existing IDP Sensors that are managed by an IDP Management Server, you should migrate those Sensors using the instructions in the *IDP-NetScreen Security Manager Migration Guide*.

- Junos devices—routers, gateways, and switches that run Junos OS:
 - EX Series devices are Ethernet switches that can be added to your network and managed through NSM.
 - J Series devices—Highly secure routers that can be added to your network and managed through NSM.

- SRX Series gateways—Firewall/VPN systems that have integrated service layer technologies such as IDP, AV, or Web Filtering.
- M Series and MX Series routers—Carrier Ethernet routers and services routers.
- Unified Access Control (Infranet Controller) devices—The policy management server of the Juniper Networks LAN access control solution.
- SSL VPN (Secure Access) devices.
- Virtual Chassis—Stacked EX Series devices functioning as one logical EX Series switch or an SRX cluster represented in NSM as a virtual chassis.
- Vsys devices—Virtual devices that exist within a physical security device. For details on adding vsys devices, see [“Adding Vsys Devices” on page 151](#).
- Clusters—A cluster is a set of multiple devices joined together in a high availability configuration to ensure continued network uptime. You can build ScreenOS clusters (NSRP), IDP clusters, Junos clusters from J Series or SRX Series devices, Secure Access clusters, or Infranet Controller clusters. For details on adding clusters, see [“Adding Clusters” on page 155](#).
- Vsys clusters—A vsys cluster device is a vsys device that has a cluster as its root device. For details on adding vsys clusters and vsys devices, see [“Adding a Vsys Cluster and Vsys Cluster Members” on page 170](#).
- Extranet devices—An extranet device is a third-party device or a device not managed by NSM. For details on adding extranet devices, see [“Adding an Extranet Device” on page 155](#).

Before adding any device type, you must determine the device status. After adding the device, you must verify the device configuration.

Determine Device Status

How you add your devices to the management system depends on the network status of the device. You can import deployed devices, or you can model devices that have not yet been deployed:

- Import deployed devices—Deployed devices are the devices you are currently using in your existing network. These devices have already been configured with a static or dynamic IP address and other basic information. For deployed devices, you can import the existing device configuration information into NSM.



NOTE: To import device configurations, the connection between NSM and the managed device must be at least 28.8 Kbps. For details on installing NSM on your network, refer to the *Network and Security Manager Installation Guide*.

- Model undeployed devices—Undeployed devices are devices that you are not currently using in your network, and typically do not have IP addresses, zones, or other basic network information. For undeployed devices, you can model a new device configuration, and later install that configuration on the device.

The Add Device wizard walks you through each step of the device creation process. It prompts you to choose a workflow. **Device is reachable** is the default option. The wizard then prompts you for specific device information, such as the device platform name, OS name and version, IP address, and device administrator name, and then uses that information to detect the device. You can then choose to modify the displayed name of the device and assign a color to the device. If the host name is not unique within NSM or is undetected, the Add Device wizard generates a validation error, forcing you to add a valid device name in order to proceed with adding the physical device to the Device Server.

After the physical device connects, it is considered to be a *managed device*, meaning it is now under the control of NSM.

Verifying Device Configuration

For managed devices that use imported device configurations, you should verify that all device information was imported correctly. To identify any discrepancies, you can generate a summary of the differences between the physical device configuration and the NSM device configuration. This summary is known as a Delta Config summary. It is also a good idea to check your imported security policies, objects, and VPNs to become familiar with how the NSM UI displays them. The Delta Config summary is available for all devices supported by NSM except IDP devices.

For managed devices that use modeled device configurations, you should verify that all device information was pushed to the physical device correctly. To identify discrepancies, generate a summary of the device configuration running on the physical device. This summary is known as a Get Running Config summary.

Managing the Device

After adding a device, you can manage its configuration, objects, and security policies in the UI. You can also view traffic log entries for your device in the Log Viewer, view administrative log entries in the Audit Log Viewer, and monitor the status of your devices in the Realtime Monitor.

You can also delete devices from NSM, and reimport them if necessary. Deleting a device removes all device configuration information from the management system, but might be the best solution if you need to perform extensive troubleshooting or reconfigure the device locally. After you have made the necessary changes locally, you can then reimport that device into NSM. However, during reimport, NSM imports all device configuration data—not just the data that was changed; any changes that exist in the modeled configuration are lost during reimport. Additionally, after reimporting a device configuration, you must reassign the imported policy to the device.

If you delete a device that was added using Rapid Deployment (see [“Using Rapid Deployment \(ScreenOS Only\)” on page 142](#)), you must also re-create the configlet and install it again on the device.

Before You Begin Adding Devices

Before adding a device to NSM, decide the following:

- Will you import or model the device?
- Will the device reside in the global domain or a subdomain?
- Will you add one or many devices?

Additionally, collect the following information about the device:

- OS and OS version running on the device
- Port Mode used by the device (some ScreenOS devices only)

The following sections provide details to help you make decisions about adding devices and determine device information.

Importing Versus Modeling

You must decide if you want to import or model your devices in NSM.

Importing Device Configurations

You can add devices in your existing network into NSM and import their configurations. Using the Add Device wizard, you configure a connection between the management system and the physical device, and then import all device parameters, policies, objects, VPNs, and so on.

After you have imported several devices, you can start using system-level management features, such as:

- The policy merge tool (for ScreenOS, J Series, and SRX Series devices) that can merge several device security policies into a single, efficient policy that is easy to maintain.
- Device groups, which group devices by function, location, or platform to make updating easier.
- The VPN Manager, for creating VPNs across multiple devices quickly.

If you modify a device that supports centralized policy management and import or reimport the device into NSM, a new policy is automatically created using the following naming syntax: `device_1`. (Each new policy increments the name.) Devices are not assigned to the new policy. If you reimport a device with no changes, then a duplicate policy is not created.

Importing and Templates

If you assign a template to a device before connecting to and importing the device, later changes to the template will change values on the device. If you assign a template to a device after importing it, changes to the template will not change set values on the device unless you specifically have the template override the existing values.

Modeling Device Configurations

For most new or undeployed device types, you can add and configure the device in NSM, and then activate the configuration when you are ready to deploy the physical device on your network.

Before connecting to the device, create a device object (using the Add Device wizard) that represents the OS and device type of the actual, physical device. Then model the device configuration in the NSM UI. Configure all device features—zones, interfaces, virtual routers, policies, logging features. Finally, activate the device (using the Activate Device wizard) by configuring a connection between the management system and the physical device, and then update the modeled configuration to the device.



NOTE: You cannot activate modeled Secure Access or Infranet Controller devices. These devices must be imported into NSM.

To quickly configure multiple ScreenOS devices, use templates (reusable, custom device settings such as DNS settings, PKI settings, and so on) and objects (reusable, custom objects such as NAT objects, CA certificates, and Address objects). For large deployments that involve multiple devices in nontechnical environments, use Rapid Deployment (RD) to bring new ScreenOS devices under NSM management for initial configuration.

Device Add Process

Although the Add Device wizard and Activate Device wizard automatically handle many of the tasks involved in adding a device to the management system, you might need to perform some steps manually after using a wizard to complete the device add process.

The amount of manual involvement when adding a device to NSM depends on several factors, such as whether you are importing a deployed device or activating a modeled device, the software version the device is running, and the type of IP address (static or dynamic) the device uses to connect to the management system. The following procedures guide you through the process.

Selecting the Domain

Determine the domain in which you want to place the device. A domain is a logical grouping of devices, device security policies, and device access privileges. NSM includes a global domain by default. You can also create additional domains, called subdomains, that exist within the global domain. Before you add the device, you must select the domain that contains the device; after the device is created, it appears only in that domain and must be managed from that domain.

When you log in to the UI for the first time after installing the management system, NSM loads the global domain by default, and the Device Manager contains no devices. To begin adding devices, ensure that you are in the domain you want to add the device to:


- Add device to the global domain—Ensure that you are in the global domain and begin the device addition process.
- Add device to an existing subdomain—From the domain menu at the top of the navigation tree, select the subdomain you want to add the device to, and then begin the device addition process. The domain menu displays only the domains you have access to.
- Add device to a new subdomain—You must first create the new subdomain in NSM before adding devices to that subdomain. For details on creating new subdomains,

see [“Configuring Role-Based Administration” on page 63](#). After you have created the subdomain, select it from the domain menu and begin the device creation process.

After you have created subdomains, you can load a specific subdomain automatically when you log in to the UI. You must have access to that subdomain, and permissions to create, edit, and view devices in that subdomain.

Domain selection is critical when using VPNs. You can create VPNs only between devices within the same domain. If you need to add a device to a VPN in a different domain, add the device as an extranet device in the domain that contains the VPN, and then add the extranet device to the VPN (as shown in [Figure 27 on page 105](#)).

Figure 27: Connecting Devices from Different Domains in VPNs



Adding Single or Multiple Devices

Determine whether you want to add devices to NSM individually or add many devices all at one time. The adding process for a single device is different than the process for adding multiple devices.

When adding a single security device, use the Add Device wizard to create the device object in NSM. To activate a modeled device or create a configlet, use the Activate Device wizard. You can import or model device configurations from a device running ScreenOS 5.0.x or later (except 6.0r1), IDP 4.0 or later, Junos 9.0 or later, SA 6.2 or later, or IC 2.2 or later.

When adding many devices, you first create a **.csv** file that defines all required and optional parameters for each device, and then use the Add Many Devices wizard to create a device object for each device in NSM. To activate modeled devices or create configlets for each device, use the Activate Many Devices wizard.

You can use the Add Many Devices wizard for the following tasks:

- Import many ScreenOS, Junos, Secure Access, or Infranet Controller devices at one time.
- Model many ScreenOS devices at one time.

- Model, create configlets for, and activate multiple ScreenOS devices at one time for use with Rapid Deployment.



NOTE: You cannot use the Add Many Devices wizard to add multiple IDP devices.

You can also add multiple Junos devices using the Device Discovery Rule wizard, which scans a range of IP addresses looking for Junos devices that match specific criteria.

Additionally, you can use the Activate Many Devices wizard to create configlets for and activate multiple ScreenOS devices at one time for use with Rapid Deployment. However, you cannot activate multiple ScreenOS devices without creating configlets. For details, see [“Using Rapid Deployment \(ScreenOS Only\)” on page 142](#).

Specifying the OS and Version

During the Add Device or Add Many Devices process, you might need to specify the operating system and version that is running on the device or devices:

- For devices that use a static IP address, you do not need to specify the operating system or the OS version. NSM automatically detects this information during the add process.
- For undeployed devices or for devices that use a dynamically assigned IP address, you must specify the operating system name and OS version of the device. NSM validates the version during the model or add process.

Additionally, ensure that the devices you are adding to NSM are running a supported version of the OS. For example, NSM no longer supports devices running 4.x or earlier versions of ScreenOS. If you are not running a supported version, you must upgrade your devices before adding them into the management system. Contact Juniper Networks customer support for details.

Determining Port Mode (ScreenOS Devices Only)

For some ScreenOS security devices, you can select a *port mode* during the model or add device process. The port mode automatically sets different port, interface, and zone bindings for the device. *Port* refers to a physical interface on the back of the physical security device; ports are referenced by their labels: Untrusted, 1-4, Console, or Modem. *Interface* refers to a logical interface that you can configure after you have added the device to the management system. You can bind each port to only one interface, but you can bind multiple ports to a single interface.

On the NetScreen-5XT and NetScreen-5GT devices, you can configure one of the following port modes:


- [Trust-Untrust Port Mode on page 107](#)
- [Home-Work Port Mode on page 107](#)
- [Dual-Untrust Port Mode on page 108](#)
- [Combined Port Mode on page 109](#)

- [Trust-Untrust-DMZ Port Mode on page 109](#)
- [Trust/Untrust/DMZ \(Extended\) Mode on page 110](#)
- [DMZ-Dual-Untrust Port Mode on page 111](#)
- [Port Mode Summary on page 112](#)
- [Changing the Port Mode on page 113](#)

Trust-Untrust Port Mode

Trust-Untrust mode is the default port mode. See [Figure 28 on page 107](#) for port, interface, and zone bindings.

Figure 28: Trust-Untrust Port Mode Bindings




This mode provides the following bindings:

- Binds the Untrusted Ethernet port to the Untrust interface, which is bound to the Untrust security zone.
- Binds the Modem port to the serial interface, which you can bind as a backup interface to the Untrust security zone.
- Binds the Ethernet ports 1 through 4 to the Trust interface, which is bound to the Trust security zone.

Home-Work Port Mode

Home-Work mode binds interfaces to the Untrust security zone and to Home and Work security zones. The Home and Work zones enable you to segregate users and resources in each zone. In this mode, default policies permit traffic flow and connections from the Work zone to the Home zone, but do not permit traffic from the Home zone to the Work zone. By default, there are no restrictions for traffic from the Home zone to the Untrust zone. See [Figure 29 on page 108](#) for port, interface, and zone bindings.

Figure 29: Home-Work Port Mode Bindings


This mode provides the following bindings:

- Binds the Ethernet ports 1 and 2 to the ethernet1 interface, which is bound to the Work security zone.
- Binds the Ethernet ports 3 and 4 to the ethernet2 interface, which is bound to the Home security zone.
- Binds the Untrusted Ethernet port to the ethernet3 interface, which is bound to the Untrust security zone.
- Binds the Modem port to the serial interface, which you can bind as a backup interface to the Untrust security zone.

Dual-Untrust Port Mode

Dual Untrust mode binds two interfaces—a primary and a backup—to the Untrust security zone. The primary interface is used to pass traffic to and from the Untrust zone, while the backup interface is used only when there is a failure on the primary interface.

See [Figure 30 on page 108](#) for port, interface, and zone bindings.

Figure 30: Dual-Untrust Port Mode Bindings

This mode provides the following bindings:

- Binds the Untrusted Ethernet port to the ethernet3 interface, which is bound to the Untrust security zone.
- Binds Ethernet port 4 to the ethernet2 interface, which is bound as a backup interface to the Untrust security zone (the ethernet3 interface is the primary interface to the Untrust security zone).

- Binds the Ethernet ports 1, 2, and 3 to the ethernet1 interface, which is bound to the Trust security zone.



NOTE: The serial interface is not available in Dual Untrust port mode.

Combined Port Mode

Combined mode enables both primary and backup interfaces to the Internet and the segregation of users and resources in Work and Home zones.




NOTE: For the NetScreen-5XT, the Combined port mode is supported only on the NetScreen-5XT Elite (unrestricted users) platform.

See [Figure 31 on page 109](#) for port, interface, and zone bindings.

Figure 31: Combined Port Mode Bindings

The ethernet4 interface is the primary interface to the Untrust zone. The ethernet3 interface (shown in gray) is the backup interface to the Untrust zone.



This mode provides the following bindings:

- Binds the Untrusted Ethernet port to the ethernet4 interface, which is bound to the Untrust zone.
- Binds Ethernet port 4 to the ethernet3 interface, which is bound as a backup interface to the Untrust zone (the ethernet4 interface is the primary interface to the Untrust security zone).
- Binds the Ethernet ports 3 and 2 to the ethernet2 interface, which is bound to the Home zone.
- Binds Ethernet port 1 to the ethernet1 interface, which is bound to the Work zone.



NOTE: The serial interface is not available in Combined port mode.

Trust-Untrust-DMZ Port Mode


Trust/Untrust/DMZ mode binds interfaces to the Untrust, Trust and DMZ security zones, enabling you to segregate Web, e-mail, or other application servers from the internal network.



NOTE: The Trust/Untrust/DMZ port mode is supported only on the NetScreen-5GT Extended platform.

See [Figure 32 on page 110](#) for port, interface, and zone bindings.

Figure 32: Trust-Untrust-DMZ Port Mode Bindings




This mode provides the following bindings:

- Binds the Ethernet ports 1 and 2 to the ethernet1 interface, which is bound to the Trust security zone.
- Binds the Ethernet ports 3 and 4 to the ethernet2 interface, which is bound to the DMZ security zone.
- Binds the Untrusted Ethernet port to the ethernet3 interface, which is bound to the Untrust security zone.
- Binds the Modem port to the serial interface, which you can bind as a backup interface to the Untrust security zone.

Trust/Untrust/DMZ (Extended) Mode

Trust/Untrust/DMZ (Extended) mode binds interfaces to the Untrust, Trust, and DMZ security zones, allowing you to segregate web, email, or other application servers from the internal network. See [Figure 33 on page 110](#).

Figure 33: Extended Port-Mode Interface to Zone Bindings



[Table 19 on page 111](#) provides the Extended mode interface-to-zone bindings.

Table 19: Extended Bindings

Port	Interface	Zone
Untrusted	Untrust	Untrust
1	ethernet1	Trust
2	ethernet1	Trust
3	ethernet2	DMZ
4	ethernet2	DMZ
Modem	serial	Untrust

DMZ-Dual-Untrust Port Mode

DMZ/Dual Untrust mode binds interfaces to the Untrust, Trust, and DMZ security zones, enabling you to pass traffic simultaneously from the internal network.




NOTE: The DMZ/Dual Untrust port mode is supported only on the NetScreen-5GT Extended platform using ScreenOS 5.1 and later.

See [Figure 34 on page 111](#) for port, interface, and zone bindings.

Figure 34: DMZ Dual Untrust Port Mode

The ethernet3 and ethernet4 interfaces are active simultaneously. In this diagram, the two interfaces are bound to the Untrust zone to allow for load balancing.



This mode provides the following bindings:

- Binds the Ethernet ports 1 and 2 to the ethernet1 interface, which is bound to the Trust security zone.
- Binds the Ethernet port 3 to the ethernet2 interface, which is bound to the DMZ security zone.
- Binds the Ethernet port 4 to the ethernet3 interface, which is bound to the Untrust security zone.
- Binds the Untrust Ethernet port to the ethernet4 interface, which is bound to the Untrust security zone.



NOTE: The serial interface is not available in DMZ-Dual-Untrust port mode.

To enable failover instead of passing traffic simultaneously, you can configure the failover settings in the device configuration after you have added the device to the management system. For details, see *Network and Security Manager: Configuring ScreenOS and IDP Devices Guide*.

Port Mode Summary

Table 20 on page 112 and Table 21 on page 112 summarize the port, interface, and zone bindings provided by the ScreenOS port modes. Port numbers are as labeled on the Juniper Networks security device chassis. The Trust-Untrust mode entries represent the default port modes.

Table 20: Security Device Port Mode Summary (Part 1)

Port	Trust-Untrust Mode		Home-Work Mode		Dual Untrust Mode	
	Interface	Zone	Interface	Zone	Interface	Zone
Untrusted	Untrust	Untrust	ethernet3	Untrust	ethernet3	Untrust
1	Trust	Trust	ethernet1	Work	ethernet1	Trust
2	Trust	Trust	ethernet1	Work	ethernet1	Trust
3	Trust	Trust	ethernet2	Home	ethernet1	Trust
4	Trust	Trust	ethernet2	Home	ethernet2	Untrust
Modem	serial	Null	serial	Null	N/A	N/A

Table 21: Security Device Port Mode Summary (Part 2)

Port	Combined Mode		Trust/Untrust/DMZ Mode		Dual Untrust Mode	
	Interface	Zone	Interface	Zone	Interface	Zone
Untrusted	ethernet4	Untrust	ethernet3	Untrust	ethernet4	Untrust
1	ethernet1	Work	ethernet1	Trust	ethernet1	Trust
2	ethernet2	Home	ethernet1	Trust	ethernet1	Trust
3	ethernet2	Home	ethernet2	DMZ	ethernet2	DMZ
4	ethernet3	Untrust	ethernet2	DMZ	ethernet3	Untrust
Modem	N/A	N/A	serial	Null	N/A	N/A

Changing the Port Mode

After you have added a device, you cannot change the port mode setting using NSM. You must delete the device from the management system, change the port mode using the WebUI or CLI, and then add the device again using the Add Device or Add Many Devices wizard.

When changing the port mode on the device, be aware that:

- Changing the port mode removes any existing configurations on the security device and requires a system reset.
- Issuing the **unset all** CLI command does not affect the port mode setting on the security device.

Supported Add Device Workflows by Device Family

Table 22 on page 113 summarizes the methods or workflows you can use to add devices from each supported device family.

Table 22: Supported Add Device Workflows by Device Family

Add Device Workflow	Device Family (Operating System name)						
	ScreenOS	IDP	Secure Access (SA)	Infranet Controller (IC)	J Series and SRX Series devices (Junos)	EX Series devices (Junos)	M Series and MX Series devices (Junos)
Device is reachable	yes	yes	yes	yes	yes	yes	yes
Device is not reachable	yes	yes	yes	yes	yes	yes	yes
Model and activate device	yes	yes	no	no	yes	yes	yes
Rapid deployment (configlets)	yes	no	no	no	no	no	no
Device discovery	no	no	no	no	yes	yes	yes

Table 22: Supported Add Device Workflows by Device Family (continued)

Add Device Workflow	Device Family (Operating System name)						
	ScreenOS	IDP	Secure Access (SA)	Infranet Controller (IC)	J Series and SRX Series devices (Junos)	EX Series devices (Junos)	M Series and MX Series devices (Junos)
Add and import many devices (CSV file) - Device is reachable	yes	no	no	no	yes	yes	yes
Add and import many devices (CSV file) - Device is not reachable	yes	no	yes	yes	yes	yes	yes
Model many devices (CSV file)	yes	no	no	no	no	no	no

Importing Devices

NSM can import device configurations from devices running ScreenOS 5.x or later, IDP 4.0 or later, Junos 9.0 or later, SA 6.2 or later, or IC 2.2 or later.

When importing from a device, the management system connects to the device and imports Data Model (DM) information that contains details of the device configuration. The connection is secured using Secure Server Protocol (SSP), a proprietary encryption method; an always-on connection exists between the management system and device.



NOTE: Importing the running configuration from a device completely overwrites all configuration information stored within NSM for that device. To help avoid accidental configuration overwriting, when you attempt to import a configuration from a currently managed security device, NSM prompts you for confirmation to import.



NOTE: IDP rulebases cannot be imported.

In some cases, you may need to configure NACN or other features on the physical device to enable the device to connect to NSM.

For details about adding multiple devices at one time, see [“Adding Many Devices Using CSV Files” on page 181](#).

Requirements

To import a single device:

- The physical device must have Telnet or SSH enabled.
- You must have the device connection information (IP address, connection method) and device administrator name and password available. For standalone IDP Sensors, you must also have the root password.



NOTE: All passwords handled by NSM are case-sensitive.

- The device must be *staged*; that is, it must be physically connected to your network with access to network resources.
- The device must have at least one interface that has an IP address. Devices that use a dynamically assigned IP address must also support NACN.
- The device must be operating in the desired port mode. You cannot change the operational mode after importing the device into NSM. Port modes apply only to some ScreenOS devices.

Adding and Importing Devices with Static IP Addresses

A static IP address is an IP address that does not change. Not all device families support static IP addresses. The following device families do:

- [ScreenOS Devices on page 116](#)
- [IDP Sensors on page 117](#)
- [Junos Devices on page 118](#)
- [SA and IC Devices on page 120](#)



NOTE: Before you add devices with IPv6 addresses running ScreenOS or SRX Series high-end devices with IPv6 addresses to NSM, install the latest device images with IPv6 NSM agent and outbound-ssh support respectively on the devices. (For supported versions, refer to the device release notes.) Otherwise, NSM fails to add devices with static IPv6 addresses.

ScreenOS Devices

ScreenOS devices must be running ScreenOS 5.0 or later release to be imported into NSM 2008.1 or later release.

To import a ScreenOS 5.0 or later device with a known IP address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device is Reachable** (default).
5. Click **Next**. The Specify Connection Settings dialog box opens.
6. Enter the following connection information:
 - Enter the IP Address of the security device.



NOTE:

- Beginning in NSM 2012.2R10, you must select the IP address format (IPv4 or IPv6) of your current IP address.
- IPv6 format is supported only for devices running ScreenOS and SRX Series high-end devices.

- Enter the username of the device administrator.
- Enter the password for the device administrator.



NOTE: All passwords handled by NSM are case-sensitive.

- Select the connection method (Telnet, SSH version 1, SSH version 2) and the port number for the selected service.

If you selected Telnet, click **Next** and go directly to step 7.

If you selected an SSH version, click **Next** and the Verify Device Authenticity dialog box opens. The device wizard displays the RSA Key FingerPrint information; to prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.

7. After the wizard displays the autodetected device information, verify that the device type, ScreenOS version, and the device serial number are correct. NSM autodetects the hostname configured on the device and uses it as the device name. You can also change the autodetected hostname.
8. Modify the autodetected device name within the device from its config editor page in NSM, and select **Update device**.

If you modified the device host name through the Junos OS CLI, SNMP, or J-Web interface, you can modify the device name again in NSM after importing the device,

using the edit option. If the device was bulk added, the name you specify in the CSV file is used.



NOTE: If you select the **Device is not reachable** or the **Model Device** workflow, NSM cannot detect the hostname automatically. You need to specify a device name.

9. Click **Next** to add the device to NSM.
10. After the device is added, click **Next** to import the device configuration.
11. Click **Finish** to complete the Add Device wizard.
12. Double-click the device in Device Manager to view the imported configuration.



NOTE: After importing a NetScreen-5GT device that uses extended port mode, NSM displays the modes as “ns5GT-Trust-Untrust-DMZ” and sets the license mode to Extended.

To check the device configuration status, mouse over the device in Device Manager (you can also check configuration status in Device Monitor). The device status displays as **Managed**, indicating that the device has connected and the management system has successfully imported the device configuration.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See [“Verifying Imported Device Configurations” on page 131](#) for details.

IDP Sensors

IDP Sensors running versions of the IDP software earlier than 4.0 cannot be imported. If those Sensors are already being managed by IDP Manager, migrate your IDP Management Server and the Sensors with it. Refer to the *IDP-NetScreen-Security Manager Migration Guide* for more information.

You need to upgrade unmanaged Sensors to 4.0 or later before adding them to NSM. See the *IDP Installer's Guide* for more information.

To import an IDP 4.0 device with a known IP address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device Is Reachable** (default).
5. Click **Next**. The Specify Connection Settings dialog box opens.
6. Enter the following connection information:

- Enter the IP Address of the Sensor.
- Enter the username of the device administrator.
- Enter the password for the device administrator.
- Enter the password for the device root user.



NOTE: All passwords handled by NSM are case-sensitive.

- Select the connection method (SSH Version 2) and the port number for the selected service.
 - Select the Port Number. The default (port 22) is recommended.
7. Click **Next**. The Verify Device Authenticity dialog box opens. The device wizard displays the RSA Key FingerPrint information. To prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.
 8. After you have verified the key, click **Next** to display the autodetect device information. This action will take a moment.
 9. Verify that the device type, OS version, device serial number, and device mode are correct. The host name from the device is also discovered. You can change this name if desired.
 10. Click **Next** to have NSM add the Sensor as a managed device.
 11. Click **Next** to have NSM import settings already present on the Sensor.
 12. Click **Finish** to complete the add operation.

An IDP 4.1 or later sensor is also updated with the Juniper Networks Recommended policy. IDP 4.0 Sensors cannot use the Recommended policy.

The Job Information dialog box shows the status of the Update Device job.

13. After the Update Device job is complete, double-click the device in Device Manager to view the imported configuration.

To check the device configuration status, mouse over the device in Device Manager (you can also check configuration status in Device Monitor). The device status displays as “Managed”, indicating that the device has connected and the management system has successfully imported the device configuration.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See [“Verifying Imported Device Configurations” on page 131](#) for details.

Junos Devices

You can add any device running Junos OS, an EX Series virtual chassis or an SRX virtual chassis to NSM using the static IP address method, so long as the following conditions are met:

- SSH v2 is enabled on the device.
- NETCONF protocol over SSH is enabled on the device.
- The device is configured with a static IP address.
- A user with full administrative privileges is created in the device for the NSM administrator.

To import a device running Junos OS, EX Series virtual chassis or an SRX virtual chassis with a known IP address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device Is Reachable** (default).
5. Click **Next** to open the Specify Connection Settings dialog box.
6. Enter the following connection information:
 - Enter the IP Address of the device.



NOTE:

- Beginning in NSM 2012.2R10, you must select the IP address format (IPv4 or IPv6) of your current IP address.
- IPv6 format is supported only for devices running ScreenOS and SRX Series high-end devices.

- Enter the username of the device administrator.
- Enter the password for the device administrator.



NOTE: All passwords handled by NSM are case-sensitive.

- Select the connection method **SSH version 2**, the only protocol for adding a Junos device to NSM. This connection method is selected by default.
 - Select a port number.
7. Click **Next**. The Verify Device Authenticity dialog box opens. The Add Device wizard displays the RSA Key FingerPrint information. To prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.
 8. Click **Next** to accept the fingerprint. The Detecting Device dialog box opens.
 9. After the wizard displays the autodetected device information, verify that the device type, OS version, and the device serial number are correct. The wizard also detects the Host name configured on the device. You can either use the Host name as the NSM device name or can enter a new name in the text box provided.
 10. Click **Next** to add the device to NSM.

11. After the device is added, click **Next** to import the device configuration.
12. Click **Finish** to complete the Add Device wizard.
13. Double-click the device in Device Manager to view the imported configuration.

To check the device configuration status, mouse over the device in Device Manager or check the configuration status in Device Monitor. The device status displays as **Managed**, indicating that the device has connected and the management system has successfully imported the device configuration.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See [“Verifying Imported Device Configurations” on page 131](#) for details.

SA and IC Devices

You can add any SA or IC device to NSM using the static IP address method, so long as the following conditions are met:

- The inbound DMI connection must be enabled in the device.
- The SSH port must be configured in the device. The default SSH port is 22.
- The DMI agent admin realms must be configured and an admin user must be mapped to a role with full admin privileges.

To import an SA or IC device with a known IP address:

1. From the Configure panel of the NSM main navigation tree, select **Device Manager > Devices**
2. Click the **Device Tree** tab, click the **New** button, and select **Device**. The New Device dialog box appears.
3. Select **Device is Reachable** and click **Next**.
4. Enter the following connection information:
 - Enter the IP address of the device.
 - Enter the username of the device administrator.
 - Enter the password for the device administrator.
 - Select SSH V2 as the connection method.
 - Ensure that the TCP port number is 22.
5. Click **Next**. The Verify Device Authenticity dialog box appears. The Add Device wizard displays the RSA Key FingerPrint information. To prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.
6. Click **Next** to accept the fingerprint. The Detecting Device dialog box opens.

7. After the wizard displays the autodetected device information, verify that the device type, OS version, and the device serial number are correct. The wizard also detects the hostname configured on the device. You can either use the hostname as the NSM device name or you can enter a new name in the text box provided.
8. Click **Next** after verifying the auto detected device information.
9. Click **Finish** to add the device to the NSM UI. The device appears in the Devices workspace.

Adding Devices with Dynamic IP Addresses

A dynamic IP address is an IP address that changes. To add a device that uses a dynamic IP address, the device must support NACN.

ScreenOS Devices

To import a ScreenOS device with an unknown IP address:

1. From the domain menu, select the domain in which you want to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device Is Not Reachable**, and then click **Next**.
5. Enter a name for the device and select a color to represent the device in the UI.
6. From the OS Name list, select **ScreenOS/IDP**. Select the device platform type and the ScreenOS version running on the device from the other pull-down menus. If desired, enable Transparent Mode.
7. Select the license key model for the device. Available selections depend on the type of security device and can include: baseline, advanced, extended, plus, 10-user.
8. Select the Device Server Connection Parameters: Use the default settings to configure the device to connect to the NSM Device Server IP address and port. Use a MIP to configure the device to connect to the NSM Device Server through a mapped IP address and port.
9. Click **Next**, and then perform the following tasks on the Specify One-Time Password screen:
 - a. Make a note of the unique external ID for the device. The device administrator will need it to connect the device to NSM. This ID number represents the device within the management system. The wizard automatically provides this value.
 - b. Specify the First Connection One Time Password (OTP) that authenticates the device.



NOTE: All passwords handled by NSM are case-sensitive.

- c. Click **Show Device Commands** to display the list of CLI commands that must be executed on the device to connect to NSM. The commands enable management and set the management IP address to the Device Server IP address, enable the Management Agent, set the Unique External ID, and set the device OTP.
 - d. Copy and paste these commands into a text file,
 - e. Click **Finish** to complete the Add Device wizard and include the new device in the Device Manager list.
10. Add the commands to the console of the device. Send the commands to the device administrator. The device administrator must make a Telnet connection to the physical device, paste the commands, and execute them to enable NSM management of the device.
 11. To check the device configuration status, mouse over the device in Device Manager or check in Device Monitor.

The status message "Waiting for 1st connect" might appear briefly.

After the device connects, the status displays "Import Needed", indicating that the device has connected but the management system has not imported the device configuration yet.
 12. Import the device configuration by right-clicking the device and selecting **Import Device**. The Job Information box displays the job type and status for the import; when the job status displays successful completion, click **Close**.

After the import finishes, double-click the device to view the imported configuration.



NOTE: After importing a NetScreen-5GT that uses extended port mode, NSM displays the modes as "ns5GT-Trust-Untrust-DMZ" and sets the license mode to Extended.

To check the device configuration status, mouse over the device in Device Manager or check in Device Monitor. The device status displays as "Managed", indicating that the device has connected and the management system has successfully imported the device configuration.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See ["Verifying Imported Device Configurations" on page 131](#) for details.

IDP Sensors

IDP Sensors running versions of the IDP software earlier than 4.0 cannot be imported. If those Sensors are already being managed by IDP Manager, migrate your IDP Management Server and the Sensors with it. See the *IDP-NetScreen-Security Manager Migration Guide* for more information.

You need to upgrade unmanaged Sensors to 4.0 or later before adding them to NSM. See the *IDP Installer's Guide* for more information.

To import an IDP 4.0 device with an unknown IP address, follow these steps:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device Is Not Reachable**, and then click **Next**.
5. Enter the following information on the Specify Name, Color, OS Name, Version, and Platform screen:
 - Enter a name and select a color to represent the device in the UI.
 - From the OS Name list, select **ScreenOS/IDP**.
 - From the Platform Name list, select the device type you want to import.
 - From the OS version list, select the OS version that is running on the device.
6. Click **Next**.

On the Specify One Time Password screen:

- a. Make a note of the unique external ID for the device. The device administrator will need it to connect the device to NSM. This ID number represents the device within the management system. The wizard automatically provides this value.
- b. Specify the First Connection One Time Password (OTP) that authenticates the device.



NOTE: All passwords handled by NSM are case-sensitive.

- c. Click **Show Device Commands** to display the list of CLI commands that must be executed on the device to connect to NSM. The commands enable management and set the management IP address to the Device Server IP address, enable the Management Agent, set the Unique External ID, and set the device OTP.
- d. Copy and paste these commands into a text file,
- e. Click **Finish** to complete the Add Device wizard and include the new device in the Device Manager list.
7. Log into the device as root and run the commands.
8. Check the device configuration status (mouse over the device in Device Manager or check in Device Monitor):
 - The status message “Waiting for 1st connect” might display briefly.
 - After the device connects, the status displays “Import Needed”, indicating that the device has connected but the management system has not imported the device configuration yet.
9. Import the device configuration by right-clicking the device and selecting **Import Device**. The Job Information box displays the job type and status for the import; when the job status displays successful completion, click **Close**.

After the import finishes, double-click the device to view the imported configuration.

To check the device configuration status, mouse over the device in Device Manager (you can also check configuration status in Device Monitor). The device status displays as “Managed”, indicating that the device has connected and the management system has successfully imported the device configuration.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See [“Verifying Imported Device Configurations” on page 131](#) for details.

Adding and Importing an Infranet Controller or Secure Access Device

The procedure for importing a device with an unknown IP address is the same for Secure Access and Infranet Controller devices. This procedure requires specific actions to be taken on the a Secure Access or Infranet Controller device in addition to those taken in the NSM UI.

Because Secure Access and Infranet Controller configurations can be large, some large binary data files are not imported with the rest of the configuration. A stub is placed in the device configuration tree instead. If you need to manage these files in NSM, you must import them later as shared objects, and then create links to those shared objects from the device configuration tree. See [“Managing Large Binary Data Files \(Secure Access and Infranet Controller Devices Only\)” on page 289](#) for details.



NOTE: Secure Access devices are more commonly configured as clusters. See [“Adding Clusters” on page 155](#) for details.

The following sections explain how to add a Secure Access or Infranet Controller device:

- [Install and Configure the Secure Access or Infranet Controller Device on page 124](#)
- [Add the Device in NSM on page 125](#)
- [Configure and Activate the NSM Agent on the Secure Access or Infranet Controller Device on page 126](#)
- [Confirm Connectivity and Import the Device Configuration into NSM on page 127](#)

Install and Configure the Secure Access or Infranet Controller Device

Before you can add a Secure Access or Infranet Controller device to NSM, the device must be installed and configured, and logon credentials for an NSM administrator must be configured for it. Perform the following steps:

1. Select **System > Network > Overview** on the device administrator's console and ensure that basic connection information is configured on the device (network interface settings, DNS settings, and password).
2. Select **Authentication > Auth. Servers** and enter the username and password of the NSM administrator in the applicable authentication server.



NOTE: Only password-based authentication servers can be used. One-time password authentication is not supported.

3. Select **Administrators > Admin Roles** and create an NSM agent role.
4. Select **Administrators > Admin Realms** and create a new NSM agent administrator realm for the NSM agent on the device. Use role mapping to associate the NSM agent role and realm. Do not apply any role or realm restrictions for the NSM agent role or realm.



NOTE: You cannot use NSM to add or manage a Secure Access device that has Host Checker policies enabled for an admin user.

For complete details on installing and configuring Secure Access devices, see the *Secure Access Administration Guide*.

For complete details on installing and configuring Infranet Controller devices, see the *Unified Access Control Administration Guide*.

Add the Device in NSM

To add the device in the NSM UI, follow these steps:

1. From the domain menu, select the domain in which you want to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device is Not Reachable**, and then click **Next**.
5. In the Specify Name, Color, OS Name, Version, and Platform screen:
 - Enter a name and select a color to represent the device in the UI.
 - From the OS Name list, select **SA** for a Secure Access device, or **IC** for an Infranet Controller device,
 - From the Platform list, select the device platform name.
 - From the Managed OS Version list, select the version of the operating system that runs on the device.
 - Select the Device Server Connection Parameters. Use the default settings to configure the device to connect to the NSM Device Server IP address and port. Use

a MIP to configure the device to connect to the NSM Device Server through a mapped IP address and port.

- Click **Next**.
6. Perform the following tasks on the Specify Device Admin User Name/Password and One Time Password screen:
 - a. Make a note of the unique external ID. The device administrator will need it to configure connectivity with NSM. The wizard automatically enters this value for the device. This ID number represents the device within the management system.
 - b. Specify the administrator user name and password for the SSH connection. This name and password must match the name and password already configured on the device.
 - c. Specify the First Connection One Time Password (OTP) that authenticates the device. Make a note of this password. The device administrator will need it to configure the connectivity with NSM.



NOTE: All passwords handled by NSM are case-sensitive.

- d. Click **Finish** to complete the Add Device wizard and include the new device in the Device Manager list.
7. Verify in the Device List tab that the new device is visible and has the connection status "Never connected."
8. Convey the unique external ID and the one-time password to the device manager.

Configure and Activate the NSM Agent on the Secure Access or Infranet Controller Device

Configure and activate the NSM agent on the Secure Access or Infranet Controller device to establish SSH communications with NSM. On successful execution of these steps, you can control the device from NSM,

1. Open the **System > Configuration > NSM Agent** screen to add the NSM management application.
2. In the Primary Server field, enter the IP address of the Device Server.
3. In the Primary Port field, enter **7804**.
4. Fill out the Backup Server and Backup Port fields if a high availability Device Server is configured.
5. In the Device ID field, enter the unique external ID provided by the NSM administrator.

6. In the HMAC field, enter the one-time password, also provided by the NSM administrator.

7. Click the **Enable** button to enable the NSM agent.

8. Click **Save Changes**, and the device attempts to establish a session with NSM.

The device software initiates the TCP connection to NSM and identifies itself using the specified device ID and HMAC. The two sides then engage in SSH transport layer interactions to set up an encrypted tunnel. NSM authenticates itself to the device based on user name and password.

Confirm Connectivity and Import the Device Configuration into NSM

In NSM, validate connectivity with the device, and then import the device configuration:

1. In the Device List, check the connection status of the newly added device. The status changes from “Never connected” to “Up.”

- If the configuration status is “platform mismatch,” you selected the wrong device platform when adding the device into NSM. Delete the device from NSM and add it again using the correct device platform.
- If the configuration status shows “device firmware mismatch,” you selected the wrong managed OS version when adding the device into NSM. Delete the device from NSM and add it again using the correct managed OS version.

2. Import the device configuration:

- a. Right-click the device in the Device Manager and select **Import Device** from the list.
- b. In the Device Import Options dialog, check **Summarize Delta Config** if desired. Click **OK**, and then click **Yes**.

The Job Information window shows progress. You can also monitor progress in Job Manager.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See [“Verifying Imported Device Configurations” on page 131](#) for details.

Adding and Importing a Junos Device with a Dynamic IP Address

The procedure for importing a device with an unknown IP address is similar for all Junos devices. This procedure requires specific actions to be taken on the device in addition to those taken in the NSM UI.

The following sections explain how to add a Junos device:

- [Install and Configure a Junos Device on page 128](#)
- [Add the Device in NSM on page 128](#)

- [Configure and Activate Connectivity on a Junos Device on page 129](#)
- [Confirm Connectivity and Import the Device Configuration into NSM on page 130](#)

Install and Configure a Junos Device

Before you can add a Junos device to NSM, the device must be installed and configured, and logon credentials for an NSM administrator must be configured for it. Perform the following steps:

1. Connect the device to the network and configure one of the interfaces so that the device can reach the NSM device server.
2. Add a user for NSM that has full administrative rights.

For complete details on installing and configuring Junos devices, see the documentation for the specific device,

Add the Device in NSM

To add the device in the NSM UI, follow these steps:

1. From the domain menu, select the domain in which you want to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device Is Not Reachable**, and then click **Next**.
5. On the Specify Name, Color, OS Name, Version, and Platform screen:
 - Enter a name and select a color to represent the device in the UI.
 - From the OS Name list, select **Junos**.
The Junos OS Type list appears.
 - Select the Junos OS type for the device you want to add:
 - To add a J Series or SRX Series device or an SRX virtual chassis, select **J/SRX Series** from the list.
 - To add an EX Series device or virtual chassis, select **EX Series**.
 - To add an M Series or MX Series device, select **M/MX Series**.
 - From the Platform list, select the device platform name.
 - Check the **Virtual Chassis** box if you are adding an EX Series virtual chassis made up of several EX Series switches (EX4200 series only) or an SRX virtual chassis..
 - From the Managed OS Version list, select the version of the operating system that runs on the device.

- Select the Device Server Connection Parameters: Use the default settings to configure the device to connect to the NSM Device Server IP address and port. Use a MIP to configure the device to connect to the NSM Device Server through a mapped IP address and port.
 - Click **Next**.
6. On the Specify Device Admin User Name/Password and One Time Password screen, perform the following tasks:
 - a. Make a note of the automatically generated Unique External ID for the device. This ID number represents the device within the management system. The device administrator will need this ID to configure connectivity between the device and NSM.
 - b. Specify the administrator user name and password for the SSH connection. This name and password must match the name and password already configured on the device.
 - c. Specify the First Connection One Time Password (OTP) that authenticates the device. Make a note of this password. The device administrator will need it to configure the connectivity between the device and NSM.



NOTE: All passwords handled by NSM are case-sensitive.

- d. Click **Finish** to complete the Add Device wizard and include the new device in the Device Manager list.
7. Verify in the Device List tab that the new device is visible and has the connection status "Never connected."
8. Give the unique external ID and the one-time password to the device manager.

Configure and Activate Connectivity on a Junos Device

The device administrator must identify the device to the management system and initiate the connectivity:

1. Log on to the Junos device.
2. At the command-line prompt, identify the management system by device name, device ID, and HMAC.

For devices running the 9.0 version of the operating system, use the following command syntax:

```
set system services outbound-ssh application-id <name> secret <string> services
netconf device-id <external-id from nsm> <nsm device server ip> port 7804
```

For example:

```
% set system services outbound-ssh application-id nsm-wei secret 123456789
services netconf device-id abcdef 10.150.42.16 port 7804
```

For devices running the 9.1 and later versions of the operating system, use the following command syntax:

```
set system services outbound-ssh client <name> secret <secret string>
services netconf device-id <external-id from nsm> <nsm device server ip> port
7804
```

For example:

```
set system services outbound-ssh client nsm-wei secret 123456789
services netconf device-id abcdef 10.150.42.16 port 7804
```

3. Establish the SSH connection with the network management system.

- For a gateway or router with a single Routing Engine, or for a single EX Series switch:

```
% commit
```

- For an EX Series virtual chassis, an SRX virtual chassis, or for a router with redundant Routing Engines:

```
% commit synchronize
```

Synchronizing the commit operation ensures that NSM connects to the backup Routing Engine following failover of the master Routing Engine.

The device software initiates the TCP connection to NSM and identifies itself using the specified device ID and HMAC. The two sides then engage in SSH transport layer interactions to set up an encrypted tunnel. NSM authenticates itself to the device based on user name and password.

Confirm Connectivity and Import the Device Configuration into NSM

In NSM, validate connectivity with the device, and then import the device configuration:

1. In the Device List, verify the connection status of the newly added device. The status changes from “Never connected” to “Up.”
 - If the configuration status is “device platform mismatch,” you selected the wrong device platform when adding the device into NSM. Delete the device from NSM and add it again using the correct device platform.
 - If the configuration status shows “device firmware mismatch,” you selected the wrong managed OS version when adding the device into NSM. Delete the device from NSM and add it again using the correct managed OS version.

2. Import the device configuration:

- a. Right-click the device in the Device Manager and select **Import Device** from the list.
- b. In the Device Import Options dialog, check **Summarize Delta Config** if desired. Click **OK**, and then click **Yes**.

The Job Information dialog shows progress, or you can monitor progress in Job Manager.

The next step is to verify the imported configuration using the Device Monitor or the Device Manager. See [“Verifying Imported Device Configurations” on page 131](#) for details.

Verifying Imported Device Configurations

After importing a device, verify that all device information was imported as you expected.

Using Device Monitor

The Device Monitor tracks the status of individual devices, systems, and their processes. After you import a device, check the status of that device in Device Monitor, located in Realtime Monitor.

The imported device should display a configured status of “Managed” and a Connection status of “Up”, indicating that the device has connected and the management system has successfully imported the device configuration.

Using Device Manager

In the security device tree, ensure that the device exists. Open the device configuration and check the following values:

- Ensure that the imported device serial number matches the serial number on the physical device.
- Ensure that the imported device IP address matches the IP address for the physical device.
- Ensure that imported device administrator name and password are correct for the physical device.



NOTE: All passwords handled by NSM are case-sensitive.

- Ensure that interfaces on the imported device are correct for the physical device.



NOTE: When importing a NetScreen-500, 5000 series, or ISG series security device, you must manually configure the network module (slot) before the imported physical interfaces appear in the NSM UI. For details on defining the Ethernet card and slot, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

- Browse the device configuration tree and ensure that the management system successfully imported all device configuration information, including zones, virtual routers, and routes.

Using Job Manager

Job Manager tracks the status of major administrative tasks, such as importing or updating a device. After you import a device, view the report for the import task to ensure that the management system imported the device configuration as you expected.



NOTE: Job Manager configuration summaries and job information details do not display passwords in the list of CLI commands for administrators that do not have the assigned activity “View Device Passwords”. By default, only the super administrator has this assigned activity.

Job Manager also tracks the status of configuration summaries, described in the following sections.

Using Configuration Summaries

NSM provides three configuration summaries to help you manage device configurations and prevent accidental misconfiguration. Use configuration summaries after you import a device to ensure that the management system imported the physical device configuration as you expected.

Configuration summaries help with ongoing device maintenance, particularly for devices on which a local device administrator has been troubleshooting using CLI commands or the Web UI. Because the device object configuration in the NSM UI can overwrite the physical device configuration, you should always confirm the commands that are sent to the device.

Configuration Summary

A configuration summary shows you the exact CLI commands that will be sent to the managed device during the next device update. To get a Configuration Summary, from the Device Manager launchpad, click **Device Config Options > Summarize Config**. You see a list of security devices to which you have access. Select the device you just imported and click **OK**. NSM analyzes the UI device object configuration and generates a summary report that lists the CLI commands or XML messages to send to the physical device during the next device update.

For a just-imported device, the configuration summary report displays the device configuration that matches the configuration currently running on the physical device.

Delta Configuration Summary (All Devices Except IDP)

A delta configuration summary shows you the differences between the configuration you see in the NSM UI and the configuration on the physical device. To get a delta configuration summary, from the Device Manager launchpad, click **Summarize Delta Config**. You see a list of devices to which you have access. Select the device you just imported and click **Apply Changes**. NSM queries the physical device to obtain a list of all CLI commands or XML messages used in the device configuration, compares that list with the UI device configuration, and generates a summary report of all differences, or deltas, discovered.

For a just-imported device, the delta config summary displays minimal deltas, meaning that very few differences exist between the configuration on the physical device and the configuration in the UI. NSM automatically imports your VPNs and displays the VPN policies; however, NSM does not create VPN abstractions for your VPN policies.

Get Running Configuration

A running configuration summary shows you the exact CLI commands or XML messages that were used to create the current device configuration on the physical device. To get the Running Config summary, from the Device Manager launchpad, click **Device Config Options > Get Running Config**. You see a list of devices to which you have access. Select the device you just imported and click **OK**. NSM queries the physical device to obtain a list of all CLI commands used in the device configuration and generates a summary report that lists those commands.

For a just-imported device, the get running config summary report displays the device configuration currently running on the physical device.

Modeling Devices

For most undeployed device types, you can create a device configuration in NSM, and then install that device configuration on the physical device. For ScreenOS devices, you can use Rapid Deployment (RD) to quickly provision multiple devices in nontechnical environments. See [“Using Rapid Deployment \(ScreenOS Only\)” on page 142](#) for details.

You cannot model and activate a Secure Access or Infranet Controller device. You must import these devices to add them to NSM.

Adding a single undeployed device to NSM is a four-stage process:

1. Model the device in the UI.
2. Create the device object configuration.
3. Activate the device.
4. Update the device configuration.

For details on modeling multiple devices at one time, see [“Adding Many Devices Using CSV Files” on page 181](#).

Requirements

To model a device, you must know the device type and OS name and version that is running on the device.

To activate a device:

- You must have the device connection information and device administrator name and password.



NOTE: All passwords handled by NSM are case-sensitive.

- The device must be staged; that is, it must be physically connected to your network and able to access network resources.
- The device must have at least one interface that has an IP address. Devices that use dynamically assigned IP address must also support NACN.

Modeling a Device

To add and model a device:

1. From the domain menu, select the domain in which you want to model the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon, and then select **Device**. The device wizard appears.
4. Select **Model Device**, and then click **Next**.
5. In the Specify Name, Color, OS Name, Version, and Platform screen, enter the following information:
 - Enter a name and select a color to represent the device in the UI.
 - In the OS Name list, select the device family that the modeled device belongs to.
 - For Junos devices, select the Junos OS type:
 - To model a J Series or SRX Series device, or an SRX virtual chassis, select **J/SRX Series** from the list.
 - To model an EX Series device or virtual chassis, select **EX Series**.
 - To model an M Series or MX Series device, select **M/MX Series**.
 - In the platform list, select the device platform name.
 - In the OS version list, select the version of the operating system or firmware that runs on the device.
6. For EX Series switches, check **Virtual Chassis** if you wish to model a virtual chassis (an array of EX4200 series switches). For an SRX virtual chassis check **Virtual Chassis** if you wish to model an SRX virtual chassis.

7. When adding a ScreenOS device that uses port modes, select the appropriate port mode from the Device subtype list, after you select the device type. NSM automatically sets the license mode to Extended.
8. Enable transparent mode, if desired (ScreenOS devices only).



NOTE: You cannot change the operational mode after the device has been modeled.

9. Click **Finish** to complete the Add Device wizard. The UI creates a corresponding device object that appears in the Device Manager list.
10. Hold your mouse cursor over the device in Device Manager to check the device configuration status, or check the configuration status in Device Monitor. The status displays “Modeled”, indicating that the management system has modeled the device, but the device is not activated and has not connected.

Creating a Device Configuration

Because undeployed devices are devices that you are not currently using in your network, they might not have a preexisting device configuration (IP addresses, zones, and interfaces) that is available for import. You can create a configuration for the device object in NSM, and then install that configuration on the device.



NOTE: When modeling a NetScreen-500, 5000 series, or ISG series security device, you must configure the network module (slot) before physical interfaces appear in the NSM UI. For details on defining the Ethernet card and slot, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Double-click the device object to display the device configuration and begin configuring the device as desired. For details on device configuration, see [“Configuring Devices” on page 199](#).

Activating a Device

After you have created a device configuration for the undeployed device, you are ready to activate the device and prompt it to connect to the management system. After that device has made contact with NSM, you can install the configuration you created on the device.

Devices with Static IP Addresses

A static IP address is an IP address that does not change.

ScreenOS Devices

To activate a ScreenOS 5.0 or later device with a static IP address:

1. Check the device configuration state by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. The device configuration state should display “Modeled”, indicating that the management system is waiting for the device to connect.
2. Right-click the device and select **Activate Device** to display the Activate Device wizard.
3. Select **Device Deployed and IP is reachable**.
4. Click **Next** and enter the connection information:
 - Enter the IP Address of the security device.
 - Enter the device administrator name and password.



NOTE: All passwords handled by NSM are case-sensitive.

- Select the connection method (Telnet, SSH version 1, SSH version 2) and the port number for the selected service.

If you selected Telnet, click **Next** and go to Step 5.

If you selected an SSH version, click **Next**. The Verify Device Authenticity dialog box opens. The device wizard displays the RSA Key FingerPrint information. To prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.

5. After NSM autodetects the device, click **Next** to activate the device in NSM.
6. Click **Update Now** to update the configuration on the device with the settings from the modelled device.

If you do not update the configuration now, you will have to do it manually later by right-clicking the device and selecting **Update Device**.

The Job Information box displays the job type and status for the update. When the job status displays successful completion, click **Close**.

After the update finishes, the device status displays as “Managed”, indicating that the device has connected and the management system has successfully pushed the device configuration.

IDP Sensors

To activate an IDP Sensor with a static IP address:

1. Check the device configuration status by holding your mouse cursor over the device in Device Manager (you can also check configuration status in Device Monitor). The device configuration status should display “Modeled”, indicating that the management system is waiting for the device to be activated.
2. Right-click the device and select **Activate Device** to display the Activate Device wizard.
3. Select **Device deployed and IP is reachable**.

4. Click **Next**. The Specify Connection Settings dialog box opens. Enter the connection information:

- Enter the IP Address of the device.
- Enter the device administrator name.
- Enter the device administrator password.
- Enter the device root password.



NOTE: All passwords handled by NSM are case-sensitive.

- Set the connection method to SSH Version 2.
5. Click **Next** to display the Verify Device Authenticity dialog box.
 6. Click **Next**.
 7. After NSM autodetects the device, click **Next** to activate the device in NSM.
 8. Click **Update Now** to update the configuration on the device with the settings from the modelled device.

If you do not update the configuration now, you will have to do it manually later by right-clicking the device and selecting **Update Device**.

Updating the device also pushes the Juniper Networks Recommended policy to the device.

After update is complete, the device status displays as “Managed”, indicating that the device has connected and the management system has successfully pushed the device configuration.

Junos Devices

To activate a Junos device (or EX Series virtual chassis) with a static IP address:

1. Check the device configuration state by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. The device configuration state should display “Modeled”, indicating that the management system is waiting for the device to connect.
2. Right-click the device and select **Activate Device**. The Activate Device wizard begins.
3. Select **Device deployed and IP is reachable**.
4. Click **Next** and enter the connection information:

- Enter the IP Address of the security device.

For a Junos device with redundant Routing Engines, provide the IP address of the master Routing Engine.

- Enter the device administrator name and password.



NOTE: All passwords handled by NSM are case-sensitive.

5. Click **Next**. The Verify Device Authenticity dialog box opens. The device wizard displays the RSA Key FingerPrint information. To prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.
6. Click **Next** to accept the fingerprint. The Detecting Device dialog box opens.
7. After the wizard displays the autodetected device information, verify that the device type, OS version, and the device serial number are correct.
8. Click **Next** to activate the device in NSM.
9. Click **Update Now** to update the configuration on the device with the settings from the modelled device, or click **Exit** to leave the wizard without updating the device.

If you do not update the configuration now, you will have to do it manually later by right-clicking the device and selecting **Update Device**.

The Job Information box displays the job type and status for the update. When the job status displays successful completion, click **Close**.

After the update is complete, the device status displays as “Managed”, indicating that the device has connected and the management system has successfully pushed the device configuration.

Devices with Dynamic IP Addresses

A dynamic IP address is an IP address that changes. To add a device that uses a dynamic IP address, the device must support NACN.

ScreenOS Devices and IDP Sensors

To activate a ScreenOS device or an IDP sensor with an unknown IP address:

1. Check the device configuration state by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. The device configuration state should display “Modeled”, indicating that the management system is waiting for the device to connect.
2. Right-click the device and select **Activate Device**. The Activate Device wizard displays.
3. Select **Device deployed, but IP is not reachable**.
4. Click **Next**. The Specify the connections settings dialog box opens.
5. Specify the First Connection One Time Password (OTP) that authenticates the device.



NOTE: All passwords handled by NSM are case-sensitive.

6. Edit the Device Server Connection parameters, if desired.
7. Click **Next**. The Specify device connections characteristics dialog box opens.

Click **Show Device Commands** to display a list of CLI commands. The commands enable management and set the management IP address to the Device Server IP address, enable the Management Agent, set the Unique External ID, and set the device OTP.

Copy and paste these commands into a text file, and then send the commands to the device administrator. The device administrator must make a Telnet connection to the physical device, paste the commands, and execute them to enable NSM management of the device.



NOTE: The device administrator can also use a console connection to execute the commands on the physical device. However, the commands must be entered three at a time to ensure that the device receives all commands.


8. Click **OK** to dismiss the Commands window and complete the Activate Device wizard.
9. Check the device configuration status by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor). When the device connects, the status displays “Update Needed”, indicating that the device has connected but the management system has not pushed the device configuration yet.
10. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update. When the job status displays successful completion, click **Close**.

After the update is complete, the device status displays as “Managed”, indicating that the device has connected and the management system has successfully updated the device configuration.

Activating a Junos Device

The procedure for activating a device with an unknown IP address is similar for all Junos devices. This procedure requires specific actions to be taken on the device in addition to those taken in the NSM UI. In NSM, you must complete the Activate Device workflow. You must also make sure the device is properly installed and configured with login credentials for the NSM administrator, and configure the device to connect with NSM. Finally, you must load the modeled configuration onto the device:

1. Install the device and configure it with logon credentials for the NSM administrator:
 - a. Connect the device to the network and configure one of the interfaces so that the device can reach the NSM device server.

- b. Add a user for NSM that has full administrative rights.
 2. Activate the device in NSM:
 - a. In Device Manager, right-click the device and then select **Activate Device** from the list.
 - b. In the Activate Device dialog box, select **Device is deployed, but not reachable**, and then click **Next**.
 - c. In the Specify Name and Device type dialog box, perform the following tasks:
 - i. Make a note of the automatically generated Unique External ID for the device. This ID number represents the device within the management system. The device administrator will need this ID to configure connectivity between the device and NSM.
 - ii. Specify the administrator user name and password for the SSH connection. This name and password must match the name and password already configured on the device.
 - iii. Specify the First Connection One Time Password (OTP) that authenticates the device. Make a note of this password. The device administrator will need it to configure the connectivity between the device and NSM.
- **NOTE:** All passwords handled by NSM are case-sensitive.
- d. Click **Finish** to complete the Add Device wizard and include the new device in the Device Manager list.
 - e. Verify in the Device List tab that the new device is visible and has the connection status "Never connected."
 - f. Give the unique external ID and the one-time password to the device manager.
 3. On the device, configure the device to connect with NSM:
 - a. Log on to the Junos device.
 - b. At the command-line prompt, identify the management system by device name, device ID, and HMAC.

For devices running the 9.0 version of the operating system, use the following command syntax:

```
set system services outbound-ssh application-id <name> secret <string>
services netconf device-id <external-id from nsm> <NSM device server ip>
port 7804
```

For example:

```
% set system services outbound-ssh application-id nsm-wei secret 123456789
services netconf device-id abcdef 10.150.42.16 port 7804
```

For devices running the 9.1 and later versions of the operating system, use the following command syntax:

```
set system services outbound-ssh client <name> secret <secret string>
services netconf device-id <external-id from nsm> <nsm device server ip>
port 7804
```

For example:

```
set system services outbound-ssh client nsm-wei secret 123456789
services netconf device-id abcdef 10.150.42.16 port 7804
```

c. Establish the SSH connection with the network management system.

- For a gateway or router with a single Routing Engine, or for a single EX Series switch:

```
# commit
```

- For an EX Series virtual chassis, an SRX virtual chassis or for a gateway or router with redundant Routing Engines:

```
# commit synchronize
```

Synchronizing the commit operation ensures that NSM connects to the backup Routing Engine following failover of the master Routing Engine.

The device software initiates the TCP connection to NSM and identifies itself using the specified device ID and HMAC. The two sides then engage in SSH transport layer interactions to set up an encrypted tunnel, and NSM authenticates itself to the device based on user name and password.

d. In the Device List, verify the connection status of the newly added device. The status changes from “Never connected” to “Up.”

If the configuration status is “platform mismatch,” you selected the wrong device platform when adding the device into NSM. Delete the device from NSM and add it again using the correct device platform.

If the configuration status shows “device firmware mismatch,” you selected the wrong managed OS version when adding the device into NSM. Delete the device from NSM and add it again using the correct managed OS version.

4. In NSM, validate connectivity with the device, and then update the device configuration:
 - a. Check the device configuration status by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. When the device connects, the status displays “Update Needed”, indicating that the device has connected but the management system has not yet pushed the device configuration.
 - b. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update. When the job status displays successful completion, click **Close**.

Using Rapid Deployment (ScreenOS Only)

Rapid Deployment (RD) enables deployment of multiple security devices in a large network environment with minimal user involvement. RD is designed to:

- Simplify the deployment of firewall devices in nontechnical environments.
- Minimize device staging or technical staff required at the deployment site.
- Enable secure and efficient deployment of a large number of firewalls.
- Bring new security devices under NSM management for initial configuration.

RD is supported on the following security devices:

ns204	ns5GTadslwan-Home-Work
ns208	ns5GTadslwan-Trust-Untrust
ns25	ns5GTwan-Combined
ns50	ns5GTwan-Dmz-Dual-Untrust
ns5GT-Combined	ns5GTwan-Dual-Untrust
ns5GT-Dmz-Dual-Untrust	ns5GTwan-Extended
ns5GT-Dual-Untrust	ns5GTwan-Home-Work
ns5GT-Extended	ns5GTwan-Trust-Untrust, ns5XP
ns5GT-Home-Work	ns5XT-Combined
ns5GT-Trust-Untrust	ns5XT-Dual-Untrust
ns5GTadsl-Extended	ns5XT-Home-Work
ns5GTadsl-Home-Work	ns5XT-Trust-Untrust
ns5GTadsl-Trust-Untrust	nsHSC-Home-Work
ns5GTadslwan-Extended	nsHSC-Trust-Untrust

RD typically involves two people: The NSM administrator, who creates the necessary device configuration for the new firewall devices in the NSM UI, and the onsite administrator, who enables the firewall device to contact NSM for configuration.

The NSM administrator performs the following tasks in the NSM UI:

1. Adds a device to the UI.
2. Creates a device configuration with specific or template-driven values.
3. Enters the basic information that defines how a security device can contact your NSM Device Server.
4. Generates a small, static command file called a *configlet*.
5. Saves the configlet in a user-defined directory, using email, CD, or another out-of-band method.
6. Sends the configlet file to the onsite administrator, who installs the configlet on the security device at its physical location.
7. After the onsite administrator installs the configlet and the device has successfully connected to the management system, the NSM administrator installs the modeled device configuration on the physical device.

The onsite administrator works locally, at the physical device and performs the following tasks:

1. Installs the configlet on a locally connected device.
2. Runs the Rapid Deployment Wizard.

The RD wizard uses the information in the configlet to establish and authenticate a secure connection to the NSM Device Server, enabling NSM to begin managing the device.



NOTE: You cannot activate devices in transparent mode using a configlet.

After the security device has connected to NSM, the NSM administrator can manage the device exactly like any other security device in NSM.



NOTE: If you delete the security device from the NSM system and then add the device again, you must also re-create the configlet and install it on the physical device.

To use rapid deployment:

- The device must be running ScreenOS 5.x or later release.
- The device must use default factory settings.
- The device must be able to reach the Internet using a static IP address, an IP address assigned with PPPoE, PPPoA, or DHCP.
- The device must be modeled in NSM system. For details on modeling a device, see [“Modeling a Device” on page 134](#).

After you have modeled the device in the management system, you can track its status using the Device Monitor. Check the device configuration status by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. The status should display “Modeled”, indicating that the management system has modeled the device, but the device is not activated and has not connected.

The following sections provide details about each stage. For details about modeling, creating configlets for multiple devices, and activating multiple devices at one time, see [“Adding Many Devices Using CSV Files” on page 181](#).

- [Creating the Configlet on page 144](#)
- [Installing the Configlet on page 147](#)
- [Updating the Device Configuration on page 149](#)

Creating the Configlet

After you have created a device configuration for the undeployed device, you are ready to activate the device and create the configlet.

1. Right-click the device and select **Activate Device**. The Activate Device wizard appears.
2. Select **Device Deployed, but IP is not Reachable**, and then click **Next**.
3. Select **Initialize Device Using Configlet**.
4. Click **Next**.
 - Specify the First Connection One Time Password that authenticates the device.



NOTE: All passwords handled by NSM are case-sensitive.

- The wizard automatically selects the interface on the device that will connect to the NSM management system. This interface is determined by the device platform and cannot be changed.
 - Select the **Device Server** connection. Use the default settings to configure the device to connect to the NSM Device Server IP address and port. Use a MIP to configure the device to connect to the NSM Device Server through a mapped IP address and port.
 - Click **Next**.
5. Specify the connection setting on the device:

- For devices with static IPs, you can predefine the IP address, mask, and gateway, or ask the onsite administrator to specify this information during configlet installation.
- For devices that use DHCP, the configlet automatically handles IP assignment during installation.
- For devices that use a PPPoE connection to the Internet, you can predefine the user name and password, or ask the onsite administrator to specify the user name and password during configlet installation.



NOTE: All passwords handled by NSM are case-sensitive.

- For devices that use a PPPoA connection to the Internet, you can predefine the following ADSL parameters:
 - VPI/VCI Pair. The Virtual Path Identifier and Virtual Channel Identifier (VPI/VCI) identify the virtual circuit on the DSLAM.
 - Multiplexing Mode, also known as the ATM encapsulation method. The multiplexing mode defines how the ADSL interface handles the multiple protocols on the virtual circuit. Your service provider must tell you the type of multiplexing used on the ADSL line.

Virtual Circuit (VC)-based multiplexing carries each protocol over a separate ATM virtual circuit.

Logical Link Control (LLC) carries several protocols to be carried on the same ATM virtual circuit. This option is the default for the ADSL1 interface on the NetScreen-5GTADSL security device.

- RFC1483 Protocol Mode. RFC 1483 describes methods of transporting bridged or routed protocol data units (PDUs) over AAL5 links.

Bridged PDUs do not require the overhead of IPsec processing, thus allowing more usable bandwidth to be available for data traffic. Because non-IPsec traffic is not secured at the IP packet layer, you should use this mode only with a private virtual circuit (the service provider assigns a static IP address for the ADSL interface).

Routed PDUs enable the NetScreen-5GT ADSL device to exchange routing information with another router through the ADSL interface.

- ADSL Operating Mode. The operating mode defines the physical line attributes for the ADSL interface.

Auto Detect (default mode) enables the ADSL interface to automatically negotiate the operating mode with the service provider DSLAM.

ANSI T1.413 Issue 2 Mode

ITU G.992.1 Mode enables the ADSL interface to use the International Telecommunications Union (ITU) G.dmt standard, which supports minimum data rates of 6.144 Mbps downstream and 640 Kbps upstream.

G.Lite Mode enables the ADSL interface to use the ITU 992.2 standard, which supports maximum data rates of 1.536 Mbps downstream and 512 Kbps upstream.

Alternatively, you can ask the onsite administrator to specify these parameters during configlet installation.

- If you don't know the ISP environment or the environment has location-specific networking requirements, prompt the onsite administrator to configure the ISP environment during configlet installation.
6. Specify the password for the configlet, or use the default device password (which is **netscreen**).
 7. Specify Device User Name and password, or use the default administrator name and passwords for the device.
 8. Check **Restrict the use of the configlet to the current device** to install the configlet only on a device with the specified serial number.
 9. Click **Next** to display the decoded configlet. To see the encoded configlet, click the **Raw Configlet** tab.
 10. Click **Save** to save the configlet (configlet files automatically use the format .cfg).



NOTE: You cannot edit a configlet file directly. To make changes to the information in a configlet file, run the Activate Device wizard to regenerate the configlet.

11. Click **Finish** to close the Activate Device wizard.
12. Send the configlet to the onsite administrator using email, CD, or another out-of-band method.

The onsite administrator must complete the configlet installation process and the device must successfully connect to the management system before you can update the device with the modeled configuration.

For help with the configlet installation process, the onsite administrator can refer to the *Rapid Deployment Getting Started Guide*. This guide provides step-by-step instructions for connecting a security device to the network, preparing the device to use a configlet, and installing and running the configlet.

You can track the connection status of the device to determine when the device connects. Check the device configuration status by holding your mouse cursor over the device in Device Manager, or by checking configuration status in Device Monitor.

- Before the device connects, the status displays “Waiting for 1st Connect”, indicating that the management system has modeled and activated the device, but the device has not connected.
- After the onsite administrator has installed the configlet, the device automatically connects to the management system and the status displays “Update Needed”,

indicating that the device has connected but the management system has not yet installed the modeled device configuration.

Installing the Configlet

The onsite administrator performs RD in two stages:

- Preparing the security device
- Installing the configlet

The following sections detail each stage. For detailed, step-by-step instructions on installing the configlet, see the *Rapid Deployment Getting Started Guide*.

Preparing the Device

Before you install the configlet, you must prepare the security device:

1. Connect the device to your network. For details on connecting the device, see the user's guide that came with your security device.
2. Connect a standalone computer, such as a laptop, to the device's eth1 port.
 - To connect directly to the device, use a crossover cable.
 - To connect to the device over a hub or switch, use a straight-through cable.

If your device has auto-sensing ports, you can use any type of Ethernet cable to connect to the device.

3. Change the IP address of the standalone computer to **192.168.1.2** and the default gateway to **192.168.1.1**. To change an IP address, see your computer's operating system documentation.
4. Ensure that the device is using the factory default settings.
 - RD works with the factory default settings of all security devices running ScreenOS 5.x or a later release. If the device does not use the factory default settings, you cannot use RD (the WebUI cannot load the configlet).
 - To restore the factory defaults on the firewall device, see the user's guide that came with your security device.
5. Ensure that the Status LED on firewall device displays green.

Installing the Configlet



NOTE: During the configlet installation process, you cannot edit the device configuration.

To install the configlet:

1. Save the configlet on the standalone computer that you connected to the security device.
2. In a Web browser, enter the IP address of the trust interface on the security device as **192.168.1.1**. The Rapid Deployment Wizard appears.
3. Select **Load configlet** file and browse to the location of the saved configlet file. Click **Next**.

The RD Wizard opens the configlet, authenticates the integrity of the configlet, and decrypts the configlet. If the configlet is valid, the RD Wizard uses the configlet information to prepare the security device for NSM management.

4. If prompted, enter the configlet password and click **Next**. The configlet password is given to you by the NSM administrator who sent you the configlet file. Click **Next**.
5. Confirm or enter the ISP information. The ISP information describes the ISP environment in which the device is deployed. If the NSM administrator included ISP information in the configlet, the RD Wizard displays that information. Ensure that all information is correct.

If the NSM administrator did not include ISP information or included only partial information, you must complete the ISP environment for the device:

- If your firewall device uses DHCP to obtain an IP address from the network, select **Using cable modem (Dynamic IP via DHCP)**.
 - If your firewall device uses a PPPoE connection to the Internet, select **Using DSL modem (Dynamic IP via PPPoE)**. Enter the username and password for your PPPoE account.
 - If your firewall device uses a static IP address, select **Using ISP-supplied Settings (Static IP)** and enter the IP address, Netmask, and Gateway for the firewall device.
 - If your security device uses a PPPoA connection to the Internet (available on NetScreen-5GT ADSL devices), select **PPPoA**. Enter the multiplexing mode, VCI/VPI pair, Multiplexing mode, RFC1483 Protocol mode, and the ADSL operating mode for your PPPoA account.
6. Click **Next** to initiate the connection to NSM.

The security device connects to the NSM Device Server. During this first connection, the device and the NSM Device Server exchange authentication information. After NSM authenticates the connection and saves the device public key, it sends a confirmation message to the device, which displays the message in the RD Wizard.



NOTE: For security, after the first successful connection, the security device erases the one-time-password (OTP) from memory.

7. Click **Close** to exit the RD Wizard.

The NSM administrator can now configure the device using NSM.



NOTE: If the configlet installation process fails, you must reset the device to factory defaults. For details, see the user's guide that came with the security device.

Updating the Device Configuration

After the onsite administrator has installed the configlet, the device has successfully connected to the management system, and the NSM administrator has modeled the device configuration, you can install the modeled device configuration on the physical device:

1. Ensure that the device is connected by viewing the device status. Check the device configuration status by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. Ensure that the configuration status for the device displays "Update Needed", which indicates that the device has connected but the management system has not updated the device configuration yet.
2. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update; when the job status displays successful completion, click **Close**.

After update is complete, the device status displays as "Managed", indicating that the device has connected and the management system has successfully updated the device configuration.

Summarize Delta Configuration

NSM allows you to perform the command **Summarize Delta Config** on a device before you update the device. You can cancel the **Update Device** directive as well as save the **Summarize Config** output. The **Update Device** has the following two phases: Summarize Delta Config and Update Device.

Users can change the default to combine these two phases so that the delta configuration summary is automatically performed before the device is updated and the results are used to update the device as an optimization. During this combined operation, both results (Delta Config and Update Device) are available to you by selecting **View Device Delta Config**, if you have the appropriate administrator rights. Otherwise, you can still update the device, but you cannot run **Summarize Delta Config**.

Example: User Successfully Selects and Updates Two Devices with Delta Option

1. In the main navigation tree, select **Device Manager > Devices**.
2. From the Device Manager launchpad, select **Update Device** to open the Update Device(s) dialog box, listing all connected and managed devices.
3. Select two devices you want to update.
4. Select **Run Summarize Delta Config** (if deselected), and then click **Apply Changes**.

NSM displays the delta configuration results for both devices.

5. Click **Update**.
6. Close the Job Information window and select **Job Manager** from the main navigation tree.
7. Select **Update Device** to see the update device job results for both devices.

Example: User Selects Two Devices with Delta Option and One Device Fails

1. In the main navigation tree, select **Device Manager > Devices**.
2. From the Device Manager launchpad, select **Update Device** to open the Update Device(s) dialog box, listing all connected and managed devices.
3. Select two devices you want to update.
4. Select **Run Summarize Delta Config** (if deselected), and then click **Apply Changes**.
NSM displays the delta configuration results. One device succeeded and the other device failed.
5. Click **Update**.
6. Close the Job Information window and select **Job Manager** from the main navigation tree.

Only the device that passed delta configuration is updated.

Example: User Selects Two Devices to Update Without the Delta Option

1. In the main navigation tree, select **Device Manager > Devices**.
2. From the Device Manager launchpad, select **Update Device** to open the Update Device(s) dialog box, listing all connected and managed devices.
3. Select two devices you want to update.
4. Deselect **Run Summarize Delta Config** (if selected), and then click **Apply Changes**.
NSM displays the updated device job results for both devices.

Example: User Selects Two Devices to Update with the Delta Option, But Has no Admin Privileges

1. In the main navigation tree, select **Device Manager > Devices**.
2. From the Device Manager launchpad, select **Update Device** to open the Update Device(s) dialog box, listing all connected and managed devices.
3. Select two devices you want to update.
4. Select **Run Summarize Delta Config** (if deselected), and then click **Apply Changes**.
NSM displays the delta configuration results for both devices.

Adding Vsys Devices

A Virtual System (vsys) is a virtual device that exists within a physical security device. The vsys device functions as a completely separate security device. The physical device, called the *root device*, can contain multiple vsys devices. The following Juniper Networks security devices can be root devices:

- NetScreen-500
- ISG1000
- ISG2000
- NetScreen-5200
- NetScreen-5400


Placing the Root Device in a Global Domain or a Subdomain

Before you begin importing or modeling a root device, determine where you want to place the vsys devices:

- To add vsys devices in the global domain and one or more subdomains, add the root device to the global domain.
- To add vsys devices in a single subdomain, add the root device to that subdomain.

An example is shown in [Figure 35 on page 152](#).


Figure 35: Connecting Vsys Devices Across Domains



New - Cluster Member

New Device
Specify Name

Cluster Member Name

Color  red

Device Exists - Import Completes Workflow

☒ Device Is Reachable (i.e. Static IP Address)

☐ Device Is Not Reachable

Device Does Not Exist - Update Completes Workflow

☐ Model Device

Importing Vsys Devices

Importing vsys devices is a two-stage process:

- Import the root device—To import the root device, use the Add Device wizard to add the root device to the appropriate domain. For details, see [“Importing Devices” on page 114](#).
- Import the vsys devices—To import a vsys device, use the Add vsys wizard to add the vsys device. If you are adding multiple vsys devices to the same domain, you can add them all at once.

To import a vsys device:

1. From the domain menu, select the domain that contains the root device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select vsys Device. The Add Device wizard appears.
 - Select the root device for the vsys.
 - Select a color to represent the vsys in the UI.

- Select **Import Existing Virtual System From Physical Device**.
4. Click **Next**. Select the domain in which to import the device.
 5. Click **Next**. Select the vsys devices to import:
 - Use **SELECT ALL** to import all vsys devices from the root device.
 - Use **SELECT NONE** to clear all checked vsys devices.
 6. Click **Finish** to complete the Add Device wizard. NSM automatically imports the selected vsys configurations, and the new vsys devices appear in the Device Manager list.
 7. To check the device configuration status, mouse over the vsys in Device Manager, or check the configuration status in Device Monitor:
 - The status message “Waiting for 1st connect” might appear briefly.
 - After the vsys connects, the status displays “Import Needed”, indicating that the vsys has connected but the management system has not imported the vsys configuration yet.

To view the imported configuration, double-click the vsys in Device Manager.

To check the vsys configuration status, mouse over the vsys device in Device Manager, or check the configuration status in Device Monitor. The device status displays as “Managed, In Sync”, indicating that the vsys has connected and the management system has successfully imported the vsys configuration.

Modeling Vsys Devices

Modeling vsys devices is a two-stage process:

- Import or model the root device.

Use the Add Device wizard to add the root device to the appropriate domain. For details, see [“Importing Devices” on page 114](#) or [“Modeling Devices” on page 133](#).

- Model the vsys device.

Use the Add vsys wizard to add the vsys device. You can model a vsys on an imported or modeled root device; however, you cannot update the vsys device configuration until you have first activated the root device. You must model one vsys device at a time.

To model a vsys device:

1. From the domain menu, select the domain that contains the root device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **vsys Device**. The Add Device wizard appears.
 - Select the root device for the vsys.
 - Select a color to represent the vsys in the UI.

- Select **Model Virtual System / Virtual System Cluster Device**.
4. Click **Next** to specify the Virtual System information:
 - In the NSM vsys Name field, enter a name for the vsys device. This name identifies the vsys device in the NSM UI. The name can contain letters, numbers, spaces, dashes, and underscores.
 - In the ScreenOS vsys Name field, enter a name for the vsys device. This name is stored in the root device. The name can contain letters and numbers and can be no longer than 20 characters.
 - In the Domain field, select the domain in which to model the device.

The wizard automatically completes the device type, and OS version of the root device.
 5. Click **Next** to select the Virtual Router for this device:
 - To use the default virtual router in the root device, select **Default Vrouter**.
 - To use a shared virtual router, select **Shared Vrouter** and select one of the virtual routers defined on the root device to be shared with vsys devices.
 - To use a user-defined virtual routers, select **User Vrouter** and enter the name of a user-defined virtual router in the root device.
 6. Click **Next**, and then click **Finish** to complete the Add vsys wizard. The new vsys device appears in the Device Manager list.
 7. Ensure that the vsys is connected by viewing the device status. Check the configuration status by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor). Ensure that the configuration status for the vsys displays “Update Needed”, which indicates that the device has connected but the management system has not yet updated the device configuration.
 8. Update the device configuration by right-clicking the vsys and selecting **Update Device**. The Job Information box displays the job type and status for the update. When the job status displays successful completion, click **Close**.

After the update finishes, the device status displays as “Managed”, indicating that the device has connected and the management system has successfully updated the device configuration.

After you have modeled the vsys device, create the vsys configuration and update the device. To check the vsys configuration status, mouse over the vsys device in Device Manager, or check the configuration status in Device Monitor. The device status displays as “Managed”, indicating that the vsys has connected and the management system has successfully updated the vsys configuration.

Adding L2V Root Systems

The NetScreen-5000 series security devices running ScreenOS 5.0 L2V also support vsys transparent mode, also known as layer 2 vsys, or L2V vsys. The VLAN Trunk vsys mode and the L2V mode are mutually exclusive; you must enable one or the other on the root system:

- When modeling an L2V root, ensure that the ScreenOS version is set to 5.0L2V and the operating mode is set to Transparent. By default, the root system is modeled as a neutral vsys, enabling you to configure the system in either L2V or VLAN Trunk mode.
- When importing an L2V root:
 - If the device is in transparent mode with L2V enabled, NSM imports those settings and creates the device in L2V mode.
 - If the device is in transparent mode with L2V disabled, NSM creates the device in neutral vsys mode. You can use the NSM UI to configure the device in VLAN or L2V mode.
 - If the device is in transparent mode with VLAN trunk enabled, NSM imports those settings and creates the device in VLAN mode. In this mode, you can add vsys devices to the root system, but you cannot import VLAN IDs to those vsys devices.



NOTE: As of Release 2007.3, NSM supports L2V on ISG1000 devices running ScreenOS 6.0 and later. L2V is still supported on ISG2000 and later.

For details on configuring these vsys modes, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Adding an Extranet Device

An extranet device is a firewall or VPN device that is not a Juniper Networks security device. If you use devices from multiple manufacturers, you can add extranet devices to NSM to represent your heterogeneous network environment. After you have added the extranet device to the NSM UI, you can use the device in groups, security policies, and VPNs.

To add a new extranet device in Device Manager, click the Add icon and select **Extranet Device**. The Extranet Device dialog box appears. Enter the extranet device information:

- Name—Enter the name of the extranet device. The name can contain letters, numbers, spaces, dashes, and underscores.
- Color—Select the color that represents the extranet device in the NSM UI.
- IP Address—Enter the IP Address of the extranet device

Click **OK** to add the extranet device to NSM.

Adding Clusters

A cluster consists of multiple devices joined together in a high availability configuration to ensure continued network uptime. The device configurations are synchronized, meaning all cluster members share the same configuration settings, enabling a device to handle traffic for another if one device fails.

Adding a cluster is a two-stage process:

1. Add the cluster device object.
2. Add the members of the cluster to the cluster device object.



NOTE: When importing cluster members, ensure that their device configurations are synchronized.

Adding a Cluster Device Object

Adding a cluster device object uses a similar procedure for every supported device family:

1. In the Device Manager, click **Devices** and click the Add icon.
2. In the drop-down menu, select **Cluster**, and enter the cluster information:
 - a. Cluster Name—Enter a name for the cluster.
 - b. Color—Select a color to represent the cluster.
 - c. OS Name—Choose the OS name that identifies the family of devices.
 - d. Platform—Select the device platform for all cluster members.
 - e. (Some ScreenOS devices only) Mode—Select the Port mode. See [“Determining Port Mode \(ScreenOS Devices Only\)” on page 106](#).
 - f. Managed OS version—Select the OS version that is to run on each member of the managed cluster.
 - g. (ScreenOS only) Transparent Mode—Enable transparent mode, if desired.
 - h. (Some ScreenOS only) License Model—Specify baseline or advanced.
3. Click **OK** to create the cluster object.

The cluster device object appears in the device tree.

Adding Members to the Cluster

After creating the cluster object, add the members of the cluster. In Device Manager, select **Devices**, right-click the Cluster device, and then select **New > Cluster Member**. The Add Cluster Member wizard appears. Follow the instructions in the wizard to import or add a new cluster member.

Adding ScreenOS or IDP Clusters

To add a ScreenOS or IDP cluster, first add the cluster object as described in [“Adding Clusters” on page 155](#). Next, add each cluster member either by importing or by modeling:

- When importing cluster members, first ensure that their configurations are synchronized. Next, right-click the cluster icon in the Device Manager and select **New > Cluster Member** from the list and select the appropriate options to import the device configurations from each physical cluster device member.
- When modeling a cluster member, ensure that both cluster members have been added to the cluster device object before configuring the cluster.

By default, the cluster propagates settings made in one device member to the other device member. However, the following settings are not propagated and must be configured on each device in the cluster: VSD group, VSD priority, authentication and encryption passwords, managed IP addresses, and IP tracking settings. All other commands are propagated among devices within the cluster.

For details on creating and configuring a ScreenOS cluster, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

To create a cluster that includes an existing device (with an existing configuration and security policy) and a new device (with no configuration or security policy), you should:

1. Create the cluster.
2. Add the existing device by importing. The Add Device Wizard automatically imports the device configuration.
3. Add the new device by modeling.
4. When the device is ready, activate the device.

Adding Secure Access or Infranet Controller Clusters

To add a Secure Access or Infranet Controller cluster in NSM, you add the cluster and then add each member. Adding a member is similar to adding a standalone device.

Secure Access clusters and Infranet Controller clusters can be configured by the device administrator to operate in active/passive mode or in active/active mode. Clusters in active/passive mode are made up of a primary member and a secondary member. All traffic flows through the primary member. If the primary member fails, then the secondary member takes over.

In active/active mode, traffic is load-balanced across all cluster members. If one member fails, then load balancing takes place among the surviving members.

The number of members permitted in a cluster is different for Secure Access and Infranet Controller clusters, and also depends on whether the cluster is configured in active/active mode or in active/passive mode. You can have no more than two cluster members in active/passive mode. In active/active mode you can have up to eight members in a Secure Access cluster, or up to four members in an Infranet Controller cluster.

Before you can activate a cluster member in NSM, the device administrator must have already created the cluster and added, configured, and enabled the physical cluster member. See the *Secure Access Administration Guide* or the *Unified Access Control Administration Guide* for details on creating and configuring these clusters.

Secure Access or Infranet Controller devices configured in a cluster must have a cluster object and member objects defined in NSM before Secure Access or Infranet Controller cluster nodes can be recognized by NSM. Nodes from this cluster that subsequently contact NSM will be represented by fully functional member icons in the Cluster Manager. Cluster members whose NSM agents do not contact NSM will be displayed in NSM device monitor as unconnected devices.

Secure Access or Infranet Controller devices use member IDs to identify each cluster member object. When importing cluster members, the member ID is imported as part of the cluster, so the Add Cluster Member wizard does not prompt for the member ID.

To add a Secure Access or Infranet Controller cluster to NSM, first add the cluster object, and then add its members. You add cluster members one at a time, in a similar manner to adding standalone devices. You can add and import devices with dynamic IP addresses. NSM does not support importing Secure Access or Infranet Controller cluster members with static IP addresses.



NOTE: Adding a cluster and adding a cluster member have no effect on the cluster itself. The cluster and cluster members must already exist.

Once a Secure Access or Infranet Controller cluster is managed by NSM, subsequent changes applied to the cluster by NSM will be synchronized by the cluster across all cluster members. Similarly, changes to Secure Access or Infranet Controller cluster membership that occur via administrator action on the native device UI will be reflected back to NSM, and NSM will display the modified cluster.

For an examples of adding clusters in NSM, see [“Example: Adding and Importing a Cluster” on page 163](#).

Adding and Importing a Secure Access or Infranet Controller Cluster through Unreachable Workflow

If the cluster is already installed and configured on the network, then you can add and import that cluster into NSM.

1. On each cluster member device, configure NSM administrator logon credentials.
2. In NSM, add the cluster object using the Add Cluster wizard.

In the Device Manager, select **Devices**, click the Add icon and select **Cluster** from the list. Provide the cluster name, color of the icon, OS name, platform, and managed OS version. The OS name, platform, and OS version must match those on the physical devices.

3. In NSM, add each cluster member.

Right-click the cluster icon in the Device Manager, select **New > Cluster Member**, and follow the instructions in the Add Cluster Member wizard. When prompted, select **Device Is Not Reachable** to add an existing device with a dynamic IP address.

The last step in adding the cluster member prompts you to continue adding cluster members. Select this option if you have more members to add; unselect it if you are done adding members.

4. On each cluster member device, configure and activate the NSM agent and establish an SSH session with NSM.

5. Import the cluster.

In the Device Manager, open the cluster icon, right-click on one cluster member and select **Import Device** from the list. You do this only once and for the entire cluster because the configuration is identical for all cluster members.

After importing, the configuration appears at the cluster level in NSM. To edit the configuration, open the cluster icon, not the individual cluster members.

Because Secure Access and Infranet Controller configurations can be large, some large binary data files are not imported with the rest of the configuration. A stub is placed in the device configuration tree instead. If you need to manage these files in NSM, you must import them later as shared objects, and then create links to those shared objects from the device configuration tree. See [“Managing Large Binary Data Files \(Secure Access and Infranet Controller Devices Only\)”](#) on page 289 for details.

Adding and Importing a Secure Access or Infranet Controller Cluster through Reachable Workflow

To add a cluster member:

1. From the left pane of the NSM UI, click **Configure**.
2. Expand Device Manager and select **Devices**. The Devices workspace appears on the right side of the screen.
3. Click the **Device Tree** tab, and select the cluster to which you want to add the members.
4. Click the **New** button and select **Cluster Member**. The New–Cluster Member dialog box appears.
5. Enter the cluster member name, select **Device Is Reachable** and click **Next**.
6. Specify the device connection settings:
 - **IP Address**—IP address of the device.

- **Admin User Name**—Administrator user name created for the device.
- **Password**—Administrator password created for the device.



NOTE: The ssh port number for cluster member is 22 by default and the port number cannot be modified.

7. Click **Next**, The device is detected and the device details are displayed.
8. Enter a new name for the device in **Device Name** to change the host name of the device. An error message is displayed if the device name is not unique.
9. Click **Finish** to add the cluster member to the NSM GUI. The cluster member is added as a child of the device cluster in the Devices workspace.

Adding Clusters of Routers Running Junos OS

You can use NSM to manage clusters made up of J Series or SRX Series devices. You cannot create clusters of EX Series, M Series, or MX Series devices.

J Series and SRX Series clusters have two members; a primary member and a secondary member. You import a configuration for the entire cluster from the primary member. You need to import the configuration only once because both members share the same configuration file. Similarly, to update the configuration on the cluster, you need to push the configuration to only the primary member.

To configure local data for one cluster only, you must use a configuration group dedicated to that purpose. See [“Configuring Devices” on page 199](#) for details.

To add a cluster of J Series or SRX Series devices to NSM, you use the Add Cluster wizard as for any other cluster type, and provide a cluster name, a color for the icon in NSM, supply Junos as the name of the operating system, J Series as the Junos OS Type, a platform name, and managed OS version.

You add cluster members one at a time, in a similar manner to adding standalone devices. You can add and import devices with dynamic or static IP addresses, or you can model devices that have not yet been configured.

When adding a modeled device, you are asked to provide a member ID. One member of the cluster should be member 0 and the other member 1. These member IDs allow configuration groups to distinguish between the two cluster members, thereby providing local configuration data in a configuration group dedicated to member 0, and local configuration data in another configuration group dedicated to member 1.

When importing cluster members, the member ID is imported as part of the cluster, so the Add Member workflow does not prompt for this information.

The Add Cluster Member wizard workflow differs depending on whether you are importing a cluster member or modeling a cluster member. Outline procedures are given here. For examples of adding clusters in NSM, see “[Example: Adding and Importing a Cluster](#)” on page 163 and “[Example: Creating, Activating, and Updating a Cluster with Modeled Cluster Members](#)” on page 166.

- [Adding and Importing a Junos Cluster on page 161](#)
- [Adding a Junos Cluster with Modeled Cluster Members on page 162](#)
- [Activating and Updating a Modeled Junos Cluster on page 162](#)

Adding and Importing a Junos Cluster

If the cluster is already installed and configured on the network, then you can add and import that cluster into NSM.

1. On each cluster member device, configure NSM administrator logon credentials.
2. In NSM, add the cluster object.

In the Device Manager, select **Devices**, right-click in the right window, and select **New > Cluster** from the list. The New Cluster dialog box appears.
3. Enter the cluster name and color of the icon. Select **Junos** as the OS name and **J/SRX Series** as the Junos OS Type. Specify the platform and managed OS version and click **OK**. The Junos OS type, platform, and OS version must match those on the physical devices.
4. In NSM, add each cluster member.

Right-click the cluster icon in the Device Manager and select **New > Cluster Member**. The New—Cluster Member dialog box appears.
5. Enter the cluster member name, select **Device Is Reachable (i.e. Static IP Address)**, and click **Next**.
6. Specify the device connection settings:
 - **IP Address**—IP address of the fxp0 interface.
 - **Admin User Name**—Administrator username created for the device.
 - **Password**—Administrator password created for the device.
7. Click **Next**.



NOTE: NSM automatically detects an SSH key of the Junos OS device. Ensure that NSM can reach the device, SSH and NETCONF SSH are enabled on the device, and no firewall disrupts the communication.

8. Click **Next** to accept the device SSH key.
9. Enter the device name. After NSM autodetects the device, click **Next**.
10. Click **Finish** to import the device.

Adding a Junos Cluster with Modeled Cluster Members

When adding a modeled cluster member, you need only provide a member name. The modeled member must be activated later when the device is ready, just as for the standalone device. At that point, you provide the remaining information necessary for managing the device through NSM, such as the first connection one-time password, the NSM administrator username and password, and the Device Server IP address.

To add a cluster with modeled cluster members:

1. In NSM, add the cluster object using the Add Cluster wizard.

In the Device Manager, select **Devices**, click the Add icon, and select **Cluster** from the list. Provide the cluster name and color of the icon. Select **Junos** as the OS name and **J/SRX Series** as the Junos OS type. Provide the platform name and managed OS version.

2. Add each cluster member.

Right-click the cluster icon in the Device Manager and select **New > Cluster Member**, and follow the instructions in the Add Cluster Member wizard. When prompted, select **Model Device**.

Adding a modeled cluster member is similar to adding a modeled standalone Junos device, except that you must specify a member ID, and you also have the option of adding a second modeled cluster member within the same workflow. You can add the second cluster members later if you prefer.

After creating the cluster and member objects, you can then model the configuration. Configuration modeling is done at the cluster level, because the configuration must be identical in both cluster members. Use the configuration group mechanism to configure any member-specific data. See [“Configuring Devices” on page 199](#) for details about configuring clusters and configuration groups.

Activating and Updating a Modeled Junos Cluster

After modeling a cluster configuration, you must perform the following steps to activate and update the cluster before you can manage it from NSM:

1. Install the cluster.
2. Configure logon credentials for the NSM administrator on each cluster member.
3. In NSM, activate each cluster member.

In the Device Manager, expand the cluster icon and right-click the cluster member. Select **Activate Device** from the list and follow the instructions in the wizard. Select **Device deployed, but IP is not reachable**, when prompted.

4. On each cluster member, configure and activate the NSM agent and establish an SSH session with NSM.
5. Push the modeled configuration to the device by right-clicking any cluster member icon and selecting **Update Device** from the list.

You need push the configuration to only the primary cluster member, because software on the cluster ensures that both cluster members are synchronized.

Example: Adding and Importing a Cluster

This example adds to NSM a Secure Access cluster that already exists on the network and imports the configuration into NSM. The cluster in this example has two members: SA-1 and SA-2. Adding and importing a cluster consists of three major steps:

- [Adding the Cluster on page 163](#)
- [Adding the Cluster Members on page 164](#)
- [Importing the Cluster configuration on page 165](#)

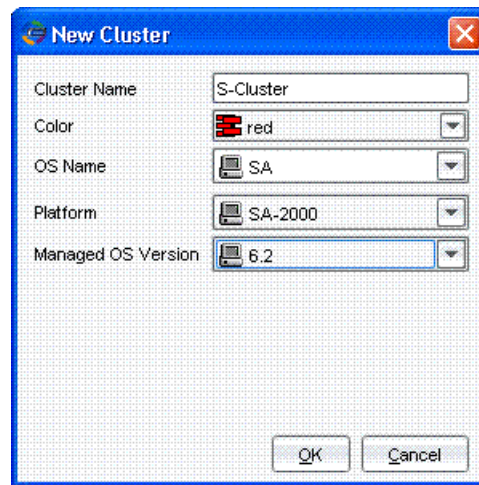
Adding the Cluster

Add a new cluster to NSM as follows:

1. Select **Device Manager > Devices**, and then click the Add icon and select **Cluster** from the list.

The Add Cluster wizard starts.

2. Enter the cluster-level information into the New Cluster dialog box as shown in [Figure 36 on page 164](#).

Figure 36: Adding a Secure Access Cluster

3. Click **OK**.

The new cluster appears in the Device Manager.

Adding the Cluster Members

1. On the device itself, configure the cluster member device with logon credentials for the NSM administrator.
2. Add the cluster member in NSM:
 - a. In the Device Manager, right-click on the **SA-Cluster** icon and select **New > Cluster Member** from the list.
 - b. In the New Cluster Member dialog box, enter a name and color for the cluster member and select **Device Is Not Reachable**.
 - c. Click **Next**. The Specify OS Name, Version, and Platform screen appears.
 - d. Specify an IP address for the NSM Device Manager server, or accept the default, and then click **Next**.
 - e. Make a note of the Unique External ID automatically displayed by NSM. The device administrator will need it later to connect the device to NSM.
 - f. Enter the NSM username and password configured on the device.
 - g. Enter a first-connection one-time password, and make a note of it. The device administrator will need it to connect the device to NSM.

- The device software initiates the TCP connection to NSM and identifies itself using the specified device ID and HMAC. The two sides then engage in SSH transport layer interactions to set up an encrypted tunnel, and NSM authenticates itself to the device based on user name and password.

NSM starts a job to import the configuration. A job window reports the progress of the job. When the job finishes, the configuration status for each cluster member changes from “Import Needed” to “Managed”.

Example: Creating, Activating, and Updating a Cluster with Modeled Cluster Members

This example creates and activates a J Series cluster named J-Cluster, with modeled members J-1 and J-2. The procedure involves four major steps:

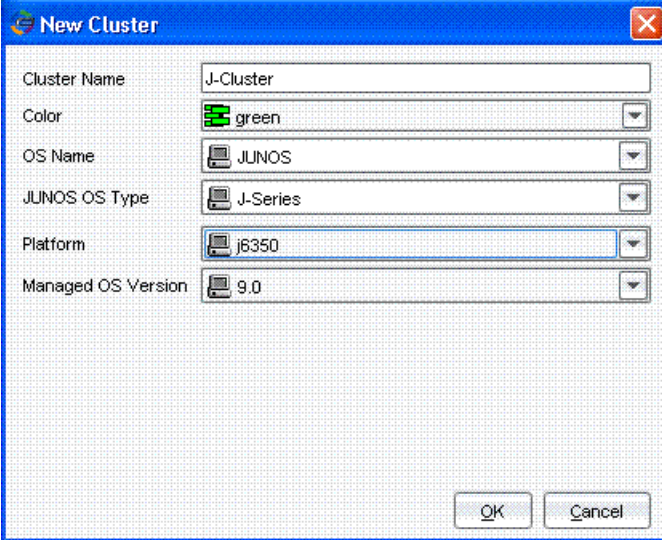
Adding the Cluster

1. Select **Device Manager > Devices**, and then click the Add icon and select **Cluster** from the list.

The add cluster wizard starts.

2. Enter the cluster-level information into the New Cluster dialog box as shown in [Figure 37 on page 166](#).

Figure 37: Adding a J Series Cluster



3. Click **OK**.

The new cluster appears in the Device Manager.

Modeling the Cluster Members

1. Right-click on the cluster icon and select **New > Cluster Member** from the list.
2. In the New Cluster Member dialog box, enter a name and color for the cluster member, and select the **Model Device** radio button.
3. Check the **Keep Adding Other Cluster Members** box and leave the Member ID as 0.

Figure 38: Adding the First Member to a J Series Cluster

The screenshot shows the 'New - Cluster Member' dialog box. The title bar is blue with the text 'New - Cluster Member' and a close button. The main area has a light blue background with a grid pattern. At the top, there's a section titled 'New Device' with a sub-label 'Specify Name'. Below this, there are two input fields: 'Cluster Member Name' with the value 'J-1' and 'Color' with a dropdown menu showing 'green'. There are two radio button options: 'Device Exists - Import Completes Workflow' (unselected) and 'Device Does Not Exist - Update Completes Workflow' (selected). Under the selected option, there's a sub-option 'Model Device' which is also selected. Below these, there's a checkbox 'Keep Adding Other Cluster Members?' which is checked. At the bottom, there's a 'Member ID' field with a spinner showing the value '0'.

4. Click **Next** to finish adding the first member.

A plus sign appears next to the cluster icon in the Device Manager indicating that the cluster now has members. The New Cluster Member dialog box re-appears for the second cluster member, as shown in Figure 35 on page 171.

5. Enter a name and color for the second member and select **Model Device**,
6. Leave the Keep Adding Other Cluster Members box unchecked.
7. Set the Member ID to 1.


Figure 39: Adding the Second Member to a J Series Cluster

The screenshot shows the 'New - Cluster Member' dialog box for the second member. The layout is identical to Figure 38, but with the following changes: the 'Cluster Member Name' field now contains 'J-2', the 'Member ID' spinner now shows '1', and the 'Keep Adding Other Cluster Members?' checkbox is now unchecked.

8. Click **Finish**.

If you expand the cluster icon in the Device Manager, you will see the new cluster members, as shown in Figure 36 on page 172.

Figure 40: Cluster Member Icons



Activating the Cluster Members

When the cluster has been properly installed, activate the cluster as follows:

1. On each cluster member device, configure logon credentials for the NSM administrator,
2. In NSM, activate each cluster member as follows:
 - a. Expand **J-Cluster** in the Device Manager to show the icons for each of the cluster members.
 - b. Right-click the cluster member icon (**J-1**) in the Device Manager and select **Activate Device** from the list.
 - c. Click the **Device deployed, but IP is not reachable** radio button.
 - d. Click **Next** to display the Specify connections setting dialog box.
 - e. Make a note the Unique External ID. The device administrator will need this ID to connect with NSM from the member device.
 - f. Enter the user name and password already set up on the device for the NSM administrator.
 - g. Enter a first connection one-time password. The device administrator will need it to connect with NSM from the member device.
 - h. Click **Finish**.
3. Repeat Step 2 for the second cluster member, J-2.

4. On cluster member J-1, configure and activate the connectivity with NSM.
 - a. Log on to the J Series router.
 - b. At the command-line prompt, identify the management system by device name, device ID, and HMAC:

For devices running the or 9.0 version of the operating system, use the following command syntax:

```
set system services outbound-ssh application-id <name> secret <string>
services netconf device-id <external-id from nsm> <NSM device server ip>
port 7804
```

For example:

```
% set system services outbound-ssh application-id nsm-wei secret 123456789
services netconf device-id abcdef 10.150.42.16 port 7804
```

For devices running the 9.1 and later versions of the operating system, use the following command syntax:

```
set system services outbound-ssh client <name> secret <secret string>
services netconf device-id <external-id from nsm> <nsm device server ip>
port 7804
```

For example:

```
set system services outbound-ssh client nsm-wei secret 123456789
services netconf device-id abcdef 10.150.42.16 port 7804
```

- c. Establish the SSH connection with the network management system.

```
% commit
```

The device software initiates the TCP connection to NSM and identifies itself using the specified device ID and HMAC. The two sides then engage in SSH transport layer interactions to set up an encrypted tunnel, and NSM authenticates itself to the device based on user name and password.

- d. In the Device List, verify the connection status of the cluster member. The status changes from "Never connected" to "Up."
 - If the configuration status is "platform mismatch," you selected the wrong device platform when adding the device into NSM. Delete the device from NSM and add it again using the correct device platform.

- If the configuration status shows “device firmware mismatch,” you selected the wrong managed OS version when adding the device into NSM. Delete the device from NSM and add it again using the correct managed OS version.
 - e. Check the device configuration status by holding your mouse cursor over the device in Device Manager, or by checking the configuration status in Device Monitor. When the device connects, the status displays “Update Needed”, indicating that the device has connected but the management system has not yet pushed the device configuration.
5. Repeat Step 4 for the second cluster member J-2.

Updating the Cluster

After you have modeled the cluster configuration, you can push the new configuration to the physical cluster using the Update Device directive.

1. In the NSM navigation tree, select **Device Manager > Devices**.
 2. Right-click **J-Cluster** (the cluster icon) and select **Update Device** from the list.
- NSM starts a job that pushes the modeled configuration to the device. A job window reports the progress. On completion, the configuration status changes from “Update Needed” to “Managed”.

Adding a Vsys Cluster and Vsys Cluster Members

A vsys cluster is a vsys device that has a cluster as its root device. Adding a vsys cluster is a three-stage process:

1. Add a vsys device that uses the cluster device as root. For details on adding a vsys device, see [“Adding Vsys Devices” on page 151](#).
2. Add cluster members to the cluster device, using the instructions in the wizard to import or add a new cluster member. A vsys cluster can have only two members.
3. Add a cluster device object. For details on adding a cluster, see [“Adding Clusters” on page 155](#). (You add members later.)

The UI also creates a vsys cluster member for each vsys device that uses the cluster as its root device. The vsys cluster member contains local information; the cluster member contains the global information. Although a cluster can have only two members, a root vsys device can support more than two vsys devices.

Example: Adding a Vsys Cluster

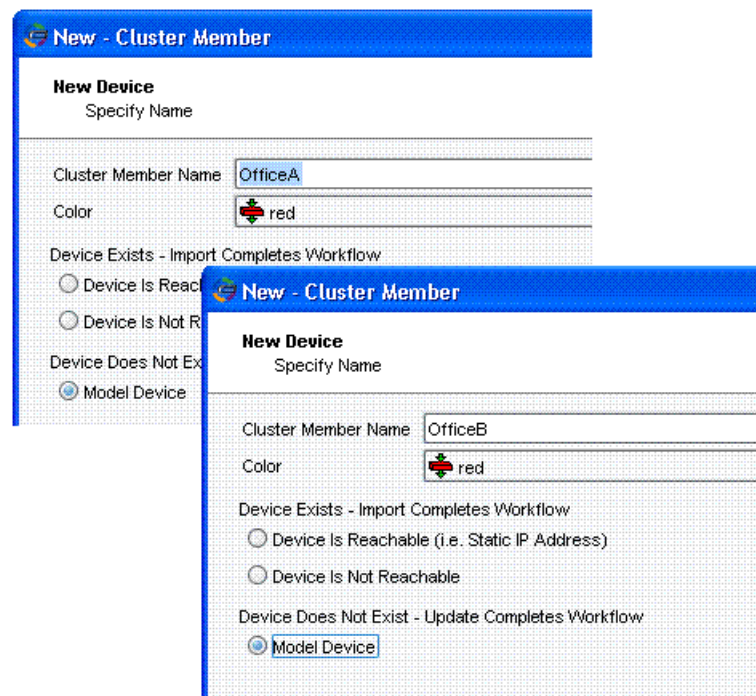
In this example, you add a vsys cluster with two members and two vsys.

1. Add the cluster device:
 - a. In the main navigation tree, select **Device Manager > Devices**.
 - b. Click the Add icon and select **Cluster**. The new cluster dialog box appears.

- c. Configure the following information:
 - For Name, enter **Paris Cluster**.
 - For OS Name, select **ScreenOS/IDP**.
 - For Platform, select **ns5400**.
 - For OS Version, select **5.1**.
 - d. Click **OK** to save the new cluster object.
2. Add cluster members:
 - a. In the main display area, right-click **Paris Cluster** and select **New > Cluster Member**. The New Cluster Member dialog box appears.
 - b. Configure the cluster members OfficeA and OfficeB as shown in [Figure 41 on page 171](#).

As you add each cluster member, NSM automatically creates both the cluster member and the vsys cluster member.

Figure 41: Configuring Cluster Members for Paris Vsys Cluster

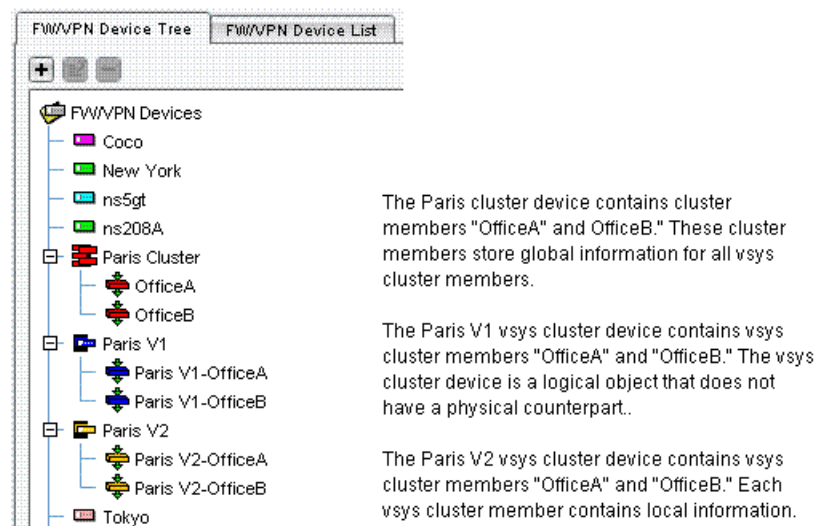


3. Add the first vsys device:
 - a. Click the Add icon and select **Vsys Device**. The new vsys device dialog box appears.
 - b. Configure the root as the Paris Cluster device, select a color, and choose **Model Virtual System/Virtual System Cluster Device**. Click **Next** to continue.

- c. Configure the NSM and ScreenOS name as Paris V1 and select global as the domain. Click **Next** to continue.
 - d. Configure the vrouter for the vsys by selecting **Default Vrouter**, and then click **Next** to continue.
 - e. Click **Finish** to add the new vsys cluster device.
4. Add the second vsys cluster device:
- a. Click the Add icon and select **Vsys Device**. The new vsys device dialog box appears.
 - b. Configure the root as the Paris Cluster device, select a color, then select **Model Virtual System/Virtual System Cluster Device**. Click **Next** to continue.
 - c. Configure the NSM and ScreenOS name as Paris V2 and select global as the domain. Click **Next** to continue.
 - d. Configure the vrouter for the vsys as the Default Vrouter, and then click **Next** to continue.
 - e. Click **Finish** to add the new vsys cluster device.

In the security device tree, the Paris Cluster (and cluster members) and Paris vsys cluster (and cluster member) appear as shown in [Figure 42 on page 172](#).

Figure 42: Paris Cluster Members and Paris Vsys Cluster Members



Importing an SRX Series Cluster into NSM

With ScreenOS devices, an HA cluster consists of two independent devices that share most parts of a configuration between them in addition to the session state. NSM manages both cluster members individually and sends cluster member-specific configuration parameters to its appropriate cluster member.

In contrast, Junos OS clusters consist of two or more Routing Engines and multiple fabrics, which are interconnected to each other and to one big virtual fabric. Individual cluster members are called nodes.

Junos OS clustering is also called virtual chassis clustering. In a virtual chassis, only one Routing Engine is active while all others replicate the exact state of the master. All Routing Engines have the same configuration installed, and replication of the configuration is automatic. Cluster node-specific parameters are configured within the groups configuration of the master configuration. All fabrics are active in all chassis at the same time, and traffic can enter or exit any chassis. Chassis are interconnected through fabric links.

From a device management perspective, you do not see individual cluster members, instead you see one virtual device that is defined by the master Routing Engine in the cluster. When you configure such a cluster from the CLI, you would always connect to the master Routing Engine.

The following options are available in NSM to manage SRX Series clusters:

- Cluster mode — A device in cluster mode is imported as a device cluster. With cluster mode, NSM talks to the fxp0 interface of the Routing Engine. Use cluster mode when you can reach NSM from both nodes through fxp0.
- Virtual chassis mode — A device in virtual chassis mode is imported as a standalone device. With virtual chassis mode, NSM talks to a payload interface. Use virtual chassis mode when you can reach NSM only from the master Routing Engine through a payload interface. Virtual chassis mode is the preferred method when NSM has to cross the firewall to gain access to the fxp0 interface.



NOTE: When you reconfigure an SRX Series cluster from virtual chassis mode to cluster mode you must reboot both chassis members before you can reimport the device in cluster mode to NSM. Otherwise, NSM assumes the device is in virtual chassis mode. Even a `commit full` will fail.

Activating Management Access

Before you import a device into NSM, you need to activate management access to NSM. Unlike ScreenOS devices, which are managed through SSH, SRX Series devices are managed through the Junos OS DMI interface, both inbound, initial contact to the device, and outbound, after a device has made initial contact to NSM.

Enabling Inbound Access

NSM automatically installs the outbound DMI interface but not the inbound interface. Unlike ScreenOS devices, where you enable only SSH, for SRX Series devices you need to enable NETCONF inbound.

To enable the NETCONF inbound, run the following commands:

```
set system services ssh;
```

```
set system services netconf ssh;
```

Importing an SRX Series Device in Cluster Mode

To manage each Routing Engine individually you need to configure IP access for each node. Both fxp0 interfaces need to be reachable by NSM. To use a hovering cluster address, the fxp0 addresses of both Routing Engines must be in the same network.

The following sample configuration imports a device in cluster mode:

```
set groups node0 interfaces fxp0 unit 0 family inet address 192.168.1.101/24;
```

```
set groups node0 interfaces fxp0 unit 0 family inet address 192.168.1.100/24  
master-only;
```

```
set groups node1 interfaces fxp0 unit 0 family inet address 192.168.1.102/24;
```

```
set groups node1 interfaces fxp0 unit 0 family inet address 192.168.1.100/24  
master-only;
```

In this example, address 192.168.1.101 is for RE0, 192.168.1.102 is for RE1, 192.168.1.100 is the cluster address, and 192.168.1.100 is the hovering cluster address that hovers within the cluster. Talking to the hovering cluster address always gets you the master on the line. Unlike ScreenOS, there is no cluster MAC address. The master-only keyword under the groups configuration, as opposed to the fxp0 interface under the general system configuration, answers ARP requests. It also updates ARP tables through gratuitous ARPs. NSM speaks individually to both nodes, 192.168.1.101 and 192.168.1.102. CLI management is done through cluster address 192.168.1.100. However, configuring the cluster address with the master-only keyword causes an interruption in the communication with NSM with every failover of RE0, because each node stops the NSM connection and restarts the communication with a new source IP. Therefore, you may consider managing the cluster through its individual fxp0 addresses even for CLI management.

Importing an SRX Series Device in Virtual Chassis Mode

The standard way to manage SRX Series devices is with out-of-band management access through the fxp0 interfaces. However, SRX Series devices can also be managed through a payload interface. The advantage of virtual chassis mode is that it is entirely transparent to the operator, or to NSM, whose node hosts the master Routing Engine while Management traffic can enter any node's fabric and is directed to the master Routing Engine through the fabric interconnect transparently and automatically. Not all models support virtual chassis mode.

To enable virtual chassis mode on the device, run the following commands in edit mode:

```
root@srx# set chassis cluster network-management cluster-master;
```

```
root@srx# set interfaces ge-0/0/0 unit 0 family inet address 192.168.1.100/24;
```




NOTE: In virtual chassis mode `fxp0` should not be configured. If both devices can reach NSM through `fxp0`, then the device connection status in NSM flaps every five seconds because both devices are trying to establish a connection and each new connection terminates the previous one.

Removing Remnants from Previous Imports

If the device was previously imported, you might be unable to import it again after you deleted the device object from NSM. Even if you were successful, the remnants in the configuration might have unexpected results. NSM does not detect whether a device was previously imported or make a new import possible. It simply adds its own configuration to the one already in place.

In virtual chassis mode, NSM adds its configuration (`outbound-ssh`) to the general configuration `set system services outbound-ssh`. In cluster mode, NSM adds its configuration `outbound-ssh` inside the specific node. If you still have `outbound-ssh` configured in the general part, neither SRX Series cluster member can make contact with NSM in cluster mode.

Before you proceed with an import, you need to remove these configurations:

```
delete groups node 0 system services outbound-ssh
```

```
delete groups node 1 system services outbound-ssh
```

```
delete system services outbound-ssh
```

Importing the Cluster Member

For cluster mode, you must first import the secondary node and then import the primary node. If you do it the other way around, the import will fail. Node 0 needs to be the master and node 1 the secondary.

Run the following commands to check the node status:

```
request chassis cluster failover redundancy-group 0 node 0
```

```
show chassis cluster status redundancy-group 0 node 0
```

In cluster mode, you must import the cluster once again, after you have added the second member. Otherwise NSM removes the NSM configuration from the second cluster member, as it was not present during the first import.

No specific action is required for virtual chassis mode.

Adding a Blade Server

A blade server consists of a server chassis, which houses SA/IC/WXC blades. The blades within the blade server are known as the blade server members, while the CMC/MSC, which controls the blade server members, is known as the management blade. NSM provides an interface that lets you logically group and manage the blade server members under a blade server through a management blade.



NOTE: SA is known as Junos Pulse Secure Access Service and IC is known as Junos Pulse Access Control Service.

The management blade can run on the following MAG-series Junos Pulse gateway platforms:

- MAG6610
- MAG6611

The blade server members can run on the following platforms:

- MAG-SM160
- MAG-SM361
- MAG-4611
- WXC-590
- WXC-2600
- WXC-3400

The following standalone devices belong to Junos Pulse Gateway platforms:

- MAG-2600
- MAG-4610



NOTE: These standalone devices can be added and managed by NSM similar to any other standalone devices.

To add a blade server:

1. In **Device Manager**, select **Devices**.
2. Click the Add icon and select **Blade Server**. The Add Device wizard appears.
3. Select **Device is Reachable** (default).
4. Click **Next**. The Specify Connection Settings dialog box appears.
5. Enter the following connection information:

- a. The IP address of the blade server.
- b. The username of the device administrator.
- c. The password for the device administrator.



NOTE: All NSM passwords are case-sensitive.

- d. Select the connection method (**Telnet**, **SSH version 1**, or **SSH version 2**) and the port number for the selected service.
 - If you selected **Telnet**, click **Next** and skip to Step 6.
 - If you selected an **SSH version**, click **Next**. The Verify Device Authenticity dialog box appears. The device wizard displays the RSA Key FingerPrint information. To prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.
6. When the auto-detected device information appears, verify that the device type, OS version, and device serial number are correct. NSM automatically detects the hostname configured on the device and uses it as the device name. You can also change the auto-detected hostname.
7. Modify the auto-detected device name within the device from its config editor page in NSM. If you previously modified the device hostname using the Junos OS CLI, SNMP, or J-Web interface, you can use the edit option to modify the device name again in NSM after importing the device.



NOTE: If you select the **Device is not reachable** workflow, NSM cannot detect the hostname automatically. You must specify a device name.

8. Click **Next** to add the device to NSM.
9. After the device is added, click **Next** to import the device configuration.
10. Click **Finish** to exit the Add Device wizard.
11. Double-click the blade server in **Device Manager** to view the imported configuration.

To check the device configuration status, mouse over the device in **Device Manager** (you can also check configuration status in Device Monitor). The device status displays as **Managed**, indicating that the device has connected and the management system has successfully imported the device configuration.

After adding a blade server, you may want to add or group the blade server members under this server. You can select manual or automatic grouping. For more information on adding the devices manually, see the section [“Manually Adding SA/IC Blade Members to the Blade Server” on page 178](#) and [“Manually Adding WXC Blade Members to the Blade Server” on page 179](#). For more information on automatic grouping, see the section [“Automatic Grouping of Device Members Under the Blade Server” on page 178](#). The

maximum number of blade server members supported is 2 on a MAG6610 chassis. On a MAG6611 chassis, the maximum number is 4.

When a blade server is deleted, the blade server members that are part of the blade server are not deleted and continue to be managed by NSM as standalone devices.

Automatic Grouping of Device Members Under the Blade Server

You may have added the blade server members that are physically part of a blade server to NSM as standalone devices instead of blade members. To automatically group these devices under the respective blade server, you must perform an import on the blade server:

- Right-click the device and click **Import**. The blade members are automatically grouped under the blade server.

Manually Adding SA/IC Blade Members to the Blade Server

To add a blade server member:

1. From the left pane of the NSM UI, click **Configure**.
2. Expand **Device Manager** and select **Devices**. The Devices workspace appears on the right side of the screen.
3. Click the **Device Tree** tab, and select the blade server to which you want to add the members.
4. Click the Add icon and select **Blade Server Member**. The New–Device dialog box appears.
5. Select **Device is Reachable** (default).
6. Click **Next**.
7. Enter the device connection settings:
 - The IP address of the device.
 - The administrator username created for the device.
 - The administrator password created for the device.
8. Click **Next**. The device is detected and the device details are displayed.
9. Enter a new name for the device in the **Device Name** text box to change the hostname of the device. An error message is displayed if the device name is not unique.



NOTE: If you select the **Device is not reachable** workflow, the device is not detected automatically. You need to specify the device name and select an OS from the **OS Name** drop-down list. Also, make sure the platform and the OS version information is accurate. Click **Next** and follow the instructions displayed in the dialog box.

10. Click **Finish** to add the blade server member to the NSM GUI. The blade server member is added as a member of the blade server in the Devices workspace.

Manually Adding WXC Blade Members to the Blade Server

To manually add a blade server member:

1. From the left pane of the NSM UI, click **Configure**.
2. Expand **Device Manager** and select **Devices**. The Devices workspace appears on the right side of the screen.
3. Click the **Device Tree** tab, and select the blade server to which you want to add the member.
4. Click the Add icon and select **Blade Server Member**. The New-Device dialog box appears.
5. Select **WXC/AAM** device under **Device Exist - Web UI Workflow**.
6. Click **Next**.
7. Enter the device connection settings:
 - Device Name—Name of the device.
 - IP Address—IP address of the device.
 - Admin User Name—Administrator username created for the device.
 - Password—Password created for the device.
 - Root User Password—Root user password created for the IDP device.
 - Connect to Devices with—Option to connect to the device.
 - Port Number—Port number of the device.
8. Click **Finish**. The device is detected and the device details are displayed.
9. Enter a new name for the device in the **Device Name** text box to change the hostname of the device. An error message is displayed if the device name is not unique.



NOTE: If you select the **Device is not reachable** workflow, the device is not detected automatically. You need to specify the device name and select an OS from the **OS Name** drop-down list. Also, make sure the platform and the OS version information is accurate. Click **Next** and follow the instructions displayed in the dialog box.

10. Click **Finish** to add the blade server member to the NSM GUI. The blade server member is added as a member of the blade server in the Devices workspace.

Adding Multiple Devices Using Automatic Discovery (Junos, SA, and IC Devices)

You can use automatic discovery to add and import multiple Junos, SA, and IC devices into NSM. You do so by configuring and running discovery rules. For a Junos, SA, or IC device to be discovered by this mechanism, it must be configured with a static IP address.

By configuring and running a discovery rule, you can search a network to discover devices in a specified subnet or within a range of IP addresses. Authentication of the devices is through administrator login SSH v2 credentials and SNMP community settings, which you also configure as part of the rule. Devices that match the rules for discovery also present an SSH key for your verification before the device is added to NSM.

Adding a Device Discovery Rule

To add a device discovery rule:

1. In the Configure pane of the NSM navigation tree, click **Device Discovery Rules**.
Any existing discovery rules appear in the main display area.
2. Click the Add icon to display the New Device Discovery Rule dialog box.
3. Give the rule a name and provide the following search criteria for the devices:
 - A prefix for device names—for example, “USA”. The prefix is used to assign names to the devices when they are added into NSM. For example, when a device at IP address **10.204.32.155** is added to NSM, its name will be **USA_10.204.32.155**.
Check the **Use Host Name if Available** checkbox, if you want the device hostname to be used as the prefix.
 - An IP subnet or range of IP addresses.
 - Administrator login name and password.
 - SNMP version and community string.
 - Select the **Run Topology Discovery** checkbox, if you want to use topology discovery to discover devices.
4. Click **Apply** to add the rules to the list of device discovery rules.

To ensure that devices are discovered correctly, it is recommended that you:

- Use SNMP V1 or V2C versions.
- Use a subnet mask narrower than 255.255.240.0 because the broadest subnet mask recommended for any device discovery rule is 255.255.240.0.
- Do not use more than 4096 IP addresses.



NOTE: Device discovery supports only IPv4 addresses. IPv6 based devices are not discovered.



NOTE: Device discovery will not add cluster members. Cluster members will have to be added to NSM manually.

Running a Device Discovery Rule

To run a device discovery rule:

1. In the Configure pane of the NSM navigation tree, click **Device Discovery Rules**.
2. Select the rule you want to run.
3. Click the Run icon in the discovery rules toolbar.

The device discovery Progress dialog box appears.

NSM pings each IP address in the specified range to see which ones it can reach, and then runs the rule against each reachable device. When it finds a device that satisfies the rule, it displays a dialog box containing an SSH key for you to verify.

4. Accept the SSH key. The device automatically connects with NSM and NSM imports the device configuration.

The Device Discovery Progress box shows which devices were added, devices that were found but could not be added, and at what stage in the discovery process the device add process failed.

Adding Many Devices Using CSV Files

If your network includes a large number of devices, you can save time by adding multiple devices in a single workflow using the Add Many Device wizard.

With the wizard, you can add up to 4000 devices at a time to a single domain (you cannot add multiple devices to different domains at one time). Additionally, for some types of ScreenOS devices, you can create configlets to activate rapidly your newly deployed security devices. However, you cannot configure Rapid Deployment for Secure Access devices, Infranet Controller devices, or devices running Junos OS. Neither can you configure

RD when adding multiple ScreenOS security devices that are systems (NetScreen-500, NetScreen-5000, ISG1000, ISG2000).



NOTE: You cannot use the Add Many Devices wizard to add clusters or cluster members, EX Series Ethernet Switches configured as a virtual chassis or an SRX virtual chassis.

Adding many devices is a three-step process:

1. Create the CSV file.

This file defines all the required and optional values for each device.

2. Use the Add Many Devices wizard to select the CSV file to import or model the devices.

The wizard validates the CSV file, notifies you of any errors, and then adds the devices for which all defined values are valid.

- When importing devices with static IP addresses, the device configuration is automatically imported during the Add Many Devices workflow.
- When importing devices with dynamic IP addresses, you must manually import the device configuration after the Add Many Devices workflow is complete.
- When modeling ScreenOS devices for Rapid Deployment, you can also create configlets during the Add Many Devices workflow, or select to skip configlet creation.

The time it takes for NSM to activate and import devices depends on the number of devices and the management system configuration.

3. Verify the device configuration.
4. Select the **Use Host Name if Available** checkbox, to detect the host name configured on the device and use it as the device name.
5. Select the **Run Topology Discovery** checkbox to trigger a topology discovery.

The following sections provide details about each step.

Creating the CSV File

Within a **.csv** file, you define the device configuration values for each device you want to add. The required and optional values depend not only on how the device is deployed on your network—static IP addresses, dynamic IP addresses, or undeployed devices—but also on the device family.

You must create a separate CSV file for the following devices:

- Devices with static IP addresses—In this CSV file, you define the device parameters required to add and import the device configurations from all supported device types except IDP.
- Devices with dynamic IP addresses—In this CSV file, you define the device parameters required to add all supported devices (except IDP) to the NSM system.

- Undeployed ScreenOS devices—In this CSV file, you define the device parameters required to add and model ScreenOS 5.x and later devices in the NSM system.



NOTE: You can model many ScreenOS devices, but you cannot activate many devices except when using the Rapid Deployment process.

Juniper Networks provides CSV templates in Microsoft Excel format for each type of CSV file. These templates are located in the `utils` subdirectory where you have stored the program files for the UI client, for example:

C:\Program Files\Network and Security Manager\utils

For each CSV file, each row defines a single device's values for each parameter. For text files, columns are separated by commas.

Devices with Static IP Addresses

For devices with static IP addresses, create a `.csv` file with the parameters shown in Table 23 on page 183.

Table 23: CSV File Information for Devices with Static IP Addresses

Field Name	Type	Required	Acceptable Values
Name	String	yes	
Color	String	yes	black, gray, blue, red, green, yellow, cyan, magenta, orange, pink
Device IP Address	String	yes	192.168.1.1, 10.1.1.10, 3.3.3.3
Device Admin Name	String	yes	<administrator>
Device Admin Password	String	yes	<password> Note: All passwords handled by NSM are case-sensitive.
Connection Protocol	String	yes	telnet, ssh_v1, ssh_v2.
Device Admin Port	Integer	no	23, 22, 4444, 7777 If null, uses 23 for Telnet and 22 for SSH
SSH Fingerprint	String	yes (when connection SSH)	<SSH fingerprint> Use any to bypass check.

Example: Using an Excel File to Add Multiple Static IP Devices

To edit the template for adding many devices with static IPs:

1. Copy and open the Microsoft Excel file **bulkadd_ipreachable-sample.csv** or **bulkadd_ipreachable-DMIDMI-sample.csv** from the **C:/Program Files/Network and Security Manager/Utils** directory.
2. Using one row for each device you want to add, enter the required values for the device. You can also provide optional values, if desired.
3. Save the file to a location on your local drive.

Example: Using a Text File to Add Multiple Static IP Devices

To add four security devices that use static IP addresses, create a text file with the following text:

```
Chicago,green,10.100.31.78,netscreen,netscreen,ssh_v2,,any
Memphis,orange,10.100.20.236,netscreen,netscreen,ssh_v2,,any
Columbus,red,10.100.20.200,netscreen,netscreen,ssh_v2,,any
Cincinnati,blue,10.100.20.2367,netscreen,netscreen,ssh_v2,,any
```

Save the file as a .csv file.

Device with Dynamic IP Addresses

For devices with dynamic IP addresses, create a .csv file with the parameters shown in [Table 24 on page 184](#).

Table 24: CSV File Information for Devices with Dynamic IP Addresses

Field Name	Type	Required	Acceptable Values
Name	String	yes	dev1, Chicago, NS-208
Color	String	yes	black, gray, blue, red, green, yellow, cyan, magenta, orange, pink
OS Name	String	Yes	ScreenOS, SA, IC, junos-es (for J Series or SRX Series devices), junos for (M Series or MX Series devices), junos-ex (for EX Series devices)

Table 24: CSV File Information for Devices with Dynamic IP Addresses (continued)

Field Name	Type	Required	Acceptable Values
Platform	String	yes	<p>With OS name ScreenOS:</p> <p>ns5GT-Combined, ns5GT-Dual-Untrust, ns5GT-Trust-Untrust, ns5GT-Dual-DMZ, ns5GT-Extended, ns5GT-Dmz-Dual-Untrust, ns5GT-Home-Work, ns5GTadsl-Home-Work, ns5GTadsl-Trust-Untrust, ns5GTadsl-Extended, ns5XP, ns5GTadslwlan-Extended, ns5GTadslwlan-Home-Work, ns5GTadslwlan-Trust-Untrust, ns5Gtwlan-Extended, ns5Gtwlan-Dmz-Dual-Untrust, ns5Gtwlan-Combined, ns5Gtwlan-Home-Work, ns5Gtwlan-Dual-Untrust, ns5Gtwlan-Trust-Untrust, ns5Gtwlan-Dual-Dmz, ns5XT-Combined, ns5XT-Dual-Untrust, ns5XT-Trust-Untrust, ns5XT-Home-Work, ns-25, ns-50, ns204, ns208, ns500, ns5200, ns5400, nsHSC-Home-Work, nsHSC-Trust-Untrust, nslSG1000, nslSG2000, SSG5-ISDN, SSG5-SB, SSG5-ISDN-WLAN, SSG5-Serial, SSG5-Serial-WLAN, SSG5-v92, SSG5-v92-WLAN, SSG-20, SSG-20-WLAN, SSG-140, SSG-320, SSG-320M, SSG-350, SSG-350M, SSG-520, SSG-520M, SSG-550, SSG-550M</p> <p>With OS name junos:</p> <p>m7i, m10i, m120, m320, m40e, m7i, m320, mx240, mx480, mx960</p>

Table 24: CSV File Information for Devices with Dynamic IP Addresses (continued)

Field Name	Type	Required	Acceptable Values
Platform (continued)	String	yes	<p>With OS name junos-es:</p> <p>j2320, j2350, j4350, j6350</p> <p>srx100, srx240, srx650, srx3400, srx3600, srx5600, srx5800</p> <p>With OS name junos-ex:</p> <p>ex3200-24p, ex3200-24t, ex3200-48p, ex3200-48t, ex4200-24f, ex4200-24p, ex4200-24t, ex4200-48p, ex4200-48t, ex8208, ex8216</p> <p>With OS name SA:</p> <p>SA-2000, SA-2500, SA-4000, SA-4000(FIPS), SA-4500, SA-6000, SA-6000(FIPS), SA-6500, SA-700</p> <p>With OS name IC:</p> <p>IC-4000, IC-4500, IC-6000, IC-6500</p>
Device subtype	String	yes	Set to "none".
Managed OS Version	String	yes	<p>With OS name ScreenOS (see Table 7 on page 14 for a list of OS versions that apply to each ScreenOS platform):</p> <p>5.0, 5.0FIPS, 5.0DSLW, 5.0WLAN, 5.0NSGP, 5.0GPRS, 5.0L2V, 5.0dial, 5.0IDPI, 5.1, 5.1GPRS, 5.1shotglass, 5.1SSG, 5.2, 5.3, 5.3TMAV, 5.4, 5.4FIPS, 6.0, 6.1, 6.2, 6.3.</p> <p>With OS name junos-es:</p> <p>9.0, 9.1, 9.2, 9.3, 9.4, 9.5, 9.6.</p> <p>With OS name Junos:</p> <p>9.0, 9.1, 9.2, 9.3, 9.4, 9.5, 9.6.</p> <p>With OS name SA:</p> <p>6.3, 6.4</p> <p>With OS name IC:</p> <p>2.2, 3.0</p>
Transparent Mode	String	yes	on, off
License Key Model	String	yes	Range of values defined in dcf file

Table 24: CSV File Information for Devices with Dynamic IP Addresses (continued)

Field Name	Type	Required	Acceptable Values
First Conn OTP	String	yes (when using ScreenOS)	
Device Admin Name	String	yes	
Device Admin Password	String	yes	Must be a minimum of 9 characters

Example: Using an Excel File to Add Multiple Dynamic IP Devices

To use an Excel file for adding many device with dynamic IP addresses:

1. Copy and open the **bulkadd_nonreachable-sample.csv** file or the **bulkadd_nonreachable-DMI-sample.csv** file located in the **C:/Program Files/Network and Security Manager/Utils** directory.
2. Using one row for each device you want to add, enter the required values for the device. You can also provide optional values, if desired.
3. Save the file to a location on your local drive.

Example: Using a Text File to Add Multiple Dynamic IP Devices

To add four devices that use dynamic IP addresses, create a text file with the following text:

```
switch1,red,junos-es,j2320,none,9.0,off,none,netscreen,root,netscreen
switch2,red,junos-es,j2320,none,9.0,off,none,netscreen,root,netscreen
switch3,red,junos-es,j2320,none,9.0,off,none,netscreen,root,netscreen
switch4,red,junos-es,j2320,none,9.0,off,none,netscreen,root,netscreen
```

Save the file as a **.csv** file.

Undeployed Devices

For undeployed devices (ScreenOS 5.x and later releases only), create a **.csv** file with the parameters shown in [Table 25 on page 187](#).

Table 25: CSV File Information for Undeployed Devices

Field Name	Type	Required	Acceptable Values
Name	String	yes	Valid character
Color	String	yes	black, gray, blue, red, green, yellow, cyan, magenta, orange, pink
OS name	String	yes	ScreenOS

Table 25: CSV File Information for Undeployed Devices (continued)

Field Name	Type	Required	Acceptable Values
Platform	String	yes	Must be a device platform that supports ScreenOS 5.x or later release and configlets (cannot be a ns5GTADSL device)
device subtype	String	yes	Set to "none".
ScreenOS Version	String	yes	5.x, 6.x
Transparent Mode	String	yes	on, off
License Key Model	String	yes	Range of values defined in dcf file
First Conn OTP	String	yes	Must be a minimum of 9 characters
Connection Type	String	yes	static, pppoe, dhcp, prompt
Device IP Address	String	yes (when connection type is static)	
Device Netmask	String	yes (when connection type is static)	8, 24, 28, 32 Any valid netmask in CIDR format
Device Gateway	String	yes (when connection type is static)	
PPPoE User Name	String	yes (when connection type is PPPoE)	
PPPoE User Password	String	yes (when connection type is PPPoE)	Must be a minimum of 9 characters
Configlet Password	String	no	Default to a random string between 9 and 256 characters
Device Admin Name	String	yes	
Device Admin Password	String	yes	Must be a minimum of 9 characters
Telnet Port	Integer	no	Default to 23
SSH Port	Integer	no	Default to 22
Restrict to Serial Number	String	yes	on, off

Table 25: CSV File Information for Undeployed Devices (continued)

Field Name	Type	Required	Acceptable Values
Device Serial Number	String	yes if restrict to serial is on	Valid device serial number

Example: Using an Excel File to Add Multiple Modeled Devices

To edit an Excel file for adding many modeled devices:

1. Copy the **bulkadd_model-sample.csv** file located in the **C:/Program Files/Network and Security Manager/Utils** directory. The header row at the top defines the settings.
2. Using one row for each device you want to add, enter the required values for the device. You can also provide optional values, if desired.
3. Save the file to a location on your local drive.

Example: Using a Text File to Add Multiple Modeled Devices

To add and model three security devices, create a text file with the following text:

```
dev13,orange,ScreenOS,ns5XP,none,5.0,off,advanced,netScreen123,static,10.10.30.5,32,10.10.30.1,,123456abc,netScreen,netScreen,,on
dev14,green,ScreenOS,ns50,none,5.0,off,advanced,netScreen123,pppoe,,,root,netScreen,,1netScreen,netScreen1,,off
dev15,red,ScreenOS,ns204,none,5.0,off,advanced,netScreen123,dhcp,,,,,2netScreen,netScreen2,,off
```

Save the file as a **.csv** file.

Validating the CSV File

When you add the device, NSM validates the configuration information in the **.csv** file and creates a Validation Report. The report lists any incorrect or duplicate configurations, and indicates the exact line that contains invalid data.



NOTE: The Validation Report displays only the first error in the line. If the line contains additional errors, those errors do not appear in the Validation Report.

Select **Cancel** to quit the Add Many Devices process, or select **Add Valid Devices** to begin adding the devices for which you have provided valid device configurations. If the Validation Report listed incorrect configurations, you can still select Add Valid Devices; however, only the devices with correct configurations are added. If the **.csv** file contains duplicate configurations, NSM ignores the duplicates.

After you have added devices, you cannot roll back or undo your changes. To edit or delete a device, select the device in the UI and make the necessary changes.

Importing Many Devices

The import process differs between devices that use static IP addresses and devices that use dynamic IP addresses:

- For devices with static IP addresses, the Add Many Devices wizard automatically imports the device configurations.
- For devices with dynamic IP addresses, you must manually import the device configurations.

In some cases, you might also need to configure NACN or other features on the physical device to enable the device to connect to NSM.

After you have added the devices, verify that the device configuration import matches your expectations. For details, see [“Verifying Imported Device Configurations” on page 131](#).

Adding and Importing Many Devices with Static IP Addresses

For devices with static IP addresses:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Many Devices**. The Add Device wizard appears.
4. In the Add Device wizard:
 - Select **Device Is Reachable** (default).
 - Specify the location of the CSV file.
5. Click **Next**. The Add Device wizard validates the CSV file and provides a Validation Report:
 - Select **Cancel** to quit the Add Many Devices process.
 - Select **Add Valid Devices** to begin adding the devices for which you have provided valid device configurations.

The Add Device wizard adds the valid devices and automatically imports their configurations.

Adding and Importing Many Devices with Dynamic IP Addresses

For devices with dynamic IP addresses:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Many Devices**. The Add Device wizard appears.
4. In the Add Device wizard:
 - Select **Device Is Not Reachable**.

- Specify the location of the CSV file.
- Specify the output directory for the `.cli` file. For each valid device configuration that uses a dynamic IP address, NSM creates a `.cli` output file. By default, the `.cli` file is saved to the following GUI Server directory:

```
/usr/netscreen/GuiSvr/var/ManyDevicesOutput/<inputFile_YYYYMMDDHHMM>./
```

Before the device can be managed by NSM, you must enter the CLI commands in the `.cli` file on the physical security device.

5. Click **Next**. The Add Device wizard validates the CSV file and provides a Validation Report:
 - Select **Cancel** to quit the Add Many Devices process.
 - Select **Add Valid Devices** to begin adding the devices for which you have provided valid device configurations.
6. The Add Device wizard adds the valid devices and automatically imports their configurations.

Modeling Many Devices

For undeployed devices, you can create device configurations in NSM in a single workflow. After you have created modeled configurations for each device, you must activate each device individually.



NOTE: The devices must be running ScreenOS 5.x or a later release.

To model many devices:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon, and then select **Many Devices**. The Add Device wizard appears.
4. In the Add Device wizard:
 - Select **Model Device**.
 - Specify the location of the CSV file.
5. Click **Next**. The Add Device wizard validates the CSV file and provides a Validation Report:
 - Select **Cancel** to quit the Add Many Devices process.
 - Select **Add Valid Devices** to begin adding the devices for which you have provided valid device configurations.

The Add Device wizard adds the valid devices to the NSM UI.
6. Model the device configuration as desired.

After you have added the device and created modeled device configurations for your undeployed device, you are ready to activate the device and prompt it to connect to the management system. After that device has made contact with NSM, you can install the modeled configuration you created on the physical device. For details on activating a device, see [“Activating a Device” on page 135](#).

Using Rapid Deployment

You can model devices, generate configlets, and activate many ScreenOS devices at one time. Alternatively, you can model multiple devices initially, and then generate configlets and activate them later. The devices must be running ScreenOS 5.x or later and support configlets; NetScreen systems (NetScreen-500, 5000 line, ISG1000, and ISG2000) do not support configlets.

Modeling and Activating Many Devices with Configlets

To model, create configlets, and activate at the same time:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Many Devices**. The Add Device wizard appears.
 - Select **Model Device**.
 - Specify the location of the CSV file.
 - Select **Activate and Create Configlet now** (ns208 and below).
 - Specify the output directory for the **.cfg** file. For each modeled ScreenOS device configuration, NSM creates a **.cfg** output file. By default, the **.cfg** file is saved to the following GUI Server directory:

`/usr/netscreen/GuiSvr/var/ManyDevicesOutput/<inputFile_YYYYMMDDHHMM> /`

4. Click **Next**. The Add Device wizard validates the CSV file and provides a Validation Report:
 - Select **Cancel** to quit the Add Many Devices process.
 - Select **Add Valid Devices** to begin adding the devices for which you have provided valid device configurations.

The Add Device wizard adds the valid devices to the NSM UI.

5. Send the **.cfg** file to the onsite administrator for the corresponding device. After the onsite administrator installs the configlet on the physical security device, the device automatically contacts the NSM Device Server, which establishes an always-on management connection. For instructions for the onsite administrator, see [“Installing the Configlet” on page 147](#), or refer to the *Rapid Deployment Getting Started Guide*.
6. Model the device configurations as desired.
7. Install the modeled configuration. After the onsite administrator has installed the configlets and the devices have successfully connected to NSM, you can install the modeled device configurations on the physical devices:

- a. Ensure that the device is connected by viewing the device status. Hold your mouse cursor over the device in Device Manager, or check the configuration status in Device Monitor. Ensure that the configuration status for the device displays “Update Needed”, which indicates that the device has connected but the management system has not yet updated the device configuration.
- b. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update; when the job status displays successful completion, click **Close**.

After the update finishes, the device status displays as “Managed”, indicating that the device has connected and the management system has successfully updated the device configuration.

For more details on the Rapid Deployment, see [“Using Rapid Deployment \(ScreenOS Only\)” on page 142](#).

Activating Many Devices with Configlets

Before activating devices and creating a configlet, you must configure a modeled configuration for the device in the NSM UI.

To create configlets and activate many devices:

1. In Device Manager, select **Devices**.
2. Click the Add icon and select **Activate Many Devices**. The Activate Device wizard appears.
3. Select the devices to activate.
4. Specify the output directory for the **.cfg** file. For each modeled ScreenOS device configuration, NSM creates a **.cfg** output file. By default, the **.cfg** file is saved to the following GUI Server directory:

```
/usr/netscreen/GuiSvr/var/ManyDevicesOutput/<inputFile_YYYYMMDDHHMM>/
```



NOTE: For security, you cannot edit a configlet file directly. To make changes to the information in any configlet file, run the Activate Many Device wizard to regenerate the configlet.

5. Send the **.cfg** file to the onsite administrator for the corresponding device. After the onsite administrator installs the configlet on the physical security device, the device automatically contacts the NSM Device Server, which establishes an always-on management connection. For instructions for the onsite administrator, see [“Installing the Configlet” on page 147](#), or refer to the *Rapid Deployment Getting Started Guide*.
6. Click **OK**. A Job Manager window displays the progress of the activation. When finished, click **Close**.
7. Update the physical device with the modeled configuration.

Adding Device Groups

You can create groups of devices to manage multiple devices at one time. Use device groups to organize your managed devices, making it easier for you to configure and manage devices within a domain. You can group devices by type (such as all the NetScreen-5GTs in a domain), by physical location (such as all the security devices in the San Jose office), or logically (such as all the security devices in sales offices throughout western Europe).

Use the groups to:

- Deploy new or updated device configurations to the entire device group.
- Deploy new or updated policies to the entire device group.
- Create reports using the log information from the entire device group.

Device groups enable you to execute certain NSM operations on multiple security devices at the same time. For example, if you have a device group of the same type of devices running similar ScreenOS versions, you can upload the firmware on all devices in the group at the same time. You can also add devices to the NSM UI, place the devices in a device group, and then import the device configurations for all devices in the device group at one time.

The devices that you add to a device group must exist; that is, you must have previously added or modeled the devices in the domain. You can group devices before configuring them. You can add a device to more than one device group. You can also add a device group to another device group.



NOTE: You cannot apply a template to a device group. You must apply templates to individual devices in a device group. If you need to apply the same set of templates to multiple devices, you can create a single template that includes all the templates that are to be applied to a device, and then apply the combined template to each device.

Example: Creating a Device Group

In this example, you create a device group that includes security devices used to protect the Sales and Marketing department of your organization.

1. Add and model the following devices to the management system:
 - Outside Sales
 - Marcom
 - Direct Marketing

- Sales
 - Marketing
2. In the navigation tree, select **Device Manager > Devices**.
 3. Click the Add icon and select **Group** from the list. The New Group dialog box displays all existing devices for the current domain in the Non-members list.
 4. In the Name field, enter **Sales**.
 5. In the Non-members list, select the devices that you want to be part of the Sales device group.
 6. Click **Add** to move the selected devices to the Member list (or drag the selected devices into the Member list), and then click **OK**.

Setting Up NSM to Work With Infranet Controller and Infranet Enforcer

A ScreenOS firewall that is managed by NSM can also be configured as an Infranet Enforcer in a UAC solution.

The Infranet Controller specifies an authorization server \$infranet for each Infranet Enforcer in its list. This name is required for correct operation between the Infranet Controller and the Infranet Enforcer. Conversely, if NSM has multiple Infranet Enforcers in its global domain, it will distinguish among them by renaming additional Infranet Enforcers \$infranet_1, \$infranet_2, and so on. To resolve this naming conflict, you must move each Infranet Controller to a separate NSM domain.

In addition, because the Infranet Controller regularly changes its NACN password with the Infranet Enforcer, you should always import the Infranet Enforcer into NSM before performing a device update to it.

The following procedures prevent these conflicts between NSM and the Infranet Controller:

- [Avoiding Naming Conflicts of the Authorization Server Object on page 195](#)
- [Avoiding NACN Password Conflicts on page 198](#)

Avoiding Naming Conflicts of the Authorization Server Object

To avoid naming conflicts with the authorization server objects, follow these steps:

1. On the Infranet Controller, create the Infranet Enforcer instances:
 - a. On the Infranet Controller, select **UAC -> Infranet Enforcer -> Connection**.
 - b. Click **New Enforcer**.
 - c. Fill out the information requested in the display.

Enter an NACN password. Remember it because you will need to use it again while setting up the Infranet Enforcer. If you are setting up a cluster instead of a single device, enter all the serial numbers in the cluster, one per line.

- d. Click **Save Changes**.
 - e. Repeat Steps b through d until all of your Infranet Enforcers have been entered.
2. If you do not have one already, create a CA certificate for each Infranet Enforcer.
 - a. Create a certificate signing request (CSR) for an Infranet Controller server certificate, and use the CA certificate to sign the server certificate.
 - b. Import the server certificate into the Infranet Controller.
 - c. Import the CA certificate into the Infranet Enforcer.

For details about setting up the certificates, see the *Unified Access Control Administration Guide*.

3. On each Infranet Enforcer, create the Infranet Controller instance:
 - a. On the Infranet Enforcer, select **Configuration > Infranet Auth > Controllers**.
 - b. Click **New**.
 - c. Enter the parameters as prompted.

The password in the second section must be the NACN password you entered in step 1.
 - d. Click **OK**.
 - e. Repeat steps a through d for all of the infranet enforcers.
 - f. On the Infranet Controller, select **UAC -> Infranet Enforcer -> Connection** and check that all the Infranet Enforcers have been added.
4. On NSM, delete the Infranet Enforcer firewalls from the global domain:
 - a. In the global domain, select **Device Manager > Devices** to list all the devices.

- b. Right-click each Infranet Enforcer firewall device in turn and select **Delete** from the list.
5. On NSM, delete the \$infranet instances from the Object Manager:
 - a. Select **Object Manager > Authentication Servers**.
 - b. Right-click each \$infranet_n object and select **Delete** from the list.
 - c. Select **VPN Manager > VPNs**, and check that you do not have any \$infra under VPN Manager. These objects are usually deleted automatically when you remove the firewall.
6. Create a new subdomain for the Infranet Enforcers:
 - a. Select **Tools > Manage Administrators and Domains**.
 - b. Select the **Subdomains** tab.
 - c. Click the **Add** icon.
 - d. In the New Subdomain dialog box, enter an appropriate name for the subdomain so you know what it will be used for, and then click **OK**.
 - e. From the drop-down list on the top left side, select your new domain.
The new domain is empty.
 - f. Add a Single Infranet Enforcer or Infranet Enforcer Cluster.
 - g. Repeat steps e and f for every Infranet Enforcer or Infranet Enforcer Cluster you need to add to NSM. When you are finished, \$infranet appears instead of \$infranet_# in each of the domains except the global domain.
7. In NSM, add the Infranet Enforcer objects to the new domain:
 - a. Select **Device Manager > Devices**.
 - b. Click the **Add** icon, and then select **Device** to start the Add Device wizard.
 - c. In the New Device window, provide a name for the device, a color for its icon in NSM, and check **Device is Reachable**.
 - d. Follow the instructions in the wizard to add and import the device.
 - e. Repeat steps b through d for each Infranet Enforcer device.

Avoiding NACN Password Conflicts

When you need to manage the Infranet Enforcers, reimport the configuration each time. Otherwise, a NACN password mismatch is possible because the Infranet Controller Dynamically changes this password periodically. Additionally, it is also good practice to issue a Summarize Delta Config directive and ensure that no \$infra policies are present. If there are, that means that the Infranet Controller has changed something on the Infranet Enforcer since you last imported the device configuration.

If you do not reimport the configuration, be sure to update the Infranet Controller and Infranet Enforcer at the same time.

CHAPTER 5

Configuring Devices

The Device Manager module in Network and Security Manager (NSM) enables you to configure the managed Juniper Networks devices in your network. You can edit configurations after you add and import a managed device, or create configurations when you model a device.

This chapter provides details of device configuration concepts and provides some examples. For instructions for configuring specific device settings, see the *Network and Security Manager Online Help* or the appropriate device-specific administration guide. This chapter also describes two important tools that you can use to simplify configuring devices: templates and configuration groups.

After you edit or create a configuration for a device object in NSM, you must update the configuration on the managed device for your changes to take effect. For details on updating devices, see [“Updating Devices” on page 257](#).

Use security policies to configure the rules that control traffic on your network. For devices that you configure to use centrally managed policies, see [“Configuring VPNs” on page 597](#). For devices that you configure to use in-device policy management, see the device-specific documentation. For details on configuring VPNs, see [“Introduction to Network and Security Manager” on page 3](#).

This chapter contains the following sections:

- [About Device Configuration on page 200](#)
- [Editing Devices Using the Device Editor on page 202](#)
- [Using Device Templates on page 210](#)
- [Using Configuration Groups on page 235](#)
- [Using Configuration Groups with Templates on page 243](#)
- [Configuring Clusters on page 248](#)
- [Configuring Junos Devices with Redundant Routing Engines on page 250](#)
- [Overview of VRRP Support in NSM on page 252](#)
- [Managing Configuration Files on page 254](#)

About Device Configuration

The device configuration contains the configuration settings for a managed device, such as interface, routing, and authentication settings. You can edit device-object configurations after you add or import a managed device, or create configurations when you model a device. When you are satisfied with your changes, you can then update the managed device with the modeled device configuration to make your changes effective.

NSM does not support all device configuration settings. You might need to make some changes to the device directly using the device's native GUI or CLI.

Each family of devices supported by NSM has different configuration requirements.

About Configuring Device Families

Through NSM, you can configure any of the following device families:

- Devices running Junos OS, including J Series routers, SRX Series gateways, EX Series switches, M Series routers, and MX Series routers
- ScreenOS or IDP
- Secure Access
- Infranet Controller

See [“Managed Devices” on page 14](#) for an overview of each of these device families and lists of supported platforms and operating system versions.

Most devices can be configured using the following interfaces:

- Native Web UI
- Native CLI
- NSM UI

All supported devices can be configured through the native Web UI or through NSM. All except Secure Access, Infranet Controller, and IDP devices have a native CLI you can use to configure the device.

When you use the native Web UI or CLI to edit the device configuration, you do so directly. That is, changes to the configuration take place immediately.

When you use NSM to edit the device configuration, you initially make the changes to a device object that models the device in NSM. When you are satisfied with your configuration changes, you use the Update Device directive to push the configuration from the device object in NSM to the device itself. At that point, the edited configuration becomes active.

About Configuring Clusters, VPNs, Vsys Devices, Policies, and Shared Objects

In addition to configuring specific devices, NSM also enables you to configure clusters, VPNs, vsys devices, policies, and shared objects:

- Clusters are made up of two or more devices from the same platform and managed OS version. You configure these as a separate entity. Configuration applied to the cluster also applies to each cluster member. See [“Configuring Clusters” on page 248](#) for details on configuring clusters.
- VPNs provide a cost-effective and secure way for routing private data through the internet. You can configure devices for inclusion in VPNs either centrally in the VPN Manager, or in the device object itself. For details about configuring VPNs, see [“Configuring VPNs” on page 597](#).
- Vsys devices are virtual devices that exist within a physical ScreenOS security device. A vsys cluster device is a vsys device that has a cluster as its root device. See [“Adding Vsys Devices” on page 151](#).
- Policies are sets of rules that provide a comprehensive plan that determines how the device behaves on your network. For ScreenOS and IDP devices, you configure policies within the NSM Policy Manager. For Secure Access, Infranet Controller, and EX Series devices, you must configure policies in the device. For J Series routers or SRX Series gateways, you can configure policies either in the NSM Policy Manager or in the device, but not both. For details about configuring policies, see [“Configuring Security Policies” on page 473](#).

Configuration Features

You can edit the device object configuration through the device editor, or you can use templates or configuration files to simplify configuration:



NOTE: These features edit only the device object in NSM. The newly configured values will not affect the device itself until you push the modeled configuration to the device using the Update Device directive. See [“Updating Devices” on page 257](#) for details about updating devices.

- [About the Device Editor on page 201](#)
- [About Device Templates on page 202](#)
- [About Configuration Groups on page 202](#)

About the Device Editor

To edit the device using the device editor, select **Device Manager > Devices**, select the device you want to edit, and then click the Edit icon. NSM displays the imported or modeled configuration parameters of the device, along with information specific to NSM, like device startup information and the color of the device icon in NSM. The layout of the screens and the information contained within them depends on the device family and the specific device platform and operating system version.

See [“Editing Devices Using the Device Editor” on page 202](#) for details.

About Device Templates

A template is a predefined set of configuration values that helps you reuse common information. A device object can refer to multiple templates, and you can use templates to configure and deploy multiple devices quickly. A device template looks like a device configuration in the device editor—the template displays panels and tables for interfaces and zones, for example, into which you can enter values. When you add a new device that uses similar information as a previously added device, you can use a device template to fill in specific configuration values so you do not have to reenter information.

Because the configuration data and layout vary depending on device family, NSM provides empty templates for each device family: ScreenOS/IDP, Secure Access, Infranet Controller, M Series and MX Series, J Series (which is also the correct template for SRX Series devices), and EX Series devices.

You can manually override any value set by a template in the configuration for a specific device.

From a device object configuration, you can reference multiple device templates. From a device template you can, in turn, reference additional device templates.

Any change applied to the template is immediately reflected in all device objects that reference the template.

For details about device templates, see [“Using Device Templates” on page 210](#).

About Configuration Groups

Configuration groups are similar to device templates in that you define configuration data to be used multiple times. In configuration groups, the configuration data is used within the same device but at several levels in the configuration. For example, a configuration group can be used to apply the same interface configuration data to multiple interfaces. A special case use of configuration groups is to apply configuration data in different members of a cluster.

Configuration groups are used only with Junos devices.

You can manually override most values set by a configuration group in the configuration for a specific device.

You can apply multiple configuration groups in multiple places in the same device object, but you cannot apply a configuration group directly to multiple device objects. You can, however, define a configuration group in a template, and apply that template to multiple device objects.

For details see [“Using Configuration Groups” on page 235](#) and [“Using Configuration Groups with Templates” on page 243](#).

Editing Devices Using the Device Editor

To configure device information in NSM, select **Device Manager > Devices**, select the device, and then click the Open icon.



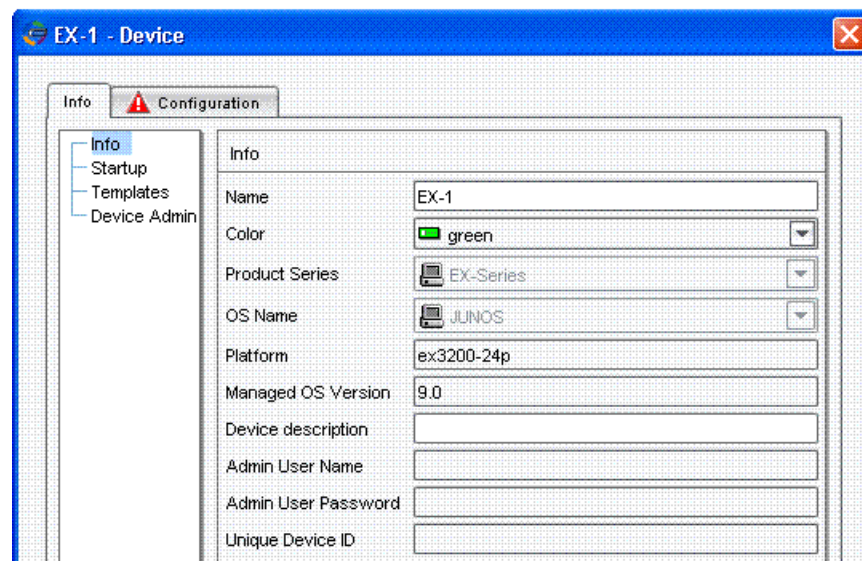
NOTE: When you open a device for viewing or editing, the NSM UI loads the entire device configuration into memory to enhance UI performance while configuring the device. When you close a device to which you made changes, the UI unloads some of the device configuration from the client memory. Although this memory optimization occurs quickly, you might see the following message appear: “Optimizing client memory usage for device”.

For Secure Access, Infranet Controller, and Junos devices, NSM displays the following tabs of information:

- Info
- Configuration

Figure 43 on page 203 shows an example.

Figure 43: Info and Configuration Tabs



The Device Info tab contains information maintained in NSM. This information can neither be imported from the device, nor is it ever pushed to the device by an Update Device directive. It contains the following information:

- Info—Basic device information such as name, OS version, and IP address.
- Startup—Startup information such as the one-time password, which is required for modeled devices and imported devices that use a dynamic IP address. The one-time password is used to authenticate the first connection between the device and NSM.
- Templates—All templates available to the device family to which the device belongs. See “Using Device Templates” on page 210 for details.
- Device Admin—Polling interval for alarm statistics.

The Quick Configuration tab

The Configuration tab contains device data that resides on the device itself when active. This information can be imported from the device, or is modeled information waiting to be pushed to the device by an Update Device directive. This information varies by device family and typically includes device-specific feature settings such as those for interfaces, routing, and authentication.

For ScreenOS or IDP devices, access to general device information and device feature configuration data is through the same tree. The top branch of that tree (the Info branch) contains similar data to that displayed through the Device Info tab for the other device families. [Figure 44 on page 204](#) shows an example.

Figure 44: ScreenOS and IDP Device Configuration Information







The screenshot shows a configuration window titled "SOS-3 - Device". On the left is a tree view with the following items: Info (selected), Network, Device Admin (with a warning icon), Auth, Report Settings, Security, VSYS, Advanced, L2TP/Auth/Local User, VPN Settings, and IDP Profiler Settings. The main pane displays the "Info" tab with the following fields:

- Name: SOS-3
- Color: green (dropdown menu)
- Platform: nslSG1000
- Mode: (empty field)
- Device description: (empty field)
- Series: NS
- OS Name: ScreenOS/IDP (dropdown menu)
- Managed OS Version: 6.2
- Running OS Version: Unknown
- Support Level: Full Support (dropdown menu)
- Serial Number: (empty field)
- IP Address: . . . (empty field)
- Transparent Mode: false
- Enable Jumbo Frame: false
- Device Root Admin: (empty field)
- Admin User Name: (empty field)
- Password: (empty field)
- Connect To Device With: SSH Version 2
- SSH Key: ...
- Policy for Device: (empty field)
- Security Policy Name: Please select... (dropdown menu)

Validation and Data Origination Icons

The device editor might display some of the icons shown in [Table 26 on page 205](#). These icons provide validation status or data origination information about the displayed data items. Data validation icons include those for errors and warnings. Data origination icons indicate whether a data item was inherited from a template or configuration group, or whether an inherited value has been overridden in the device edit dialog box. For details about using device templates, see ["Using Device Templates" on page 210](#). For information about configuration groups, see ["Using Configuration Groups" on page 235](#).

Table 26: Validation Icons

Icon	Message Type	Meaning	Priority
	Error	A configuration or parameter is not configured correctly in the NSM UI. Updating a device with this modeled configuration will cause problems on the device.	Highest
	Warning	A configuration or parameter is not configured correctly in the NSM UI. Updating a device with this modeled configuration might cause problems on the device.	
	Override	The displayed value was set manually and that the value overrides whatever value might come from a template or configuration group. The icon can also indicate an override of a VPN-provided value or a cluster-provided value. Changes to a template will not change this value unless "Remove conflicting device values" is selected in the Template Operations dialog box.	
	Template Value	The value was inherited from a template. Changes to the template are also shown in the device edit dialog box.	
	Configuration Group Values	The value was inherited from a configuration group. Changes to the configuration group are also shown in the device editor.	
	From Object	A value is set for a field in a template or configuration group definition. This icon is shown only in a template or configuration group definition. From Object messages appear only when you view template objects to help find fields set in the template.	Lowest

When more than one type of icon appears within a panel, the highest priority icon appears next to the icon in the tree and the panel title bar.

Configuring Device Features

To configure a device that has been added, imported, or modeled in NSM:

1. In the navigation tree, select **Device Manager > Devices**.
2. Open the device configuration using one of the following methods:
 - Double-click the device object in the security device tree or the device list.
 - Select the device object and then click the Edit icon.
 - Right-click the device object and select **Edit**.

For ScreenOS and IDP devices, the device navigation tree appears on the left, listing the device configuration parameters by function.

- For Secure Access, Infranet Controller, and Junos devices, select the **Configuration** tab.

The device configuration tree appears in the left pane.

- In the device navigation tree, select a function heading to see device parameters, and then select the configuration parameter you want to configure.
- Make your changes to the device configuration, then choose one of the following:
 - Click **OK** to save your changes and close the device configuration.
 - Click **Apply** to save your changes and continue making changes.
 - Click **Cancel** to discard all changes and close the device configuration.

To reset a device feature to its default value, right-click on the feature name in the device editor and select **Revert to template/default value**.

A brief overview of each device family follows. For details, see the referenced device-specific documentation.

Configuring ScreenOS/IDP Device Features

The device configuration tree for a ScreenOS or IDP device looks similar to the example in [Figure 45 on page 206](#).

Figure 45: ScreenOS Device Object Configuration Data

The screenshot shows the 'SOS-2 - Device' configuration window. On the left is a tree view with the following nodes: Info, Network, Device Admin (with a warning icon), Auth, Report Settings, Security, Deep Inspection, Attack DB, VSYS, Advanced, L2TP/XAuth/Local User, and VPN Settings. The 'Info' node is selected. The main area on the right displays configuration parameters for the selected device:

Name	SOS-2
Color	green
Platform	ns5400
Mode	
Series	NS
OS Name	ScreenOS/IDP
Managed OS Version	5.0FIPS
Running OS Version	Unknown
Support Level	Full Support
Serial Number	
IP Address	. . .
Transparent Mode	false
Device Root Admin	
Admin User Name	
Password	
Connect To Device With	SSH Version 2
SSH Key	...
Policy for Device	
Security Policy Name	Please select...

At the bottom right are buttons for OK, Cancel, and Apply.

For details about configuring the device features for all supported ScreenOS and IDP platforms, see the *Configuring Screen OS Devices Guide* (http://www.juniper.net/techpubs/en_US/nsm2011.1/information-products/pathway-pages/screenos-devices/index.html) or the *Configuring Intrusion Detection Prevention Devices Guide* (http://www.juniper.net/techpubs/en_US/nsm2011.1/information-products/pathway-pages/intrusion-detection-prevention-devices/index.html).

Unsupported Changes

Some device configurations can be performed only by the device administrator using the CLI or Web UI. An NSM administrator cannot perform the following device configurations in the Device Manager:

- Configuring functions that are only applicable for the device administrator, such as setting initial IKE contact, audible alarms, MAC addresses, or console operations.
- On standalone IDP Sensors, configuring Sensor mode (sniffer, transparent, and so on), port speed and duplex settings, virtual routers, and other settings. See the *IDP Installer's Guide*, *IDP Concepts and Examples Guide*, and *IDP ACM Help* for more information.
- Configuring functions that require device administrator intervention, such as Secure Command Shell (SCS) and Secure Shell (SSH) client operation.
- Executing debugging commands.

Changes that Affect the Management Connection

Some configuration changes to a managed device can affect the NSM connection to the device when you update the device, such as:

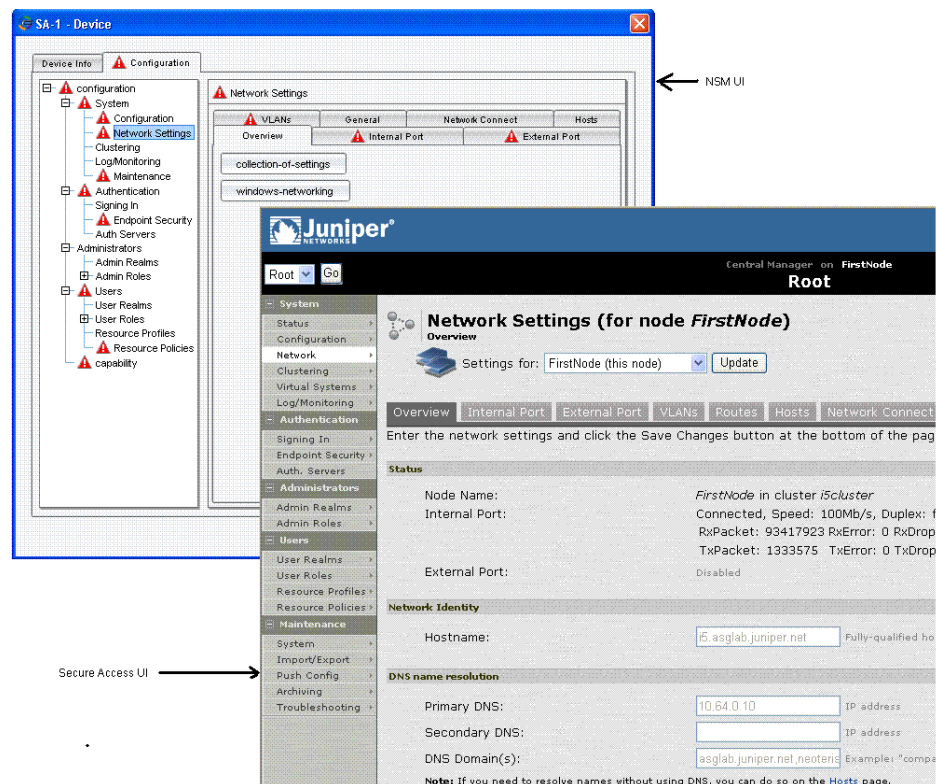
- Changing the connection method (Telnet or SSH) used between the NSM Device Server and the managed device.
- Disabling the ability of the managed device to communicate with the NSM Device Server.
- Changing the IP address of the NSM Device Server on the managed device.
- Changing the interface on the managed device that is permitted to receive NSM management traffic.
- Changing the VPN that handles traffic between the managed device and the NSM Device Server.
- Modifying router information on the managed device.
- Changing security policy rules on the managed device that cause NSM traffic to be dropped.

If you need to make any of the above changes to the managed device, use the Web UI or CLI to make the changes locally, and then reimport the device configuration into the NSM UI.

Configuring Secure Access or Infranet Controller Device Features

For Secure Access Devices and Infranet Controller Devices you can configure most of the same parameters through NSM that you can configure through the native device GUI or with the device CLI. The configuration screen rendered in NSM for any Secure Access or Infranet Controller device looks similar to that of the native GUI. [Figure 46 on page 208](#) compares the two user interfaces. In this example, the view is of the Network Settings screen.

Figure 46: Secure Access Device Object



For details about configuring Secure Access devices, see the *Configuring Secure Access Devices Guide* (http://www.juniper.net/techpubs/en_US/nsm2011.1/information-products/pathway-pages/secure-access-devices/index.html). For details about configuring Infranet Controller devices, see the *Configuring Infranet Controllers Guide* (http://www.juniper.net/techpubs/en_US/nsm2011.1/information-products/pathway-pages/infranet-controller-devices/index.html).

You can edit the same configuration parameters from NSM as you can in the native GUI, except that you cannot perform the following operations:

- View the system status as presented in the System Status screen in either a Secure Access device or an Infranet controller device. However, you can view status information on the Investigate panel of NSM.
- Edit Secure Access or Infranet Controller device licensing information, although you can view it.
- Create clusters, join nodes to clusters, or enable or disable cluster nodes.
- Manage the device configuration as a text-based file.
- Execute device-specific troubleshooting commands.
- Use the technical support service that allows packaged collections of information for remote analysis by Juniper Networks Technical Assistance Center (JTAC).
- Reboot the device.

The view of the configuration from NSM might also be missing data configured in large binary files. When you initially import a Secure Access or Infranet Controller configuration into NSM, large binary data files are replaced with stubs. If you want to manage these files on NSM, you must import them separately as shared objects, and then link to those objects from the stubs in the device configuration. See "[Managing Large Binary Data Files \(Secure Access and Infranet Controller Devices Only\)](#)" on page 289 for details.

Configuring Junos Device Features

You can configure Junos device features in NSM. Although the configuration screens rendered in NSM look different, the top-level configuration elements present are mostly the same as for the native GUI configuration screens, and correspond to commands in the CLI.

See the NSM documentation for the following specific device families for details about configuring specific device parameters:

- J Series Services Routers and SRX Series Services Gateways

Configuring J Series Services Routers and SRX Series Services Gateways Guide
(http://www.juniper.net/techpubs/en_US/nsm2011.1/information-products/pathway-pages/J-Series-SRX-Series-devices/index.html)

- EX Series switches

Configuration Guide for EX Series Devices
(<http://www.juniper.net/techpubs/software/management/security-manager/nsm2011.1/nsm-ex-series-book.pdf>)

- M Series and MX Series devices

M Series and MX Series Devices Guide
(http://www.juniper.net/techpubs/en_US/nsm2011.1/information-products/pathway-pages/m-mx-devices/m-mx-series.html)

To edit the configuration of a Junos device, you can edit the device object configuration itself, create and apply a template, or use configuration groups.

For more information about configuration groups, see [“Using Configuration Groups” on page 235](#).

To edit Junos device configuration data in NSM, double-click the device in the Device Manager and select the **Configuration** tab. The configuration tree appears in the main display area with all parameters viewable or configurable from NSM.

Updating the Configuration on the Device

The primary difference between editing the configuration in NSM and editing the configuration as described in the Secure Access, Infranet Controller, or Junos device documentation is that edits described in those documents apply changes directly to the device. Edits done in NSM apply to the device object in NSM, which is not pushed to the device until you perform an Update Device directive.

To perform an update device operation, right-click the device in the Device Manager, and then select **Update Device** from the list.

Using Device Templates

Use templates to define a common device configuration and then reuse that configuration information across multiple devices. In a template, you need define only those configuration parameters that you want to set; you do not need to specify a complete device configuration. Templates provide these benefits:

- You can configure parameter values for a device by referring to one or more templates when configuring the device.
- When you change a parameter value in a template and save the template, the value also changes for all device configurations that refer to that template, unless specifically overridden in the device object.

When you apply a template to a device, NSM applies the template settings to the device. For example, you can create a template that specifies the IP address of the NTP server to which all managed security devices synchronize their clocks. You can apply this template to the configuration of each device in your subdomain (or all devices if defined in the global domain) so that all devices use the same NTP server.

A template contains all possible fields for all possible devices within a device family. NSM provides different templates for:

- ScreenOS/IDP devices
- Secure Access devices
- Infranet controller devices
- J Series devices (includes SRX Series devices)
- M Series and MX Series devices
- EX Series devices

The templates for each family are different because the configuration fields for each family are different.

Some devices might not have all fields prescribed in the template for that device family. You can apply a template to any device in its family. NSM will ignore any fields that do not apply to the given device.

A template can refer to other templates, enabling you to combine multiple templates into a single template. When you make changes to any of the referenced templates, those changes propagate through the combined template. For instructions for creating and applying templates, see the *Network and Security Manager Online Help* topics, “Applying Templates.”

Junos device templates can also contain configuration groups. A configuration group is similar to a template in that it specifies configuration data for reuse, but only within the same device, and not across devices. See [“Using Configuration Groups” on page 235](#).

The EX Series device template has the following configuration enhancements:

- **Template categories:** These are logical groupings of configuration nodes based on product functionality, such as VLAN, STP, and PoE for EX Series devices. You can select one or more of these template categories while creating an EX Series template. The configuration tree displays only the nodes associated with the selected template category, which enhances the usability of the template. If template categories are not selected, the default display is a full tree view. You can also view the associated template categories in the **Device Template** table view.
- **Customized configuration tree views:** Instead of numerous configuration nodes being always displayed irrespective of whether the node is configured, the EX Series device template allows you to select either a full or partial configuration tree view. Check the **Display all** box, at the upper left of the **Configuration** window, to have a full tree view. The default setting is the partial configuration tree view, which shows the first level configuration node and its children. You can switch between the two views, even after the device template is migrated.



NOTE: Conflicting configuration created using different configuration nodes, for example, VLAN created through the VLAN configuration template, as well as the Interface Configuration, can result in wrong device configuration during template application and device update.

Template data applied to a device object does not affect the device itself until you push the device object configuration to the device using the Update Device directive as described in [“Updating Devices” on page 257](#).

Modifying Values in Templates

You can modify a template that has already been applied to one or more device configurations. When you change a field value in a template, the device object that references the template also changes.



NOTE: When you change a template, one or more devices that use the template might become invalid. For example, the change could cause a required field to be missing or a field value to be outside the allowed range.

Example: Creating and Applying a Device Template for DNS Settings

In this example, you create and apply a template that configures the IP addresses of primary and secondary DNS servers for ScreenOS devices.

Creating the Template

Create the template as follows:

1. In the navigation tree, select **Device Manager>Device Templates**.
2. Click **Add** in the Device Template Tree or the Device Template List and select **ScreenOS/IDP Template** from the list.

The New Device Template dialog box displays the template navigation tree in the left pane and the Info screen in the right pane.

3. In the Info screen, enter **DNS** in the Name field.
4. From the template navigation tree, select **Network>DNS>Settings**.
5. Configure the following:
 - **Primary DNS Server IP**— Enter **IP address** as **1.1.1.1**, and select **Src interface**.
 - **Secondary DNS Server IP**— Enter **IP address** as **2.2.2.2**, and select **Src interface**.
 - **Tertiary DNS Server IP**— Enter **IP address** as **3.3.3.3**, and select **Src interface**.



NOTE: If **Is IPv6** is selected for Primary, secondary and Tertiary DNS Server, then enter IP address as **1:1:1:1:1:1:1:1**, **2:2:2:2:2:2:2:2** and **3:3:3:3:3:3:3:3** respectively.

- **Static Host**— Enter **Host Name**, **Is IPv6**, select **Is IPv6** and enter **Host IP**.
 - **DNS Refresh Schedule**— Select **Refresh Daily**.
6. Click **OK** to save the template.

You can now use this template when configuring security devices.



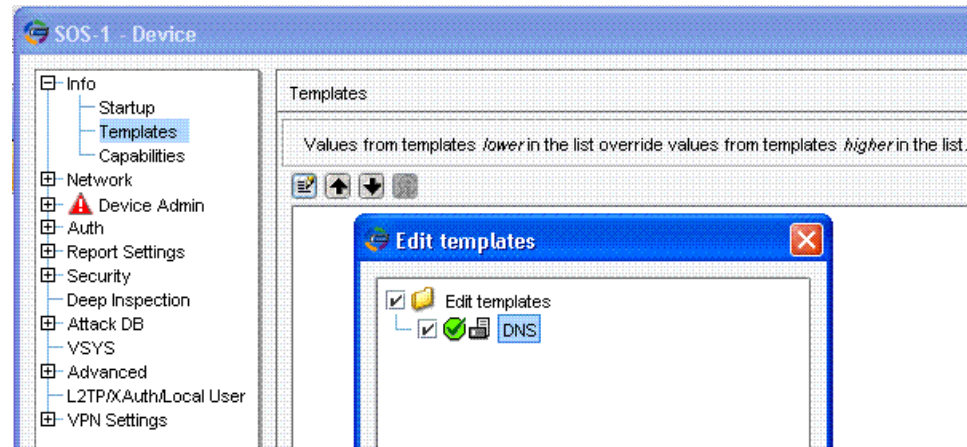
NOTE: In an SA, IC or Junos templates, default values for the configuration parameters are the default values from the schema. Default values are not displayed for configuration parameters that are based on match conditions such as device platform or release version.

Applying the Template

Apply the template as follows:

1. Ensure that the device you want to apply the template to has been added or modeled in the management system.
2. In the navigation tree, select **Device Manager > Devices**, and then double-click the device to open the device editor.
3. In the device navigation tree, select **Info > Templates**. The templates configuration screen appears.
4. Click the Edit icon. The Edit Templates dialog box appears.
5. Select the **DNS** template.

Figure 47: Applying a Template



6. Click **OK** in the Edit Templates dialog box. Then click **Apply** to save your changes to the device configuration.

The template icon appears next to “Network” in the device navigation tree. To confirm your settings, select **Network > DNS > Settings**.

To apply the settings to the device itself, invoke the Update Device directive to push the configuration to the device. See [“Updating Devices” on page 257](#).



NOTE: Select **Retain Template values on Removal** in the templates configuration screen for SA, IC or Junos templates, to retain template values if a template is removed from the device.

Templates and Importing Devices

You can set device values using a template or directly on the device object. If you import a device that already has certain values set, then those values are also stored by NSM.

Where field keys match, imported values override values inherited from the template so that the effective device object configuration matches the device. The live relationship with the template is preserved, however, so that reverting to the previous value removes the latest value inherited from the template.

You can override device settings manually or by using the Template Operations directive.

Promoting a Device Configuration to a Template

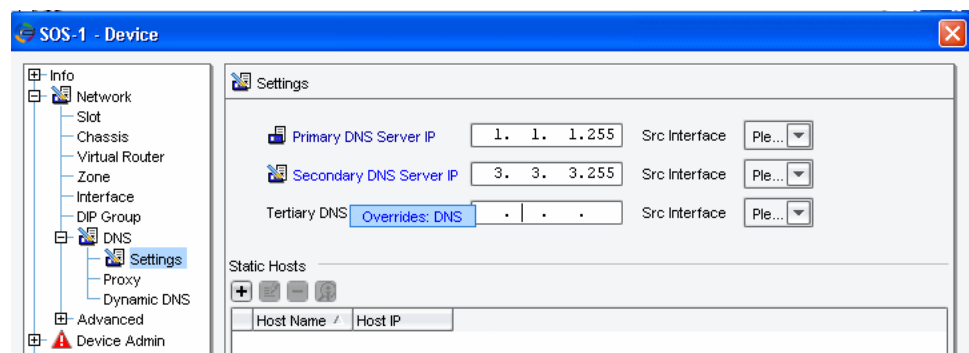
NSM allows you to import the configuration of any Secure Access, Infranet Controller, or Junos device and then convert (promote) it into a template. You can then use that template to make identical configurations on other devices.

To promote a device configuration to a template, in the left panel of the device editor, right-click on the configuration node you want to promote to a template, and select **Promote Template**. In the Select Templates dialog box, select the template to which you want to apply the selected part of the configuration.

Changing Values Inherited from Templates

You can manually override any value inherited from a template in the individual device configuration. All fields inherited from templates appear with a blue template icon next to them. All fields that were inherited from a template but manually overridden have a pale blue marker added to the template icon as shown in [Figure 48 on page 214](#). In this example, the secondary DNS server IP address has had its value overridden, but the primary has not. This example also shows the effect of moving the mouse cursor over the field name of an overridden value; a tool tip message appears showing the name of the template whose value has been overridden.

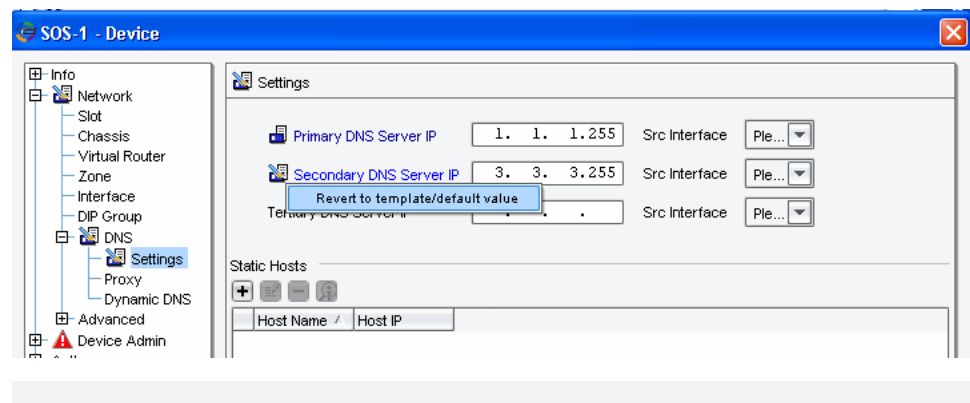
Figure 48: Template Override Icon



For values inherited from the template, the message “From: *template-name*” appears when you move the mouse cursor over the field name.

For any value in the device configuration that was set by a template and overridden, right-click the value and select **Revert to template/default value** to change the device-specific value to the template-defined value (this also changes non-template values back to the default value). An example is shown in [Figure 49 on page 215](#).

Figure 49: Revert to a Template or Default Value



A device-specific configuration value always overrides a template value.

Reverting a Configuration to Default Values of a Template

NSM allows you to revert the device configuration values to that of the template default values for SA, IC and Junos devices. The default value is inherited from the template based on the order of priority in which the templates are applied to the device.

To revert a device configuration to the default values in a template:

1. Select **Configure > Device Manager > Devices** and double-click the device whose configuration settings you want to revert. The Device dialog box appears.
2. Select the **Configuration** tab. The Device configuration tree appears.
3. From the Device configuration tree, right-click the configuration node that you want to revert and select **Revert to template/default value**.
4. Click **OK**. The configuration reverts to the template default values.

Templates and Validation

You can apply a single template to different device types that run different OS versions within the same device family. In some cases, the field values you specify in the template might not be appropriate for all OS versions and device types:

- If the template specifies a field that a device does not support, then the field does not appear in the device editor and is not updated to the device. No validation message appears.
- If the template specifies a field that the device supports, but the value is outside the permitted range for the device, a validation message appears in the Device dialog box. A template value might be valid for one device but invalid for other devices.

As you create and edit template values and fields, NSM validates the values, and might display validation messages. For example, you can configure an IP address in one template

and the netmask for that IP address in another template. However, a validation message might appear when you enter the IP address because the netmask is not specified within that same template.

You can safely ignore a validation message if the missing value is derived from another template that is applied to the device, or if you manually entered the value in the specific device configuration.

Applying Multiple Templates

When applying multiple templates to a single device, you determine the order in which the templates are applied. The highest-priority template is at the end of the template list, and can override values set in any of the lower-priority templates. If more than one template specifies a value for the same field, the value in the highest-priority template takes precedence. The lower the template appears in the template list, the higher priority it has when applying values to a device configuration.

Example: Using Multiple Device Templates

In this example, you create two templates that each configure different values for the same firewall SCREEN option for the untrust zone. The first template, DoS, sets several values in the SCREEN options, including setting the source-based IP session threshold limit to 128 for the untrust zone. The second template, DoS2, sets the source-based IP session threshold limit to 256 for the untrust zone. When you apply these templates to a device, the template with the highest priority overrides the values in the lower-priority template.

1. Create a template that sets SCREEN options for the untrust zone, and then apply the template to a NetScreen-208 device running ScreenOS 5.0:
 - a. In the navigation tree, select **Device Templates**, click the Add icon, and then select **ScreenOS/IDP Template**. The New Device Template dialog box appears.
 - b. In the Info screen, enter **DoS** in the Name field.
 - c. In the template navigation tree, select **Network > Zone**. The Zone configuration screen appears.
 - d. Click the Add icon in the Zone configuration screen and select **Pre-Defined Security Zone — trust|untrust|dmz|global**. The Predefined Zone dialog box appears.



NOTE: Because the untrust security zone is predefined for the device, you must select the Predefined Security Zone option. You can select the Security Zone or Tunnel Zone option only when adding or configuring a user-defined zone.

- e. In the General Properties screen, enter **untrust** in the Name field.
- f. In the zone navigation tree, select **Screen > Denial of Service Defense**. The Denial of Service Defense screen appears.
- g. Select and configure the following options:

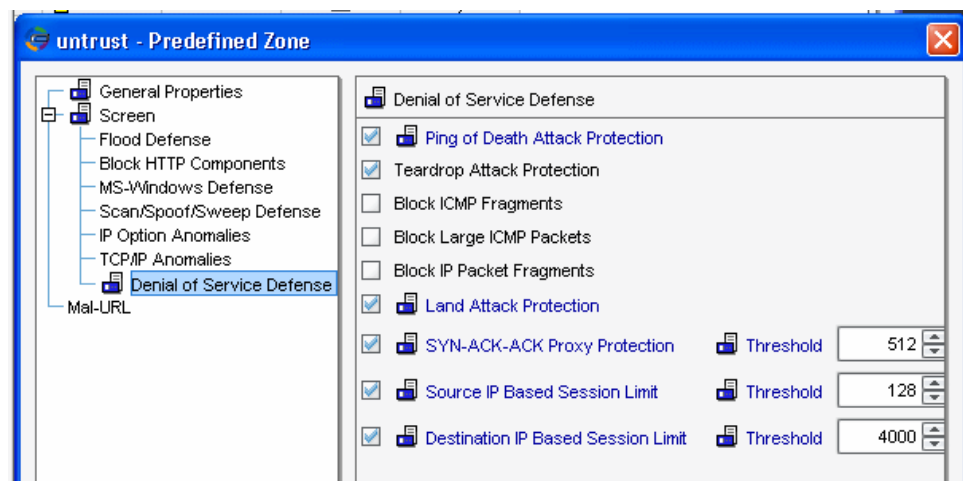
- Select **Ping of Death Attack Protection**, **Teardrop Attack Protection**, and **Land Attack Protection**.
- Select **SYN-ACK-ACK Proxy Protection** and set the Threshold to **512**.
- Select **Source IP Based Session Limit** and set the Threshold to **128**.
- Select **Destination IP Based Session Limit** and set the Threshold to **4000**.

Click **OK** to save the new zone.

h. Click **OK** to save the new device template.

2. Apply the DoS template to a device configuration for a NetScreen-208 running ScreenOS 5.0:
 - a. Add a NetScreen-208 security device to the management system, and model the configuration. Be sure to configure the device as running ScreenOS 5.0.
 - b. In the navigation tree, select **Device Manager > Devices**. Double-click the NetScreen-208 device icon to open the device editor.
 - c. Select **Info > Templates** in the device navigation tree. Click the Edit icon in the Templates screen. The Edit Templates dialog box appears.
 - d. Select the **DoS** template.
 - e. Click **OK** in the Edit Templates dialog box.
3. Verify that the DoS template values have been applied to the device:
 - a. Select **Network > Zone** in the device navigation tree. Double-click the untrust zone. The untrust-Predefined Zone dialog box appears.
 - b. Select **Screen > Denial of Service Defense** and review the values applied by the template, as shown in [Figure 50 on page 217](#).

Figure 50: View Denial of Service Defense Values from DoS Template

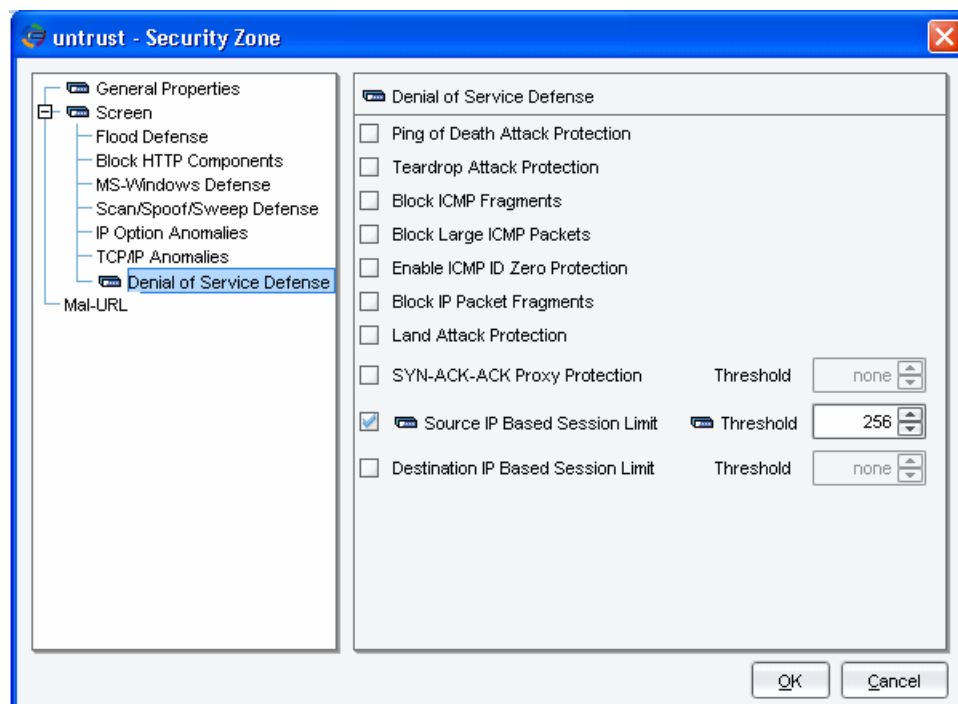


4. Create a second template that sets a different value for a SCREEN option than was set in the DoS template:

- a. In the navigation tree, select **Device Manager > Devices Templates**, click the Add icon, and then select **ScreenOS/IDP Template**. The New Device Template dialog box appears.
- b. In the Info screen, enter **DoS2** in the Name field.
- c. In the template navigation tree, select **Network > Zone**. The Zone configuration screen appears.
- d. Click the Add icon in the Zone configuration screen and select **Pre-Defined Security Zone—trust|untrust|dmz|global**. The Predefined Zone dialog box appears.
- e. In the General Properties screen, enter **untrust** in the Name field.
- f. In the zone navigation tree, select **Screen > Denial of Service Defense**. The Denial of Service Defense screen appears. Select and set the Source IP Based Session Limit Threshold to **256**.

Your settings appear as shown in [Figure 51 on page 218](#).

Figure 51: Configure DoS Defense Settings for the DoS2 Template



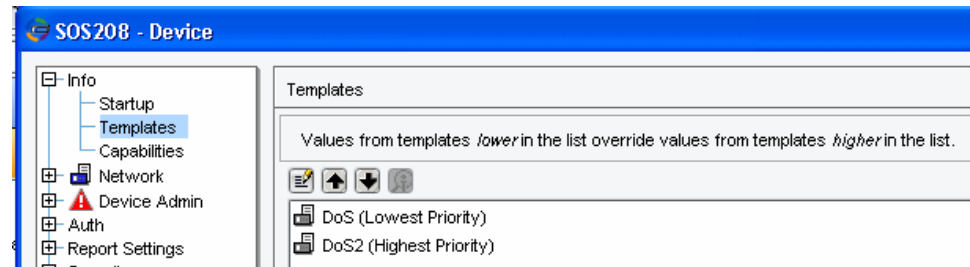
Click **OK** in the New — Predefined Zone dialog box, and then click **OK** in the New Device Template configuration dialog box.

5. Apply the DoS2 template to the NetScreen-208 device:
 - a. In the navigation tree, select **Device Manager > Devices**. Double-click the NetScreen-208 device icon to open the Device dialog box.
 - b. Select **Info > Templates** in the device navigation tree. Click the Edit icon in the Templates configuration screen. The Edit Templates dialog box appears.

- c. Select the **DoS2** template (and keep the DoS template selected).
 - d. Click **OK** in the Edit Templates dialog box.
6. Set the template priority.

Currently, the DoS2 template has the higher priority, which enables it to override any similar values set by the DoS template, as shown in [Figure 52 on page 219](#). The DoS2 template overrides similar values set in the DoS template.

Figure 52: View Template Priority (DoS Highest)

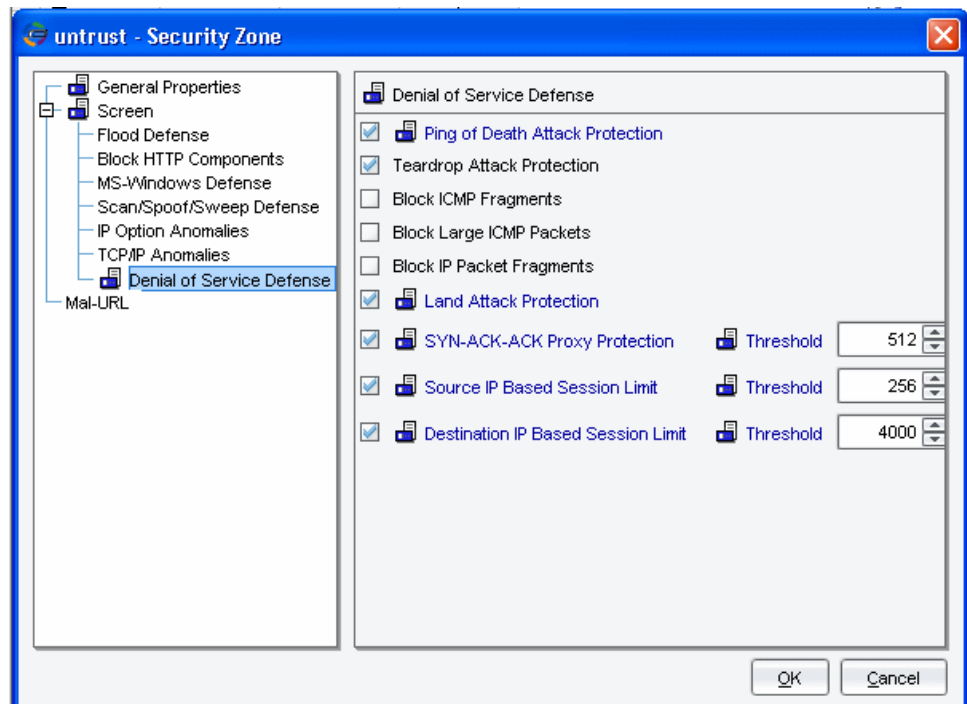


Use the up and down arrows to change the template priority order. This example continues with DoS2 having the higher priority.

7. Verify that the configuration values from the DoS and DoS2 templates have been applied in the device configuration:
 - a. Select **Network > Zone** in the device navigation tree. Double-click the untrust zone. The untrust-Predefined Zone dialog box appears.
 - b. Select **Screen > Denial of Service Defense** and review the values applied by the template, as shown in [Figure 53 on page 220](#).

Although both the DoS and DoS2 templates configured threshold values for the Source IP Based Session Limit field, the higher threshold value from DoS2 appears in the device configuration because you assigned the DoS2 template a higher priority than the DoS template.


Figure 53: View Values from DoS and DoS2 Templates



- c. Verify the origin of each value by pressing the Shift key and moving the mouse cursor over the field name.

For the Source IP Based Session Limit, the message “From template: DoS2” appears, as shown in Figure 54 on page 220.

Figure 54: View DoS2 Value for Source IP Based Session Limit



For the SYN-ACK-ACK Proxy Protection and for Destination IP Based Session Limit, the message “From template: DoS” appears, as shown in Figure 55 on page 220.

Figure 55: View DoS Value for SYN-ACK-ACK Proxy Protection Setting



8. Manually override the SYN-ACK-ACK Proxy Protection value that is set by the template DoS:
 - Clear the SYN-ACK-ACK Proxy Protection check box.
 - Select and set the SYN-ACK-ACK Proxy Protection threshold to be **1000**.
 - The override icon appears next to the field name. Verify that the setting is derived from the device configuration itself and not a template by moving the cursor over

the field name. The message “From object” appears, as shown in [Figure 56 on page 221](#).

Figure 56: View Default SYN-ACK-ACK Proxy Protection Setting



Template Limitations

When configuring and using templates in NSM, be aware of the following limitations.

Maximum of 63 Templates

You can apply a maximum of 63 templates to a single device. However, configuring certain features reduces the maximum number of templates you can apply to a device:

- Cluster or vsys member—Configuring a device as a vsys device or as a member of a cluster reduces the maximum number of templates by one.
- VPNs—Each centrally managed VPN that the device belongs to also reduces the maximum number of templates by one.
- Referenced templates—Each referenced template (a template referred to by another template) reduces the maximum number of templates by one. For example, a device that uses template A, which in turn refers to templates B and C, counts as three templates.

Device Groups

You cannot apply a template to a device group. To use the same template for multiple devices, you must apply the template to each device individually with the Template Operations directive. See [“Adding Device Groups” on page 194](#) for details about device groups.

Default Values

Default values do not appear when editing a template because many default values depend on the OS version and device platform.

Predefined Device Data

Templates do not automatically include any predefined device data, such as zones, interfaces, or virtual routers. To create a template that refers to a specific predefined entity, you must create the entity in the template.

For example, to create a template that refers to the ethernet1 interface on a ScreenOS device:

1. In the template navigation tree, select **Network > Interface**.
2. Click the Add icon and select **Predefined Interface**. The Physical Interface dialog box appears.

3. For Name, enter **ethernet1**.



NOTE: When creating or editing predefined interfaces in a template, you must use the exact name for each interface.

When adding an entity in a template, ensure that the menu option you select is appropriate for the predefined entity. Choose the menu option that includes the name of the predefined entity you are creating.

For example, to create a template that refers to the mgt zone on a ScreenOS device:

1. In the template navigation tree, select **Network > Zone**.
2. Click the Add icon and select **Predefined Functional Zone — mgt/vlan**. The Zone dialog box appears.
3. Enter **mgt**.

List Key Fields

List key fields are used for matching a configuration object in a list of similar objects. They are read-only. You cannot edit list key parameters that are derived from a template. For example, a zone name uniquely identifies a zone in the list of zones for the device. If you create a zone in a template and apply the template to a device, you cannot change the zone name in the device configuration. You must first delete the template-derived zone, and then create a new zone.

A list or table entry in the configuration can contain multiple list key fields. For example, in the routing table for a ScreenOS device, multiple fields (including IP address/netmask, interface, next-hop, vsys, and so on) uniquely identify a particular route entry.

Specifying the Order of List Entries

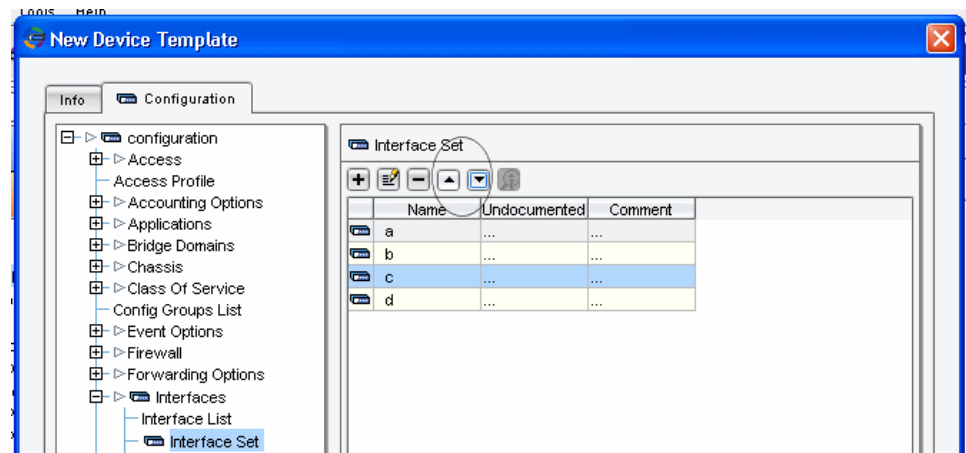
When configuring a device through NSM, you can enter most list or table entries in any order. However, in some cases the order of list or table entries is significant to device behavior. For example, in a routing policy or firewall filter, you define terms that are analyzed sequentially. Also, when you create a named path in dynamic MPLS, you define an ordered list of the transit routers in the path, starting with the first transit router and ending with the last one.



NOTE: The ordering of list entries is a detailed point and of low significance to most users. Skip this section if ordering of list entries is not significant to you.

To specify a sequence in which the list or table entry order matters, select the entry in the template and then use the up and down arrows at the top of the dialog box. The up and down arrows appear on any table where the order of parameters is important and can be altered. [Figure 57 on page 223](#) shows an example.

Figure 57: Up and Down Arrows for Changing the Sequence of a List



NSM allows you to add, insert, or delete entries in ordered lists. NSM uses some simple rules to establish the order of configuration entries when they are changed on the device and then imported into NSM, or changed in the NSM template, and then updated on the device.

The default order of entries in a list depends on their order in the template, their order in the device object, their order in a configuration group, and their order on the device itself.

For details about how the order of list entries in a template is affected by configuration groups, see [“Using Configuration Groups” on page 235](#). The following sections describe how the order of lists is affected by template position, device object data, the order on the device itself, and the operations that affect the sequence of ordered list entries.

- [Combining Template Data with Device Object Data on page 223](#)
- [Operations That Change the Sequence of Ordered Lists on page 224](#)
- [Rules for Reordering Lists on page 225](#)
- [Examples of Reordered Lists on page 225](#)
- [Identifying Ordered List Entries That Do Not Match the Template or Configuration Group Order on page 228](#)

Combining Template Data with Device Object Data

If a template and the device each have data for an ordered list, by default, NSM places the template entries after the device entries. For example, suppose a device object has two entries for an ordered list:

```
D1 D2
```

A template has two list entries for the same list:

```
T1 T2
```

When you apply the template to the device, NSM applies the default ordering, which is to place the template entries after the device entries:

```
D1 D2 T1 T2
```

Now push the configuration to the device, and then connect to the Web UI of the device and reorder the list entries, such that the list entries that came from the template are reversed:

```
D1 D2 T2 T1
```

Now consider what happens when you reimport the configuration from the device. To preserve the relationship between the template and the device object, the T1 and T2 entries must continue to refer to the template. Yet the order must be preserved to correspond to the device. Because another device using the template might give the entries a different order, NSM must keep the ordering information on a per-device basis.

Operations That Change the Sequence of Ordered Lists

By default, template entries appear at the end of each ordered list in template priority order, followed by entries specified in the device object data. Through changes on the device outside of NSM, or through user action, template-provided list entries can be reordered and intermixed with list entries provided in the device object. The following operations can affect the order of list entries in the device object:

- Adding a template
- Deleting a template from the device object
- Adding or inserting new list entries in the template
- Adding or inserting new list entries in the device
- Deleting entries from an ordered list in the device
- Deleting entries from an ordered list in the template
- Deleting template-provided entries in the device
- Reordering entries in the device
- Reordering entries in the template



TIP: Where possible, avoid mixing reordering operations with other operations that affect the order of list entries. The results might not be what you intended. When reordering is combined with adding and deleting entries, there is no single obvious way to define the behavior so it always matches your intention.

Rules for Reordering Lists

When the template order changes, NSM uses best effort to apply the new sequence to the device object using the following rules:

- NSM uses the device order before the operation as the starting point.
- If, prior to applying the change, a contiguous subsequence of parameters in the template matches a contiguous subsequence of parameters in the device, then NSM applies the new template order for the subsequence to the device,
- Entries added in a template are placed in the same sequence in the device; that is, an entry follows the entry in the device that precedes it in the template, even if that entry has been moved in the device.

Examples of Reordered Lists

Examples of the rules for reordering lists follow.

Examples That Reorder a Common Sequence

Example 1: In the following example, the device has inherited the template entries and maintained the same order as in the template. The user then changes the template order:

Before:

Template Sequence	A	B	C		
Device Sequence	A	B	C	1	2
Matching Subsequence	A	B	C		

Change:

Now reverse the first three items in the template sequence. Because the reordering takes place within what was the matching subsequence, the new sequence is transferred to the device:

After:

Template Sequence		C	B	A	
Device Sequence		C	B	A	1 2

Example 2: In the next example, the device has again inherited entries from the template, and then inserted an additional entry within the inherited list. The user then changes the template order.

Before:

Template Sequence	A	B	C		
Device Sequence	A	B	1	C	2
Matching Subsequence	A	B			

Change:

Now reverse the first two items in the template sequence. Because the reordering takes place within what was a matching subsequence, the new sequence is transferred to the device:

After:

Template Sequence	B	A	C		
Device Sequence	B	A	1	C	2

Example 3: The following example shows entries inserted into the list on the device such that there is no matching subsequence. The user then reorders the entries in the template:

Template Sequence	A	B	C		
Device Sequence	A	1	B	2	C
Matching Subsequence	None				

Change:

Now reverse the first two entries in the template sequence. Because there is no matching subsequence, the order in the device remains unchanged. The rule is implemented this way because the user's intention is unclear; for example, should NSM reverse A and B? If so, should 1 come before B, after A, or between B and A?

After:

Template Sequence		B	A	C	
Device Sequence (no change)	A	1	B	2	C

Examples Showing Entries Inserted in the Template

Example 1: In the following example, the device has inherited the template entries and maintained the same order as in the template. The user then inserts a new entry into the template.

Template Sequence	A	B	C		
Device Sequence	A	B	C	1	2
Matching Subsequence	A	B	C		

Change:

Now add an entry to the template. The new entry is added to the device in the same sequence as it was added in the template. That is, the new entry follows entry C in the template, so it follows entry C in the device.



NOTE: The inserted entry can be anywhere in the list. It does not have to go at the end. The same rule still applies.

After:

Template Sequence	A	B	C	D		
Device Sequence	A	B	C	D	1	2

Example 2: In the following example, the device has reordered the entries that it inherited from the template. The user then inserts a new entry into the template.

Template Sequence	A	B	C		
Device Sequence	C	1	2	A	B

Change:

Now add an entry to the template. The new entry is added to the device in the same sequence as it was added in the template. That is, the new entry follows entry C in the template, and it still follows entry C in the device, even though the device order has changed.

In this example, it is not obvious what the user intended. The user might have wanted to place D after 1 or 2. In this case, NSM makes a reasonable attempt and places it after C.

After:

Template Sequence	A	B	C	D		
Device Sequence	C	D	1	2	A	B

Example 3: In the following example, the device has inserted entries 1 and 2 before the template entries, and deleted entry C:

Template Sequence	A	B	C	
Device Sequence	1	2	A	B

Change:

Now add an entry to the template. The new entry is added to the device in the same sequence as it was added in the template. In this case, however, entry C has been deleted from the device, so the inserted entry follows entry B.

After:

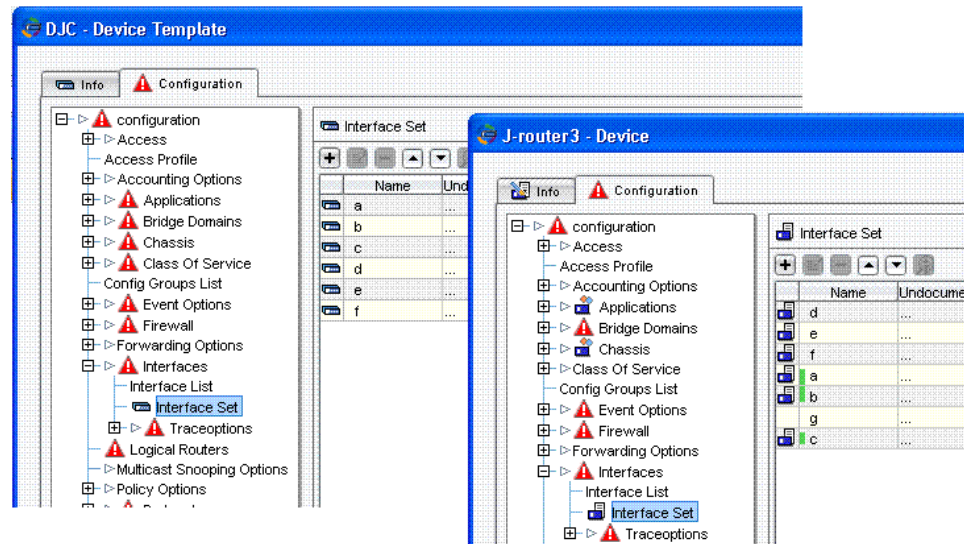
Template Sequence	A	B	C	D	
Device Sequence	1	2	A	B	D

Identifying Ordered List Entries That Do Not Match the Template or Configuration Group Order

In ordered lists, green highlights in the first data column indicate that entries in the regular configuration are not in the order specified in the template. NSM finds the longest common subsequence between the template order and the list order. Entries that are part of the longest common subsequence are not highlighted. Entries from the template that are not part of the longest common subsequence have green highlights.

Figure 58 on page 229 shows a template and a device configuration to which it has been applied. The template provides six entries in the order a, b, c, d, e, f. In the regular device configuration, list entry g has been added directly (as shown by the lack of any data origin icon), the template has been added, and then the list has been reordered.

Figure 58: Identifying Ordered List Entries that Do Not Match the Template Order



In the device configuration in Figure 58 on page 229:

- The subsequence d, e, f is not highlighted because it is the longest common subsequence between the device configuration and the template.
- g is not highlighted because it does not come from the template.
- a and b have a continuous green highlight because they represent a common subsequence, though not the longest.
- c has a single-entry out-of-order mark because it is adjacent to neither of its neighbors in the template.



NOTE: If multiple subsequences tie for the longest common subsequence, then NSM picks either one but not both.

NSM recomputes the longest common subsequence each time the list is reordered and makes changes to out-of-sequence highlighting accordingly.

Using the Template Operations Directive

The Template Operations directive allows you to add templates or remove templates for multiple devices at one time, and to validate configurations after changes.



NOTE: The Template Operations directive only updates the configuration database. To apply changes to devices, you must use the Update Device directive.

Figure 59: Template Operations Directive

Template Operations

Select OS Name: J-Series

Select Device(s)

- ☒ J-1
- ☒ J-2

Select Templates

☐ Add templates with lowest priority
☐ Add templates with highest priority
☐ Remove templates
☒ Don't change templates

☐ Remove conflicting device values
☐ Report irrelevant template values
☐ Report conflicts with other templates
☐ Validate

Previous Apply Changes

The Template Operations directive displays a dialog box that is divided into the following sections:

Select OS Name Section

Select a device family from the Select OS Name list to determine which set of templates and devices to show.

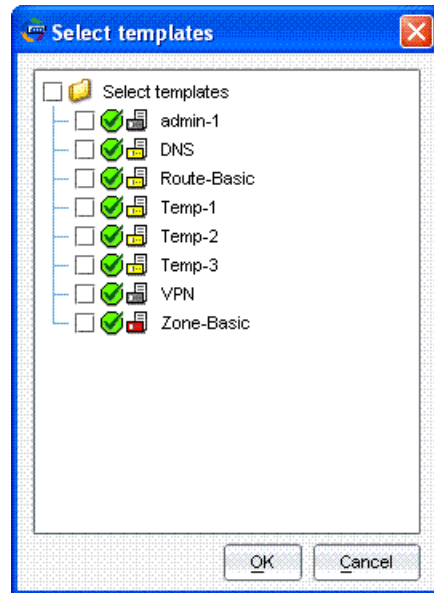
Select Devices Section

In this section, select one or more devices for template operations.

Select Template Section

Select one or more templates to apply to the selected devices. Use the edit button to open the Select Templates dialog box. Check one or more check boxes to select templates.

Figure 60: Select Template Dialog Box



After you have selected templates, click **OK**. You can use the up and down arrow buttons in the main dialog box to order the selected templates as you wish.

Template Operation Section

The next section features four buttons:

- **Add templates with lowest priority**—Adds the templates to each device's template list with the lowest priority. Added templates will be lower priority than templates previously assigned to the device. Templates with higher priority can override the values applied by these templates.
- **Add templates with highest priority**—Adds the templates to each device's template list with highest priority. Added templates will be higher priority than templates previously assigned to the device. Values in these templates will override values applied by lower-priority templates.
- **Remove templates**—Removes all selected templates from each selected device.
- **Don't change templates**—Makes no changes to devices in the database. This setting is useful if you want to perform a validation, run a report, or clear overrides without changing the template assignments.

Options Section

The last section provides optional operations:

- **Remove conflicting device values**—Overrides any device settings that override template values provided by the selected templates. Normally, template values do not override manually set values.
- **Report irrelevant template values**—Reports any values that are set in templates but that are not used on the selected devices. A template might provide values for features

that aren't available on every device. For example, wireless configuration information is not relevant to devices that do not provide wireless functionality.



NOTE: If the template specifies a field that a device does not support, the field does not appear in the Device dialog box and is not applied to the device. No validation message appears. You can see these values by checking the Report irrelevant template values check box.

If the template specifies a field that the device supports, but the value is outside the permitted range for the device, a validation message appears in the Device dialog box.

- **Report conflicts with other templates**—Reports any values that conflict between the selected templates and existing templates assigned to the device. If undesired conflicts exist, you might need to modify the templates to get the configuration you want.
- **Validate**—Checks that the configured device (after any changes) is a valid configuration and reports any errors.

Template Operations Box Recommended Workflow

The Template Operations dialog box can be used in many ways. This section describes one recommended workflow.

Step 1: Look at the Effect of Planned Changes Before Making Them

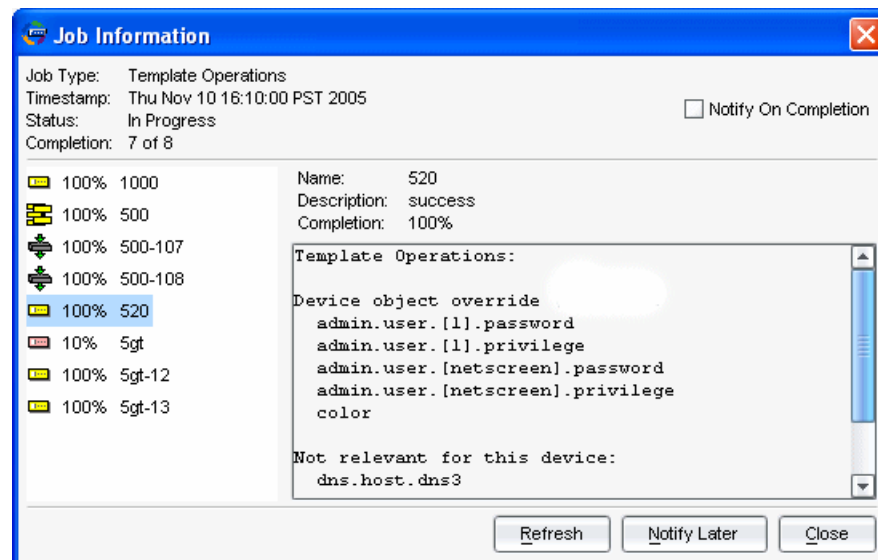
1. From the Device Manager launchpad, click the **Template Operations** icon, and select a device family from the Select OS Name list.
2. Select one or more devices to validate templates against.
Example: 5gt, 5gt-12, 5gt-13, 500, 520, 1000
3. Select the Edit icon under the Edit Templates header to display a list of templates.
4. Select one or more templates to validate against the devices.
Example: VPN, DNS, Route-Basic, Zone-Basic, admin-1
5. Click **OK**.
6. If you have selected more than one template, use the up and down arrows to order the templates. Templates higher in the list will be assigned first. Templates lower in the list will be assigned later. Later templates can override the settings of earlier templates.
7. Select **Don't change templates**.
8. If you want, you can select one or more of the following validation and reporting options:

- Report irrelevant template values (optional)—Reports template settings that are not relevant to the selected device or devices.
- Report conflicts with other templates (optional)—Reports settings that conflict between two or more templates.
- Validate (recommended)—Validates the device configuration after appropriate changes to the templates and the removal of conflicting values (if needed).
- Do not select Remove conflicting device values because it might alter device values.

9. Click **Apply Changes**.

Step 2: Review Results in Job Information Dialog Box

Figure 61: Template Operations Job Information Dialog Box



Consult the reports and error messages generated in Step 1. Resolve any conflicts, missing assignments, or other errors as desired. Repeat steps 1 and 2 until you are satisfied with your planned changes.

Step 3: Apply Templates and Clear Overrides

This step updates the NSM database, but does not push the new settings to the device.



CAUTION: You can easily reverse adding templates, but there is no automatic way to restore conflicting device values that have been removed. Be very careful not to remove values you want to keep.

Repeat the operations specified in Step 1, but specify one of the Add templates buttons. If desired, also check the Remove conflicting device values check box.

Removing Templates with the Template Operations Directive

To remove one or more templates from one or more devices, follow these steps:

1. From the Device Manager launchpad, click the **Template Operations** icon, and select a device family from the Select OS Name list.

2. Select one or more devices to remove templates from.

Example: **5gt**

3. Select the Edit icon under the Edit Templates header to display a list of templates.

4. Select one or more templates to remove.

Example: DNS

5. Click **OK**.

6. Select **Remove templates**.

7. Click **Apply changes**.

Exporting and Importing Device Templates

Templates can be exported to a text file, and then reimported into another template. When an exported template is imported into a template, the values in the exported template overwrite existing values.

The exported template contains only device settings. It does not contain any policies or objects.

Exporting a Device Template

To export a device template:

1. From the Device Manager launch pad, select **Export/Import**, and then select **Export Device Template to File**
2. In the Export Config to File dialog box, select the template you want to export, and then click **OK**.
3. When you see the file generated in the Job Information window, select **Save Selected** and give the file a name.

You do not have to highlight the configuration file text. You can use any file extension.

Importing a Device Template

To import device template configurations into a template, follow these steps:

1. From the Device Manager launch pad, select **Export/Import**, and then select **Import Device Template Config From File**.
2. Select the templates you want the saved template settings to be applied to.
3. Select the saved template you want to import.

The settings in the saved template are imported into the NSM template.

Refer to the *Network and Security Manager Online Help* for detailed procedures.

Using Configuration Groups

Configuration groups are a Junos concept that you can use in NSM to help speed configuration of Junos devices.

In Junos, configuration groups allow you to create a group containing Junos device configuration statements and to direct the inheritance of that group's statements in the rest of the configuration. The same group can be applied to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.

In NSM, you can create configuration groups in the UI. The configuration data you provide in the UI becomes configuration statements when you push the configuration to the Junos device,

Configuration groups allow you to create smaller, more logically constructed configuration files, making it easier to configure and maintain the Junos device configuration. For example, you can group configuration information that is repeated in many places in the configuration, such as when configuring interfaces, and thereby limit updates to just the group.

You can also use wildcards in a configuration group to allow configuration data to be inherited by any object that matches a wildcard expression.

The configuration group mechanism is separate from the grouping mechanisms used elsewhere in the Junos configuration, such as Border Gateway Protocol (BGP) groups. Configuration groups provide a generic mechanism that can be used throughout the configuration.

Configuration groups use true inheritance, which involves a dynamic, ongoing relationship between the source of the configuration data and the target of that data. That is, a live relationship exists between the configuration group and the device object. Any change in the configuration group immediately takes effect in the device object.

Data values changed in the configuration group are automatically inherited by the target. The target need not contain the inherited information, although the inherited values can be overridden in the target without affecting the source from which they were inherited.

NSM enables you to create, edit, and view configuration group definitions. Specifically, NSM supports the following operations related to configuration groups:

- Create, edit, and view configuration group definitions.
- Display the effective value of applying the configuration group.

- Identify the origin of values derived from configuration groups; that is, identify which configuration group a value came from.
- Support the specification of configuration groups in templates. See [“Using Configuration Groups with Templates” on page 243](#) for details.

This section contains the following topics:

- [Creating and Editing Configuration Groups on page 236](#)
- [Applying a Configuration Group on page 239](#)
- [Excluding a Configuration Group on page 240](#)
- [Editing a Device Object That Uses Configuration Groups on page 241](#)
- [Deleting a Configuration Group on page 242](#)
- [Adding Ordered List Entries Using Configuration Groups on page 242](#)
- [Reordering Lists on page 243](#)

Creating and Editing Configuration Groups

You access configuration groups through the device editor. The main panel shows the names of existing configuration groups. Use the NSM buttons at the top of the display area to create, edit, or delete configuration groups.

Creating a Configuration Group

The following example shows how to create a configuration group. When you apply this configuration group, it sets the speed of a specific interface to 100 Mbps, and the speed of all other configured interfaces to 1 Gbps by using a wildcard mechanism.

1. Double-click the device in the Device Manager and select the **Configuration** tab.
2. In the configuration tree, select **Config Groups List**
3. Click the Add icon and select **Regular**.

The New dialog box appears. It looks like the device configuration tree, except that it does not have a Config Groups List branch, because you cannot define configuration groups recursively.

4. Give the configuration group a name, for example **set-speed**.
5. Enter some configuration data. For example, configure some interface parameters to set the interface speed:
 - a. Expand **Interfaces**, and then select **Interface**.
 - b. Click the Add icon to display the interface configuration screen.

- c. Name the interface, for example **ge-0/0/1**, which identifies the ge interface in slot 0, PIC 0, port 1.

- d. In the Speed field, set the speed to **100m**.

A tooltip icon appears next to the Speed field. This icon indicates that its value has been set in the configuration group.

- e. Click **OK** to save the interface definition.



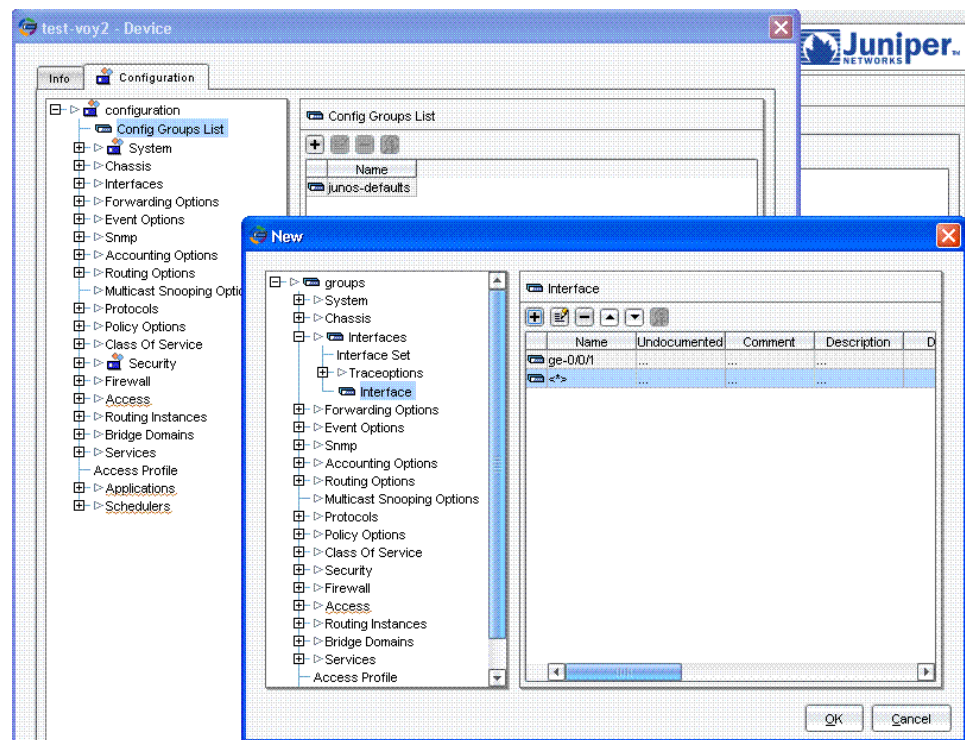
NOTE: After saving a configuration group entity, you cannot change its key field. Specifically, after saving an interface definition, you cannot change its name.

- f. Click the Add icon again and name the second entry with a wildcard character **<*>**. Include the angle brackets, because they are required by the Junos syntax.

- g. In the Speed field, set the speed to **1g** and click **OK**.

The configuration group icon appears next to the two interface entries in the group, and next to each element in the tree above the interface entries. See [Figure 62 on page 238](#). Mouse over the icons to see a summary of what has been set and where the information came from.

Figure 62: Adding a Configuration Group



- Click **OK** to save the configuration group. The new configuration group appears in the Config Groups List.



NOTE: After saving a configuration group, you cannot change its name.

Editing a Configuration Group

You can edit configuration groups before or after applying them. If you edit the group after applying it, then the new edits take effect immediately in the device object.

When you create or edit a configuration group, values specified for the configuration group are marked with icons and tooltips similar to those used to highlight template values. Move your mouse cursor over a tooltip icon to gather information about the entity it describes. Right-click the item and select **Revert to template/default value** to undo the effect of setting the value in the configuration group.

You cannot create groups that start with “junos.”

Validating a Configuration Group

Validation within the configuration group does not apply required field constraints (similar to when editing a template), but the operating system version and platform are available, so more validation is done for configuration groups in devices than for templates.

Ordered Lists and Wildcard Matching

In configuration groups, the order of list entries is significant. All lists in configuration groups are displayed in the order in which they are defined. To change the order of the list, you must do so explicitly by using the up and down arrows at the top of the main display area.

The order of lists is significant because configuration group wildcard matching is done starting from the first configuration group entry and stopping after the first match. Consider a configuration group containing the following list of interface definitions, specified in the order shown. For each list item, the first entry is the interface name, and the second an assigned value for MTU:

<ge-0/0/0>	mtu 4k	<ge-*>	mtu 5k
<*>	mtu 8k		

When you apply this configuration group to a device object that already has **ge-0/0/0** and **ge-0/0/1** interfaces configured in it, but with no defined value for MTU, then the first match for **ge-0/0/0** is the **ge-0/0/0** entry in the configuration group, so the effective MTU is 4k. The first match for the **ge-0/0/1** entry is the **<ge-*>** entry in the configuration group, so the effective MTU is 5k.

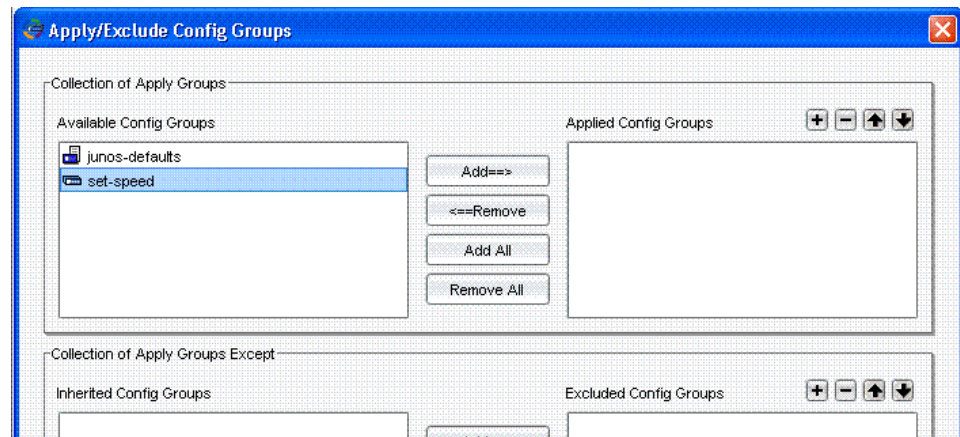
Applying a Configuration Group

You can apply a configuration group at any level in the configuration tree. A triangle next to a branch represents a point at which a configuration group can be applied. By selecting that branch and applying a configuration group to it, all parameters set under that branch in the configuration group are applied to the device object configuration,

The following example applies the configuration group defined in “[Creating a Configuration Group](#)” on page 236 to the device object configuration.

1. In the device object configuration tree, right-click **Interfaces**.
2. Select **Apply/Exclude Config Groups** from the list. The Apply/Exclude Config Groups dialog appears, as shown in [Figure 63](#) on page 240.

Figure 63: Applying a Configuration Group



3. Select the desired configuration groups from the Available Config Groups list, and then click **Add**. The Available Config Groups list includes all configuration groups created in the device object.

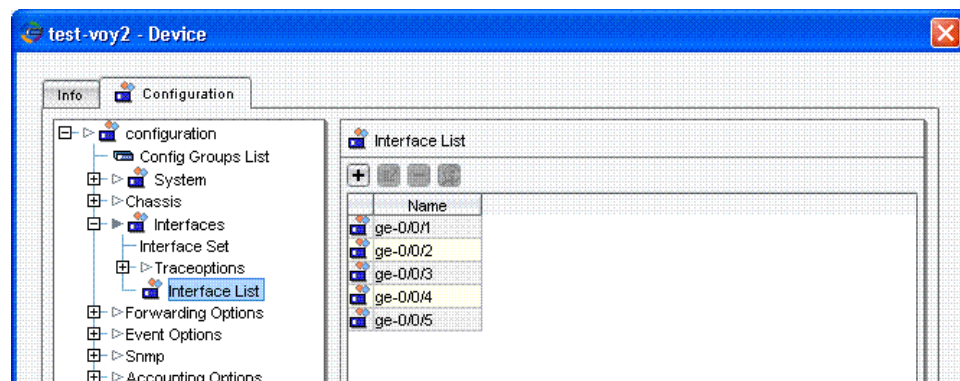
The configuration group and icon move to the Applied Config Groups list.

4. Click **OK** to apply the configuration group.



NOTE: The first configuration group in the list has the highest priority. This convention is the reverse of the ordering for templates, where the last template in the list has the highest priority.

Figure 64: Configuration Group Applied



Excluding a Configuration Group

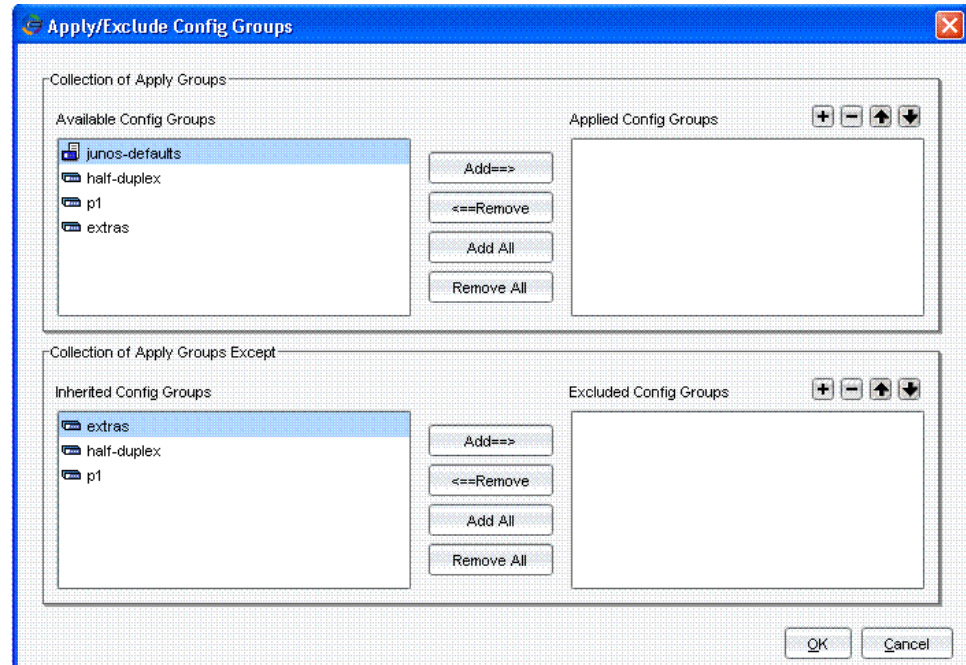
Configuration groups applied at a higher level in the configuration can be explicitly excluded at lower levels. To exclude a configuration group from a branch of the configuration, follow these steps:

1. Right-click on the branch and select **Apply/Exclude Config Groups** from the list.

The Apply/Exclude Config Groups dialog box appears. All configuration groups applied at higher levels in the configuration appear in the Inherited Config Groups list.

Figure 65 on page 241 shows an example.

Figure 65: Excluding a Configuration Group



2. In the Inherited Config Groups list, select the configuration groups you want to exclude, and then click **Add**,

The selected configuration group names move to the Excluded list.

3. Click **OK** to exclude those groups from that part of the configuration.

Editing a Device Object That Uses Configuration Groups

After you apply the configuration group, tooltip icons identify where configuration groups have affected the configuration. You can mouse over these items to display information about them.

When you edit an entity that was derived from a configuration group, the new value overrides the value derived from the configuration group. The tooltip icon changes so you can easily identify which entries have been overridden,

To revert to the configuration group value for a specific field, right-click on the field, and select **Revert to template/default value**.



NOTE: Nonwildcard list entries derived from configuration groups cannot be deleted in the device object.

Deleting a Configuration Group

When you delete a configuration group, all applied and excluded data is removed from the configuration.

To delete a configuration group, follow these steps:

1. In the Device Manager, open the device in which you want to delete a configuration group and select **Config Group List** from the device configuration tree.
2. Select the configuration group you want to delete and click the Delete icon.

The configuration group is deleted.

Adding Ordered List Entries Using Configuration Groups

List entries added by configuration groups into ordered lists in the device object appear in a specific order determined by Junos convention. By default, entries from templates appear first, followed by regular configuration data, followed by entries created by the configuration groups in the order in which the configuration groups are applied.

Consider two configuration groups J and K:

- Configuration group J contains the following list items in the stated order:

```
c
a
b
```

- Configuration group K contains the following list items in the stated order:

```
r
a
z
```

- The device object already has the following list items configured in the stated order:

```
x
z
```

After applying groups J and K, the entries appear in the following order in the device configuration:

```
x    # from the regular configuration
z    # merger of data from regular configuration and configuration group K
```

```
(regular configuration takes precedence)
c  # from configuration group J
a  # merger of data from configuration groups J and K (J takes precedence)
b  # from configuration group J
r  # from configuration group K
```

For unordered lists, this rule is unimportant. For ordered lists such as policies and configuration group definitions, the rule is important.

Reordering Lists

Ordered lists that are inherited from a configuration group follow the same rules as ordered lists inherited from a template:

- NSM uses the device order before the operation as the starting point.
- If, prior to applying the change, a contiguous subsequence of parameters in the configuration group matches a contiguous subsequence of parameters in the device, then NSM applies the new configuration group order for the subsequence to the device.
- Entries added in a configuration group are placed in the same sequence in the device; that is, an entry follows the parameter in the device that precedes it in the configuration group, even if that entry has been moved in the device.

For details and examples, see [“Specifying the Order of List Entries” on page 222](#).

NSM uses the same mechanism for identifying ordered list entries that do not match the configuration group order as is used for templates. Green highlights in the first data column indicate entries in the regular configuration that are not in the order specified in the configuration group. See [“Identifying Ordered List Entries That Do Not Match the Template or Configuration Group Order” on page 228](#) for details and examples.

Using Configuration Groups with Templates

If a field in a device object can inherit from both a template and a configuration group, then the template value is used. NSM first expands the template, and then expands the configuration group. Regular configuration data has precedence over template or configuration group data.

For simplicity, we recommend that you use either templates or configuration groups for each part of the configuration, but not both. Avoid applying a configuration group in a device object to part of the configuration that also has values applied from a configuration group that is part of a referenced template.

In some cases, however, it can be desirable to mix templates with configuration groups.

One practical use of mixing configuration groups with templates is so that you can use a wildcard mechanism in a configuration group to assign a common value to like parameters across multiple list entries such as interfaces.

Sharing Configuration Group Definitions Across Multiple Devices

Templates provide a natural mechanism for sharing configuration group definitions across multiple devices. The use of configuration groups enhances the capabilities of templates, for example, by allowing you to set the same field value on all interfaces across multiple devices by using the configuration group wildcard feature.

This example shows defining the configuration in NSM using a configuration group in a template.

The example uses a template to apply an MTU value of 3K to all interfaces on devices to which it is applied, except those interfaces explicitly assigned a different MTU value in the template itself or in the device object. Settings in the regular configuration take precedence over those set in the template which, in turn, take precedence over those set in the configuration group. The following table shows the MTU values assigned by each mechanism in this example, and the final outcome in the device itself.

Interface Name	Value Set in Configuration Group by Wildcard	Value Set in Template	Value Set in Regular configuration	Final Value Applied to Device
	Order of increasing precedence —>			
fe-0/0/0	3K	6K	5K	5K
fe-0/0/1	3K	4K		4K
fe-0/0/2	3K			3K
All other interfaces	3K			3K

To create this configuration, follow these steps:

1. Create a template containing a configuration group that will apply an MTU value of 3K to all devices to which the configuration group is applied:
 - a. In the Device Manager, select **Device Templates**.
 - b. Click the Add icon and select **Junos Template**.
 - c. From the Junos Product Series list, select **Junos J Series** to create a new template for J Series devices.

- d. Click **Next** and then **Finish** to create a new template for J Series devices.
- e. In the Name field, enter a new name for the template, for example, **set-mtu**.
2. Open the template by double-clicking it in the Device Template window, and enter the following information:
 - a. In the Configuration tab of the new template, select **Config Group List**.
 - b. Click the Add icon and select **Regular** from the list.
 - c. In the New dialog box, name the new group, for example, **group1**.
 - d. Expand **Interfaces**, select **Interface**, and click the Add icon.
 - e. Name the interface with the wildcard character by typing `<*>` in the Name field.



NOTE: Be sure to include the angle brackets, because they are required by the Junos syntax.

- f. Set the Mtu field to **3072**. Click **OK** to finish the interface definition.
- g. Click **OK** again to finish creating the configuration group.
3. Apply the configuration group to the template:
 - a. In the Configuration tab of the template, right-click **Configuration** and select **Apply/Exclude Config Groups** from the list.
 - b. Select **Group1** in the Available Config Groups list and click **Add** to move it to the Applied Config Groups list.
 - c. Click **OK** to apply the configuration group.
4. Configure some interfaces explicitly in the template:
 - a. Expand **Interfaces** and select **Interface List**.
 - b. Click the Add icon and select **Physical Interfaces** from the list.
 - c. In the Set Slot Configuration dialog box, set the slot range to **0**, the PIC range to **0**, the port range to **0-1**, and click **OK**.

The new interfaces show in the Interface List for the template.

- d. Set the MTU for **fe-0/0/0** to 6K:
 - i. Click on the **fe-0/0/0** interface in the Interface List and click the Edit icon to open the interface.
 - ii. Set the Mtu field to **6144** and click **OK**.
- e. Repeat the previous step for **fe-0/0/1** and set the MTU value to 4096.

The equivalent Junos configuration syntax for the template looks like this:

```
groups {
  group1 {
    interfaces {
      <*> { mtu 3k; } # wildcard matches all interfaces
    }
  }
}
interfaces {
  apply-groups group1; # apply-groups takes a list
  fe-0/0/0 { mtu 6k; }
  fe-0/0/1 { mtu 4k; }
}
```

5. Configure some interfaces in the device object.

In this example, we set the MTU for **fe-0/0/0** to 5120, and create **fe-0/0/1** and **fe-0/0/2** with blank MTU values:

- a. In the Device Manager, select **Devices**.
- b. Select the device and click the Edit icon.
- c. Select the **Configuration** tab, and then expand **Interfaces**, and select **Interface List**.
- d. Click the Add icon and select **fe Physical Interfaces** from the list.
- e. In the Set Slot Range Configuration dialog box, set Slot Range to **0**, set Pic Range to **0**, Port Range to **0-2**, and click **OK**.
Three new interfaces populate the interface list.
- f. Select **fe-0/0/0** and click the Edit icon.
- g. Set the Mtu field to **5120** and click **OK**.

The equivalent Junos configuration syntax for the device object looks like this:


```
# regular config
interfaces {
  fe-0/0/0 { mtu 5k; }
  fe-0/0/1 { ... }
  fe-0/0/2 { ... }
}
```

6. Apply the template to the device:
 - a. Click the **Info** tab of the device, and click **Templates**.
 - b. Click the Edit icon to display the Edit Templates dialog box.
 - c. Check the box next to the template you just created and click **OK** to apply the template to the device.
7. Check the device object configuration:
 - a. Select the Configuration tab.
 - b. Expand Interfaces, if necessary, and click Interface List.
 - c. **fe-0/0/0** has an MTU of 5120, because the regular configuration takes precedence over both the value in the template and the value in the configuration group.
fe-0/0/1 has an MTU of 4096, because the template value takes precedence over the value in the configuration group.
fe-0/0/2 has an MTU of 3072 because it was not explicitly defined in the regular device object configuration or the template. Therefore, it takes the value from the wildcard setting in the configuration group.
8. Push the configuration to the device using the Update Device directive:
 - a. In the Device Manager, click **Devices**.
 - b. Right-click the device, and select **Update Device** from the list.

The equivalent Junos configuration syntax received by the device looks like this:

```
groups {      # CG defn from template
  group1 {
    interfaces {
      <*> { mtu 3k; } # wildcard matches all interfaces
    }
  }
}
# regular config
interfaces {
  apply-groups group1; # apply-groups from template
  fe-0/0/0 { mtu 5k; }
  fe-0/0/1 { ... }
```

```
fe-0/0/2 { ... }  
}
```

Configuring Clusters

Configuring clusters has many similarities to configuring standalone devices. You can use NSM to configure cluster objects either directly through regular configuration or through templates. Junos device clusters can also use configuration groups to define cluster data, just as for standalone Junos devices.

For Screen OS/IDP clusters, Secure Access clusters, and Infranet Controller clusters, most of the configuration information is the same among the various members. However, there are some differences that are configured separately on each cluster member. Junos clusters are different in that the configuration on each Junos cluster member is identical. These clusters use a special implementation of the configuration group mechanism to maintain differences between the members, but within the same configuration file.

Although you cannot edit the configuration of a cluster member, you can view its configuration. In the Info tab of the open cluster, select **Members**. Icons representing the members of the cluster appear in the main display area. Select the cluster member you want to view, and click the Edit icon to display the cluster member configuration.

After editing the cluster configuration, you push the edited configuration to the cluster using the Update Device directive. The cluster ensures that all members are synchronized.

Configuring Cluster Objects Directly by Editing the Configuration

For all device families, you can edit the cluster configuration by selecting the cluster icon in the Device Manager and clicking the Edit icon. You then apply edits to the cluster as you would to a standalone device. See [“Editing Devices Using the Device Editor” on page 202](#) for details about editing a configuration.

Configuring Cluster Objects Using Templates

To configure a cluster object using a template:

1. In the Device Manager, select **Devices**.
2. In the main display area, select the cluster you want to edit, and then click the Edit icon.
3. In the Cluster Info tab, click **Templates**, and then click the Edit icon. The Edit templates dialog box appears.
4. Select the templates you want to apply to the cluster, and then click **OK**.
5. The selected templates appear in the main display area. Arrange them in order of preference using the up and down arrows.

For more information about templates, see [“Using Device Templates” on page 210](#).

Configuring Global Cluster Data with Configuration Groups (Junos Clusters Only)

You can apply configuration groups to a Junos cluster object just as you can to a standalone Junos device. See [“Using Configuration Groups” on page 235](#).

You can cluster J Series routers or SRX Series gateways. You cannot cluster EX Series devices, M Series devices, or MX Series devices.

You can include configuration groups within templates when configuring cluster objects. Exactly the same rules apply as when configuring a standalone device. See [“Using Configuration Groups with Templates” on page 243](#) for details.

Configuring Member-Level Data in a Junos Cluster

To provide configuration data for a specific cluster member, such as the node name, NSM implements a special form of the wildcard mechanism to designate a configuration group to a specific cluster member. For ease of management, we recommend placing all your member-specific configuration data in one configuration group for each member. You can apply multiple configuration groups to each member.



NOTE: Imported configurations already have the member-specific configuration groups created and applied. Use the procedure described here only for modeled configurations.

We recommend using `node0` and `node1` as the names of the configuration groups that correspond to member 0 and member 1 of the cluster, although you can use any name containing the strings “node0” and “node1”. We recommend you do not use `node0` or `node1` as the names of configuration groups that contain cluster-level data.

To configure member-level data in a J Series cluster, follow these steps:

1. In the Device Manager, select **Devices**.
2. From the list of devices, select the cluster whose member you want to configure, and then click the Edit icon.
3. In the Configuration tab, select **Config Groups**.
4. Click the Add icon and select **Config Group for HA Node (node0|node1 etc)** from the list.
5. Configure the group as desired and click **OK**.
A configuration group called “node” appears in the Config Group List.
6. Right-click **Configuration** in the cluster member tree and select **Apply/Exclude Config Group** from the list.

The Apply/Exclude Config Groups dialog appears with the configuration group named node already highlighted in the Available Config Groups list.

7. Click the Add icon above the Applied Config Groups list (and not the Add button).

A dialog box appears and requests you to enter a string.

8. Type **apply \${node}** in the box, and then click **OK**.

The `${node}` is automatically expanded by NSM to create and apply configuration groups node0 and node1 to each member node.

9. Click **OK** to apply the configuration group.

If you later need to edit the local data for a cluster member, you do so by editing the configuration group for that member.

Configuring Junos Devices with Redundant Routing Engines

Configuring a device with dual Routing Engines differs from configuring a device with a single Routing Engine in that you can configure features for a specific Routing Engine. Two special configuration groups are used for this purpose:

- Configuration group re0 for the Routing Engine in slot 0
- Configuration group re1 for the Routing Engine in slot 1

Features configured in these special Routing Engine configuration groups appear only in the Routing Engine configuration to which they were applied. They do not appear in the global configuration, regardless of which Routing Engine is the master.

All other configuration groups applied to the device apply to the global configuration and not to individual Routing Engines.

Configuring a Routing Engine

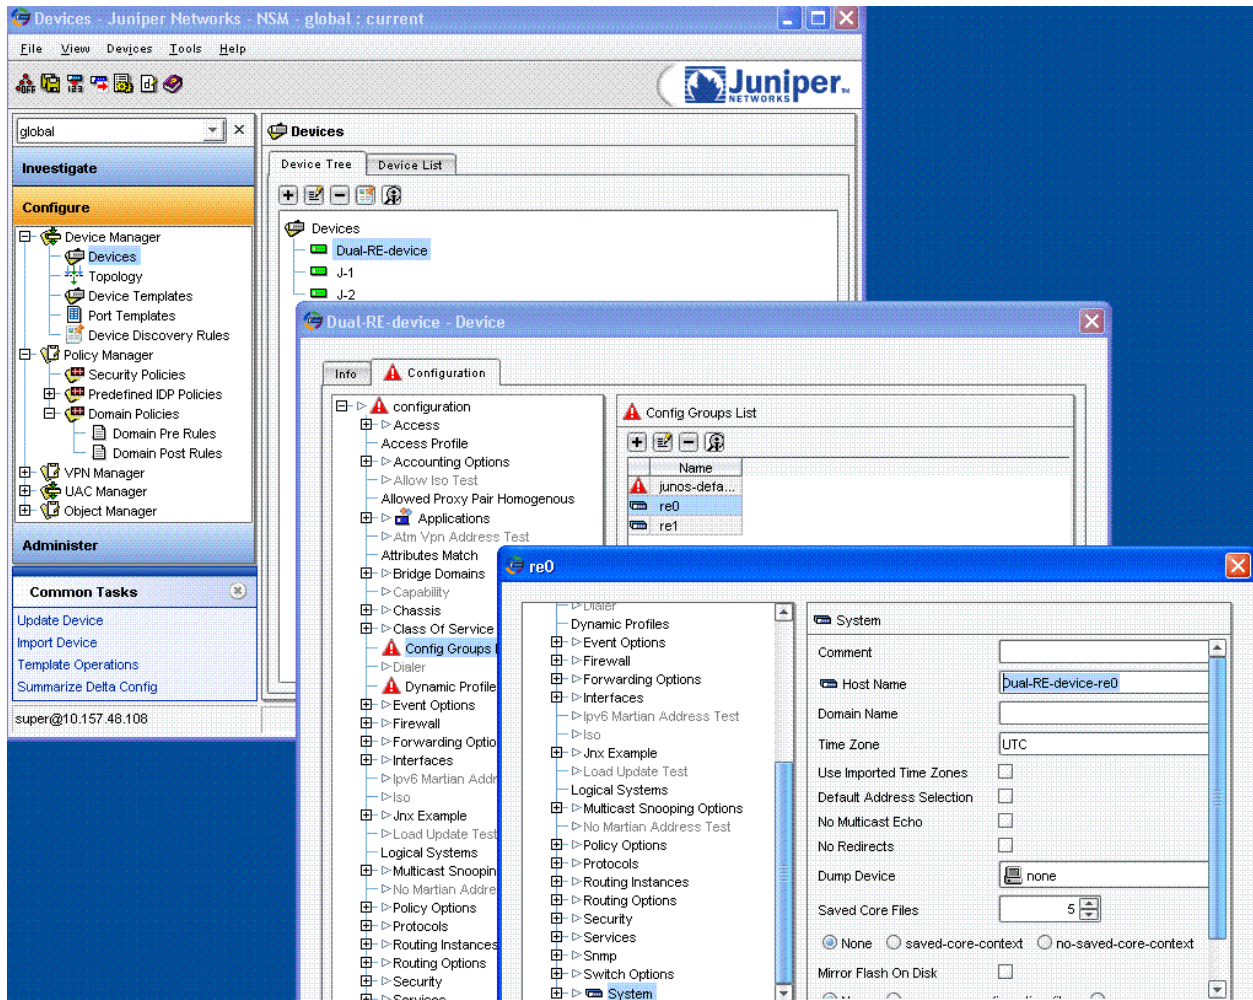
The following example configures a separate hostname for the Routing Engine in slot 0.

To configure a separate hostname for a Routing Engine in slot 0, see [Figure 66 on page 251](#) and follow these steps:

1. In the navigation tree, select **Device Manager > Devices**.
2. In the Device Tree, double-click the Junos router with redundant Routing Engines.
3. In the Configuration tab of the device editor, select **Config Groups List**.
4. If the config group re0 exists, open it by double-clicking its icon. If it does not already exist, click the Add icon, name the new configuration group re0, and then save it.

5. In the navigation tree for re0, select **System**.
6. In the Host Name field, assign a name to the Routing Engine, for example, **Dual-RE-re0**.
7. Click **OK** twice.

Figure 66: Configuring Routing Engine Specific Parameters



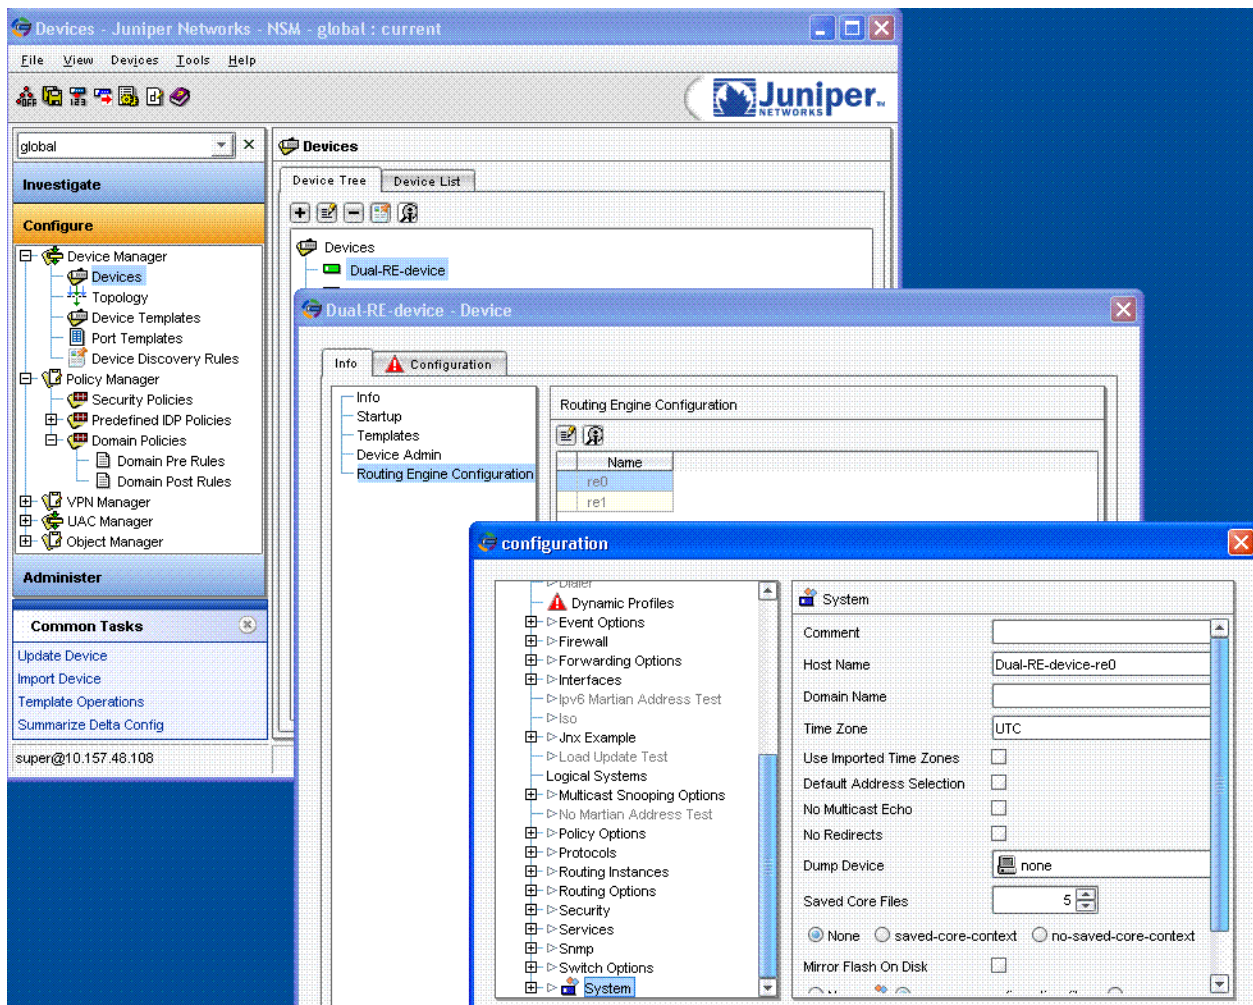
Viewing a Routing Engine Configuration

The following example shows how to display the hostname assigned to a specific Routing Engine. See [Figure 67 on page 252](#) and follow these steps:

1. In the navigation tree, select **Device Manager > Devices**.
2. In the Device Tree, double-click the Junos router with redundant Routing Engines.

3. In the Info tab of the device editor, select **Routing Engine Configuration**.
4. Double-click on the configuration group name to show the configuration for the corresponding Routing Engine.
5. In the navigation tree, select **System**. The configured Routing Engine name appears in the Host Name field.

Figure 67: Viewing the Routing Engine Configuration



Overview of VRRP Support in NSM

The VRRP feature allows you to use redundant routers on a LAN by configuring a single default route on the host. All VRRP routers share the IP address corresponding to the configured default route. One of the VRRP routers acts as the master and the others are backups. If the master fails, one of the backup routers becomes the new master, providing a virtual default router, and ensuring that traffic on the LAN is continuously routed.

The NSM implementation of VRRP has the following limitations:

- WAN and serial interfaces cannot support VRRP. VRRP requires Ethernet or gigabit interfaces.
- A single interface can support only two VRRP groups.
- You can enable a VRRP group only on a single interface.
- VRRP and NSRP are mutually exclusive on same device. Therefore, a device supporting VRRP cannot support NSRP and vice versa.
- All parameters available to VSI are applicable to the VRRP interface.
- You can only activate VRRP on either a regular interface or a sub-interface generated from one.
- You can configure VRRP parameters only on a VSI interface generated from either a regular interface or a sub-interface generated from one.

Platforms on Which NSM Supports VRRP

Release versions of ScreenOS 6.1 and NSM 2008.2 and later support VRRP fully.

NSM 2008.2 and later, support the VRRP feature in the following devices: SSG5, SSG20, SSG140, SSG320, SSG320M, SSG350, SSG350M, SSG520, SSG520M, SSG550, SSG550M.

NSM does not support VRRP in the following devices: ISG1000, ISG1000 with SM, ISG2000, ISG2000 with SM, NS5200M2/SPM2, NS5400M2/SPM2, Vsys devices.

Activating VRRP on a Device Interface

You can enable VRRP on an Ethernet Interface only if VRRP has already been activated on the device. You can only enable VRRP on a regular interface or a subinterface generated from one. You can also configure VRRP parameters only on a VSI interface generated from either a regular interface or a subinterface generated from a regular interface.

The VRRP interface follows the naming conventions of the NSRP VSI interface and is defined as “interface: group-id”. You must select a group-id between 1 and 7.

To enable VRRP, from the Physical Interface screen of the device, select the check boxes to activate VRRP, as well as to enable VRRP, and click OK.

Defining a VSI as a VRRP interface

A Virtual Switch Interface (VSI) can also function as a VRRP interface. NSM treats both VSI and VRRP interfaces similarly. You can create a VRRP through a VSI only if VRRP has already been activated on the interface. All the VSI parameters in the interface are editable for the VRRP. You can enter the IP mask for the VRRP. A new page is added under the VSI Protocol subtree for configuring VRRP parameters. NSRP cluster settings are disabled on the new VRRP page.

Managing Configuration Files

Configuration file management is available for devices running Junos OS and ScreenOS. You can access the configuration in its native text format, typically ASCII text, and process it in the NSM database, where you can view it, compare it with other configuration file versions, or use it to update a device.

You can fetch a configuration file from a device interactively from the NSM menus, or you can set up a cron job to import the configuration text file periodically and put these configuration file versions into the NSM database. Setting up cron job is not supported on devices running ScreenOS.

For a given device, you can see the different configuration file versions that are available, as well as compare different versions. Right-click the device in the Device Manager. select **Config File Management** from the list, and use the menu items. ScreenOS configuration file management is configuration driven and can be enabled using the option **Preferences System Properties Enable SOS CFM**.

Viewing and Comparing Configuration File Versions

If you choose **Config File Management > Show Config File Versions**, the Config File Versions dialog box appears. This dialog box shows all available versions of the configuration file saved in the NSM database and, for each version, shows the version number, the timestamp of when the version was imported and placed in the NSM database, and a comment.

Click on an entry in the table to view the contents of a specific version. The text file appears in the main part of the display. You can edit the comment that appears at the top of the display.

You can select any two versions listed in the Config File Versions dialog box and compare them. Click **Compare**. The UI highlights any differences between the two files to more easily allow you to navigate to the detected differences.

Updating the Device with a Configuration File Version

You can select a version listed in the Config File Versions dialog box and use it to update a specific device. Click **Update** and answer **Yes** to the confirmation prompts. NSM launches a job that updates the device with the selected configuration file.

Updating the Archived Configuration File Version

NSM service provides a **devSvrTFTP** option for updating the archived configuration files. NSM configuration file is hosted on a TFTP server. Because the device does not support any NSP messages for pushing the configuration file to device flash, the device gets the new configuration file from TFTP root directory and saves the new configuration to device flash. NSM will execute a reboot device command to load the new configuration on a device after the configuration file is successfully saved.

**NOTE:**

- Updating the files with an archived configuration file version is not supported on Vsy Devices.
- If preference is not enabled, the configuration file cannot be updated.

Importing or Viewing the Current Version of the Configuration File for SRX Series Devices

Select **Config File Management > Import Config File** to fetch the current version of the configuration file from the device. NSM launches a job to complete the task.

Select **Config File Management > View Running Config File** to fetch the current configuration and display it in the UI.

Select **Config File Management > Diff Running Config File** to compare the running configuration on the device with the latest available version in the database. The display highlights the differences.

Importing or Viewing the Current Version of the Configuration File for Devices Running ScreenOS

Select **ScreenOS Config File Management > Import Config File** to fetch the current version of the configuration file from the device. NSM launches a job to complete the task.

Select **ScreenOS Config File Management > Diff Running Config File** to compare the running configuration on the device with the latest available version in the database. The display highlights the differences.

Automatic Import of Configuration Files

When enabled, NSM by default automatically deletes oldest versions of config files to accommodate newer versions of the config files that are being imported. You can set a system preference for the maximum number of config file versions to be preserved. The default is 25 versions. The Config File Manager can automatically import config files from managed devices running Junos OS when configuration changes are committed on these devices, enabling NSM to have different versions of the device configuration. You can enable or disable the automatic import of config files and track those devices on which the feature is enabled. You can also see status of the config file versions.



NOTE: Automatic import of the configuration files is not supported on devices running ScreenOS.

CHAPTER 6

Updating Devices

This chapter explains how to update the running configuration (the configuration on the device) with the modeled configuration (the configuration in the Juniper Networks Network and Security Manager (NSM) UI). This chapter also describes the events that can require you to update your device, as well as NSM tools that help you to track, verify, and preview the update process.

After you model or make changes to a device configuration in the NSM UI, you must install that device configuration on the physical Juniper Networks device before those changes can take effect. NSM supports *atomic configuration*, a fail-safe feature that ensures successful updates occur without errors or the update is not performed. Atomic configuration is always enabled and occurs automatically when a device update causes the device to lose its connection to the management server.

This chapter contains the following sections:

- [About Updating on page 257](#)
- [Knowing When to Update on page 262](#)
- [Using Preview Tools on page 266](#)
- [Performing an Update on page 270](#)
- [Tracking Device Updates on page 273](#)

About Updating

When you update a managed device, you modify the running device configuration (the configuration currently installed on the physical device) with the modeled device configuration (the configuration currently modeled in NSM).

You can update a single device, multiple devices, vsys devices, clusters, virtual chassis or device groups simultaneously. For example, if you have created a device group that includes only NetScreen-5GT devices, you can update the entire device group in a single update procedure. During the update, NSM displays the progress of the update on each individual device so you can see exactly what is happening. Simultaneous updating also reduces downtime to unaffected devices and areas of your network.

Updating a device is a three-step process.

1. Ensure that you have configured the device correctly, created and assigned a policy to the device, and established a connection between the device and the management server.
2. From the Device Manager launchpad, select **Update Device**. The launchpad displays the Update Device(s) dialog box.

All connected and managed devices appear in the device list. Modeled devices and devices awaiting import for the first time do not appear.

3. Select the devices or device groups you want to update and click **Apply Changes**. NSM updates the selected devices or device groups with the modeled configuration.



NOTE: For IDP devices, the firmware version on the IDP device must match the version recorded in NSM, or the update will fail and the Config state of the device will change to "OS Version Adjustment needed". This situation can happen if the device firmware is upgraded from outside of NSM without importing the new configuration into NSM.

NSM uses centralized control and tracking to indicate when you need to update a device, and to follow the progress of the device configuration you are updating. Before updating your managed devices, you can use other NSM modules and tools to identify devices that need to be updated, validate their modeled configurations, and preview how those devices accept the new configuration. After updating, you can use the same tools to verify a successful update. These tools include:

- **Audit Log Viewer**—This NSM module records changes made to a device configuration. The audit log entry also identifies the administrator who performed the change, shows when the change was updated on the device, and provides a history of change details.
- **Report Manager**—This NSM module collects data from traffic logs on various events that occur over your network and provides a visual representation of them. You can customize reports to display and filter parameters.
- **Configuration Summaries**—These tools provide a preview of the modeled configuration, enabling you to compare it with the configuration that is running on the device. Use configuration summaries to ensure the modeled configuration is consistent with what you want to update on the device.
- **Job Manager**—This NSM module tracks the status of running and completed update processes. The Job Manager displays details of the update process in a dedicated information window and includes the update's success or failure and errors involved in a failed update.

How the Update Process Works

The managed device is functioning normally. You have successfully added the device to NSM, reviewed the device configuration, and updated the device. An event occurs on the managed device that requires a change to the device configuration. For example, malicious traffic might have entered your network, requiring you to update the security policy for the device to detect and prevent that attack.

1. Using the NSM monitoring tools, you learn of the attack and locate the cause of the event. Using NSM modules such as the Realtime Monitor and Log Viewer, you determine the exact attack that penetrated the device. From the Report Manager, you also determine what rule in the security policy was ineffective in blocking the attack.
2. You update the modeled device configuration, editing the security policy to detect and prevent the attack from entering your network again.
3. Before updating the running configuration, you review the modeled device configuration. Using a delta configuration summary, compare the modeled configuration with the running configuration on the device to confirm the differences. Fine-tune the modeled configuration, if needed.
4. When you are confident that the modeled configuration is valid, update the device. NSM updates the running configuration with only the new changes (delta). During the update, you track the update progress using Job Manager in real time and observe the transfer of the configuration from NSM to the device.

If the update is unsuccessful, use the information in the Job information window to correct the problems in the modeled configuration.

5. After updating, run a second Delta Configuration Summary to identify any remaining differences between the modeled configuration and the running configuration on the device. When the Delta Configuration Summary reveals no differences between the new configuration and the old configuration on the device, you have successfully updated the running configuration.

About Atomic Configuration—ScreenOS Devices

NSM uses atomic configuration, a fail-safe feature for updating devices. Atomic configuration ensures that a current valid configuration is not overwritten by a flawed configuration in flash memory. The update must finish without errors and the device connection to the management system must remain active, or the update is aborted to prevent an invalid, error-prone, or flawed configuration from being installed on the device.

Atomic configuration is always on. During an update:

1. NSM saves and locks the active configuration on the device, and then starts a timer for the update process. While the active configuration is locked, it cannot be changed.
2. NSM sends the modeled configuration to the device.
3. As the device receives the modeled configuration, it updates its existing active configuration with each command as the command is received:
 - If the device executes the entire modeled configuration (all commands) and the connection to the management system remains up, NSM unlocks the active configuration and saves the new active configuration.
 - If the device cannot execute a command, NSM resets the device, unlocks the active configuration, and restores the saved active configuration to the device (the device reboots). After rebooting, the device sends a final error message to the management

system; the contents of this message, which include any CLI errors in the failed configuration, appear in the Job Manager status window for this upgrade.

- If the device connection to the management system is down after all commands have been executed, the update timer expires and the device automatically resets. The device unlocks the active configuration and restores the saved active configuration (the device reboots). The connection might be down due to a command in the modeled configuration that causes the device to lose connection with the NSM Device Server.



NOTE: When updating vsys devices, atomic configuration occurs only for the root vsys.

About Atomic Updating—ScreenOS Devices

In addition to atomic configuration, devices running ScreenOS 5.1 and later also support *atomic updating*, which enables the device to receive the entire modeled configuration (all commands) before executing those commands (instead of executing commands as they are received from the management system). Because NSM sends all commands at one time, the performance of the management connection is enhanced.

Atomic updating also enables the device to temporarily lose connection to NSM during the update process. If the management connection is down when the device has finished executing the commands in the modeled configuration, the device reestablishes the connection. Because the device no longer needs to maintain a constant connection to the management system during updating, you can configure changes to the management connection from the NSM UI.

During an atomic update:

1. NSM saves and locks the active configuration on the device, and then starts a timer for the update process. The timer expires after 40 minutes. While the active configuration is locked, it cannot be changed.
2. NSM sends the modeled configuration to the device.
3. The device receives all commands before executing the commands on the active configuration. During the update, the device sends progress messages to the management system every 15 seconds; these messages appear in the Job Manager status window for the update.

During the update, the Job Manager status window displays other messages, depending on the success of the update:

- **Updates Without Errors**—If the device executes the entire modeled configuration (all commands) and the connection to the management system remains up or can be reestablished, NSM unlocks the active configuration and saves the new active configuration. The device sends a final message to the management system; this message appears in the Job Manager status window for this update.
- **Updates With Errors**—If the device cannot execute a command, it notifies the management system, which makes a decision whether to ignore and proceed, abort, or revert.

For ignore and proceed decisions, the device continues the update.

For abort and revert decisions, the device automatically resets. The device unlocks the active configuration and restores the saved active configuration (the device reboots). After rebooting, the device sends a final error message to the management system; this message, which includes any CLI errors in the failed configuration, appears in the Job Manager status window for this update.

- **Re-establish Management Connection**—If the device connection to the management system is down after all commands have been executed, the device attempts to reestablish connectivity.

If successful, NSM unlocks the active configuration and restores the saved active configuration to the device. The device sends a final message to the management system; this message appears in the Job Manager status window for this update.

If attempts to reconnect are unsuccessful for two hours, the update timer expires and the device automatically resets. The device unlocks the active configuration and restores the saved active configuration (the device reboots). After rebooting and reestablishing the connection to the management system, the device sends a final error message to the management system; this message, which includes any CLI errors in the failed configuration, appears in the Job Manager status window for this update.

About Atomic Configuration and Atomic Update—DMI-Compatible Devices

The device management interface (DMI) provides built-in support for atomic configuration and atomic update for Junos devices, Secure Access devices, and Infranet Controller devices. The update to the device does not get committed until the operation is complete, so a failure during an Update Device operation can never leave an inconsistent configuration on the device.

- If the Update Device directive finishes and the connection between the device and NSM remains up, the device transitions to use the new configuration. The Update Device operation is successful.
- If the connection between the device and NSM is lost during the Update Device operation, the job status will report failure, and the device will rollback to the original configuration. Unlike ScreenOS devices, however, DMI-compatible devices do not need to reboot in order to rollback.
- If the connection between the device and NSM remains up throughout the Update Device operation, but the update itself fails, the DMI device will keep the original

configuration, because update to the device does not get committed until the operation finishes.

About Implicit Updates (Secure Access and Infranet Controller Devices Only)

Secure Access and Infranet Controller configuration data is structured such that the creation or change of some configuration data can cause implicit change in other configuration data. For example:

- Resource policies can be automatically created as a result of Resource Profile configuration.
- Encrypted passwords are created on a Secure Access device as a result of cleartext password configuration.
- Bookmarks created in the Web Access Policies list (resource policies or allow rules for certain Web sites) are created when the "auto-allow" checkbox is checked when a Role bookmark is created.
- Configuration of a HostChecker role restriction triggers automatic configuration of a realm-level host checker evaluation setting.

As a result, an Update Device directive to a Secure Access or Infranet Controller device can result in configuration data changing on the device which was not set in NSM. To synchronize the configuration data, NSM imports the configuration after the update.

If an Update Device directive causes implicit configuration changes on one or more devices, each device reports the event to NSM in the update device response. On receipt of this message, NSM performs the following actions:

- Shows the Update Device job as "Done" in the Job window.
- Changes the configuration state on devices with implicit changes to "Managed, Device Changed".
- Displays the Device Import Options popup that lists the devices with implicit configuration changes and informs you that the configuration is being imported.

Click **OK** to close the dialog box.

- Starts a job to import the configuration from each affected device.

Knowing When to Update

Typically, you update a device after changing the device configuration or after modifying the security policy that is assigned to the device.

- To overwrite the existing configuration on the physical device, update the physical device with the modeled configuration in NSM.
- To overwrite the modeled configuration in NSM, import the existing configuration from the physical device. NSM does not support delta updates from the device.

Using NSM, you can identify the changes made to the device or to the modeled configuration, and then update the device. For significant changes to the network that the security device is deployed in, you might also need to change the assigned policy.

The following sections explain how to detect configuration or policy changes:

- [Verifying Device Status in Device Monitor on page 263](#)
- [Verifying Device Status in Device Manager on page 265](#)
- [Reviewing Logs on page 265](#)
- [Identifying Administrative Changes on page 266](#)
- [Reviewing Reports on page 266](#)

Verifying Device Status in Device Monitor

Within the management system, a managed device has an associated connection status and configuration status. NSM displays each status for each managed device in RealTime Monitor > Device Monitor.

For more details on using the Device Monitor, see [“Monitoring Managed Devices” on page 706](#).

Connection Status

The connection status indicates the status of the connection between the managed device and the Device Server. NSM uses heartbeat packets to continually test the connection between the Device Server and the physical device. The connection status column in the Device Monitor displays the current status of the device:

- Up status—The device is connected to the Device Server and is running properly. Before you can update a device, it must be in the Up state.
- Down status—An event has occurred, either manually by an administrator or automatically by the flow of a type of traffic, that has stopped the device from running.
- Never Connected status—The device has not made an initial connection to Device Server. Typically, this state appears for modeled devices that have not been activated, or for devices waiting to be activated using Rapid Deployment.

Configuration Status

The configuration status indicates the status of the device configuration on the physical device. Some common configuration states include:

- Managed—The running configuration is the same as the modeled configuration (the device is using a “managed” configuration).
- Modeled—The running configuration is not the same as the modeled configuration, and the device has not yet connected to NSM.
- Import Needed—The running configuration is not the same as the modeled configuration, but the device has connected to NSM and is awaiting manual import.

- **Update Needed**—Indicates that the running configuration is not the same as the modeled configuration, and the device is connected to NSM. You must update the managed device before the changes you made in the modeled configuration can take effect.
- **Platform Mismatch**—The device platform selected when adding a DMI device in NSM does not match the device itself.
- **Device Firmware Mismatch**—The OS version selected when adding a DMI-compatible device does not match the OS version running on the device itself.
- **Device changed**—For Junos devices with dual Routing Engines, indicates that a switchover has occurred, because the configuration commit timestamps on the master and backup Routing Engines are not synchronized.

For devices running ScreenOS 5.1 and later, NSM supports additional configuration states that indicate the status of the physical device configuration in relation to the modeled configuration in NSM. In addition to the states listed above, a device running ScreenOS 5.1 and later can have one of the configuration states shown in [Table 27 on page 264](#).

Table 27: Additional Configuration States for Devices Running ScreenOS 5.1 and Later

Detail State	Details
Managed, In Sync	The physical device configuration is synchronized with the modeled configuration in NSM.
Managed, Device Changed	<p>The physical device configuration is not synchronized with the modeled configuration in NSM. Changes were made to the physical device configuration (the configuration on the physical device is newer than the modeled configuration).</p> <p>To synchronize the two configurations, import the configuration from the physical device.</p>
Managed, NSM Changed	<p>The modeled device configuration is not synchronized with the physical device configuration. Changes were made to the modeled configuration (the configuration on the NSM is newer than the physical device configuration).</p> <p>Any change made in the UI automatically causes the NSM configuration state to change, even when the change is canceled or undone. For example, if you change a value in the UI to a different value, and then undo the change by entering the original value, the NSM configuration state is still considered not synchronized with the physical device.</p> <p>To synchronize the two configurations, update the configuration on the physical device.</p>
Managed, NSM and Device Changed	<p>Both device configurations (physical and modeled) are not synchronized with each other. Changes were made to the physical device configuration and to the modeled configuration.</p> <p>Although you cannot synchronize delta changes, you can run a delta configuration summary (see “Using a Delta Configuration Summary” on page 267) to identify the differences, then manually make the changes to the modeled configuration, and then update the device.</p>

Table 27: Additional Configuration States for Devices Running ScreenOS 5.1 and Later (continued)

Detail State	Details
Managed, Sync Pending	<p>Completion of an Update Device directive is suspended, waiting for the device to reconnect.</p> <p>This state occurs only for ScreenOS devices that have the Update When Device Connects option selected during device update.</p>

Changing the name, color, or NSM port on a device causes the configuration state to be out of sync, even though the management system and device do not share these parameters (these parameters are not transmitted to or from the device during an update).

For details on all states, see “[Viewing Device Status](#)” on page 706.

Verifying Device Status in Device Manager

You can view the connection and configuration status for each managed device in Device Manager.

NSM automatically updates the device status and displays the state of each device in the UI. To view device status, place your mouse cursor over the device name. A tooltip shows the device name, device type and OS version, IP address, domain, the Attack Db version if it is a Firewall\IDP device, and the connection and configuration states.

To manually verify the configuration status for devices :

- For a single device—Right-click the device and select **Check Config Sync Status**.
- For multiple devices—From the Device Manager launch pad, select **Device config Options > Check Config Sync Status**. Select the devices for which you want to view configuration status, then click **OK**.



NOTE: You can use the directive Check Config Sync Status from any location in the NSM UI. (You do not need to select the Device Manager.)

Reviewing Logs

The Log Viewer can help you identify event patterns on your network. To see patterns in the Log Viewer, create a custom view using filters that display log entries based on specific criteria. To set a column or cell filter, right-click the column or cell that you want to use as the matching criteria and specify the value.

For example:

- To track all events for a specific time period, create a filter on the timestamp column; when applied, the filter displays only the log entries that meet the specified time period.

- To track all events from a specific source IP address, create a filter on source address column; when applied, the filter displays only the log entries that use the specified source address.
- To track all events for a specific category or subcategory of log entries, such as Configuration or Attack log entries, create a filter on the category or subcategory column; when applied, the filter displays only the log entries with the specified category or subcategory designation.

For more details on using the Log Viewer, see [“Logging” on page 783](#). For step-by-step instructions on creating a filter, see the *Network and Security Manager Online Help* topic “Setting Filters”.

Identifying Administrative Changes

Use the Audit Log Viewer to identify administrative changes made to your managed devices. Audit log entries also identify the administrator who made the change, the action performed, and the date and time of the change. You can track changes by time of logging, administrator name, action, targets, and devices. If an administrator made a change to a device or an object, you might want to update the affected devices.

For details on using the Audit Log Viewer, see [“Using the Audit Log Viewer” on page 832](#).

Reviewing Reports

Use Report Manager to determine when you are receiving too many attacks of a certain type and order them by an IP address. For example, if you determine that the current device configuration and security policy cannot block scans, you might want to create a new rule in the security policy that guards against those attacks, and then update the device.

Report Manager provides three types of reports: time-based reports, event-based reports, and severity-based reports.

- To identify common events, select an event-based report to see the frequency of events in a bar graph or pie chart. To see details for a specific event, right-click the event and select **View in Log Viewer** to display a custom view in a new window. You can save a detailed view as a custom report. For example, when viewing the Top Alarms, expand a location to view the data that makes up this location.
- To examine how specific rules in your security policy are performing, select the **Administrative > Top Rules** report. You might need to fine-tune an inefficient rule to better handle events in your network traffic.

For details on using Report Manager, see [“Reporting” on page 855](#).

Using Preview Tools

When you update a managed device, you overwrite the existing configuration that is running on the physical device. Therefore, it is important to verify a configuration before sending it to the device.

Using preview tools, you can preview how the modeled configuration looks in CLI command form or XML message form to predict the success of the update and anticipate errors. NSM supports these types of preview tools:

- **Configuration Summary**—Displays the modeled configuration using ScreenOS CLI commands, or XML for DMI-compatible devices.
- **Delta Configuration Summary**—Displays the modeled configuration and running configuration in CLI command form or XML script form, and lists the differences between the two configurations.
- **Running Configuration**—Displays the configuration installed on the physical device.

The configuration and delta configuration summaries help you ensure that the modeled configuration is correct before you update your managed devices, while the running configuration helps you identify settings already on the managed device.

Running a Configuration Summary

When you update a managed device using NSM, the management system generates CLI commands or XML script that map to the settings in the NSM UI. To verify that the configuration you are installing on the device generates the correct commands, run a configuration summary.

1. From the launchpad, select **Devices Config Options > Summarize Config**. The launchpad displays the Summarize Config dialog box.
2. Select the devices or device groups for which you want to run a configuration summary and click **OK**. A Job Information window displays the progress of the summary.
3. When the job completes, review the CLI commands or XML script in the Job Information window. When you update the device, these are the commands that NSM uses to overwrite the running configuration.

For some settings, the CLI commands for a UI settings do not map one-to-one. For example, a single vsys configuration in the NSM UI generates multiple ScreenOS commands.

Because the management system generates all information (UI settings, XML, and CLI commands) for a configuration summary, you can run a configuration summary on a modeled device, even if no corresponding physical device is connected.

Using a Delta Configuration Summary

A delta configuration summary compares the active configuration on the ScreenOS or DMI-compatible device with the modeled configuration in NSM and displays the differences between the two configurations. The delta configuration summary produces four sets of data. See [Table 28 on page 268](#).

Table 28: Delta Configuration Summary Information

Delta Config Data	Description
Config on Device but not on NSM	Displays (in CLI command form or XML script form) the commands detected on the device that do not map to NSM settings. Use this information to identify any out-of-band updates (made by the local device administrator) to the running configuration; you might not want to overwrite these settings.
Config on NSM but not on Device	Displays (in CLI command form or XML script form) the commands (as mapped to NSM settings in the modeled configuration) detected in NSM but not on the device. Use this information to identify the changes you have made to the modeled configuration since the last update.
Config on both NSM and Device but reordered	Displays (in CLI command form or XML script form) the commands for configuration settings present on both the device and NSM, for which the CLI command sequence has been reordered.
Config to be sent to device on next Update Device	Displays (in CLI command form or XML script form) the commands that NSM will send to the device on the next update.

You should run a delta configuration summary twice:

- Before updating—Because you are overwriting the running configuration with the modeled configuration, you might want to identify and verify the configuration you are installing on the device.
- After updating—Ensure that the device received the configuration as you expected, and that no differences exist between the running configuration and the modeled configuration.

Delta configuration summaries are helpful tools for ongoing device maintenance, particularly for devices that are managed both locally by a device administrator using CLI commands or the Web UI and remotely by a NSM administrator using the NSM UI. Because the modeled configuration can overwrite the running configuration, you should always confirm the commands that are sent to the device.

To run a delta configuration summary:

1. From the Device Manager, select **Summarize Delta Config**. The launchpad displays the Summarize Delta Configuration dialog box.
2. Select the devices or device groups for which you want to run a delta configuration summary and click **Apply Changes**. A Job Information window displays the progress of the summary.
3. When the job completes, review the CLI commands or XML script in the Job Information window. Specifically, review the commands in the section “Config to be sent to device on next Update Device”; when you update the device, these are the commands that NSM uses to overwrite the running configuration.

A sample delta configuration summary for a ScreenOS device is shown in [Figure 68 on page 269](#).

Figure 68: Delta Configuration Summary Example

<p>Config on Device but not on NSM:</p> <pre>unset interface serial manage telnet unset interface serial manage web unset interface serial manage ssl unset interface serial manage snmp unset interface serial manage scs unset interface serial manage ping unset av http trickling set av all max-connections 0 set zone untrust screen syn-flood set zone untrust screen syn-flood alarm-threshold 512 set zone untrust screen syn-flood attack-threshold 200 set zone untrust screen syn-flood queue-size 1024 set zone untrust screen syn-flood timeout 20 set zone untrust screen syn-flood source-threshold 512 set zone untrust screen syn-flood destination-threshold 1024 set policy id 1 from trust to untrust Any Any ANY permit set nsrp ha-link probe threshold 5 set nsrp ha-link probe interval 1 set nsrp vsd-group init-hold 5</pre>	<p>Commands for objects already configured on the device.</p>
<p>Config on NSM but not on Device:</p> <pre>set pppoe name untrust set zone untrust screen syn-flood set zone untrust screen syn-flood alarm-threshold 512 set zone untrust screen syn-flood attack-threshold 200 set zone untrust screen syn-flood queue-size 1000 set zone untrust screen syn-flood timeout 20 set zone untrust screen syn-flood source-threshold 512 set zone untrust screen syn-flood destination-threshold 1000 set policy id 700029 from trust to untrust Any Any ANY permit</pre>	<p>Commands for objects configured in NSM.</p>
<p>Config on both Device and NSM but reordered:</p>	<p>Commands on the device and NSM that have been reordered.</p>
<p>Config on NSM but not on Device:</p> <pre>set pppoe name untrust set zone untrust screen syn-flood set zone untrust screen syn-flood alarm-threshold 512 set zone untrust screen syn-flood attack-threshold 200 set zone untrust screen syn-flood queue-size 1000 set zone untrust screen syn-flood timeout 20 set zone untrust screen syn-flood source-threshold 512 set zone untrust screen syn-flood destination-threshold 1000 set policy id 700029 from trust to untrust Any Any ANY permit</pre>	<p>Commands to send during the next update. These commands overwrite the active configuration.</p>

Occasionally, the delta configuration report might display discrepancies that do not actually exist between the running configuration and the modeled configuration. In some specific situations, the running configuration includes CLI commands that do not appear as pending changes in NSM, yet the two configurations are actually in sync (no delta exists). This situation can occur when:

- Some settings for a feature have been configured in NSM, but the feature itself is not enabled. For example, if you configure NSRP settings but do not deploy the device in

NSRP mode, the CLI commands for NSRP settings appear in the running configuration but are not managed by NSM (because the feature is not active).

- DHCP settings, such as interface IP addresses, are not assigned by NSM, and are not included in the CLI commands sent to the device. The CLI commands do appear, however, in the running configuration.
- Default, unconfigured settings might not be managed by NSM. For example, if the running configuration includes the domain name mycompany.net, but that domain name is not configured in NSM, the management system leaves the value unchanged.

Performing an Update

You can update a single device, multiple devices, vsys devices, clusters, virtual chassis or device groups using the same process.

Before updating:

- Ensure that you have configured the device correctly, created and assigned a policy to the device, and established a connection between the device and the management server.
- Run a configuration summary on the device to view the CLI commands for the modeled configuration. Review these commands to ensure that you have configured the device as desired.
- Run a delta configuration summary to view the differences between the modeled configuration and the running configuration in CLI command or XML script format.

To update the device:

1. From the Device Manager launchpad, select **Update Device**. The launchpad displays the Update Device(s) dialog box.
2. Select the devices or device groups you want to update.
3. In the lower portion of the dialog box, check the Run Summarize Delta Config box if desired, and then click **Apply Changes**.

NSM begins updating the selected devices or device groups with the modeled configuration.

After updating:

- Review the information in the Job Information window to determine if the update was successful.
- If you chose to run a delta configuration summary, review the summary to ensure that no conflict exists between the running configuration and the modeled configuration.

Retrying a Failed Update

When updating your managed security devices, the update fails for each device that is not connected to the management system at the time of update. For devices running ScreenOS 5.1 or later, you can configure NSM to save the pending changes for an unconnected device, and then install those changes when the device finally connects to the management system.

NSM automatically changes the configuration state of an unconnected device that is waiting for changes to the “Sync Pending” state. When a device in this state connects to the management system, pending changes are immediately installed on the device and the configuration state is changed to “In-Sync”.

You can also configure the management system to abort update attempts for previously unconnected devices to which out-of-band changes have been made. For example, you attempt to update all your managed NS-5GT security devices, but device NS-5GT-25 is disconnected from the management system for troubleshooting at the time of update. When troubleshooting is finished and the device reconnects, to prevent NSM from overwriting any out-of-band changes made, enable the option “Do not Update If Device Has Changed”.

Configuring Update Options

You can configure device update and retry options on a systemwide basis (in the UI preferences), on a per-update basis for multiple devices (in the Update Device(s) dialog box), and on a per-update basis for a single device (in the device options dialog box). The systemwide settings appear as the default settings for the per-update settings, which you can change as needed for each update.

When configuring systemwide update options, you can enable or disable any option independently; when configuring per-update options, dependencies apply.

Update Options for ScreenOS and IDP

Unconnected Device options include:

- **Show Unconnected Devices in Device Selection Dialog**—When this option is enabled, the NSM UI displays devices that are not connected to the management system in the Update Devices dialog box (which appears when you attempt to update the configuration for a managed device).

When this option is disabled, unconnected devices do not appear in the Update Devices dialog box, preventing administrators from selecting an unconnected device for updating.

When you configure this option on a per-update basis, you must enable it in order to make the “Update When Device Connects” option available.

- **Update When Device Connects**—When this option is enabled, NSM attempts to update a previously unconnected device that has pending changes stored in the management system.

When this option is disabled, NSM does not update a previously unconnected device, and the configuration state of the device remains as “Sync Pending”.

When you configure this option on a per-update basis, you must enable it in order to make the “Do not Update If Device Has Changed” option available.

- **Do not Update If Device Has Changed**—When this option is enabled, NSM does not update a previously unconnected device if out-of-band changes have been made to the device. The configuration state of the device remains either **NSM and Device Changed**, or **Device Changed** when the update device job is canceled as a result of a change on the device.

Firewall Device options include:

- **Rematch, Session Treatment when modifying a policy rule**—When this option is enabled, NSM preserves the existing sessions that are being tracked by the installed security policy during the policy update procedure (devices running ScreenOS 5.1 or later only). At the end of the update, NSM restores all valid sessions on the managed device and deletes all invalid sessions.

When this option is disabled, NSM does not preserve and restore existing sessions for an updated managed device.

Standalone IDP Device options include:

- **Update IDP Rulebase Only**—When this option is enabled, only the IDP rulebase is updated. When changes are made in the summarize delta configuration, only the firewall policies are displayed. This option is applicable only for ISG and SRX Series devices.
- **Update Firewall Configuration Only**—When this option is enabled, firewall-specific configurations are pushed to the device. Pushing large IDP configurations along with firewall configurations takes a very long time. For this reason, all configurations other than IDP and IDP policies are pushed to the device. This option reduces the time required for updating devices with large configurations and improves performance.



NOTE: Deleting an address-book or security-zone in NSM requires full device update. A firewall-only update does not delete the address book or the security zone. This is applicable only on SRX Series devices.

For details on tracking update status, see the next section, [“Tracking Device Updates” on page 273](#). For details on troubleshooting failed updates, see [“Understanding Updating Errors” on page 276](#).

Update Options for DMI-Compatible Devices

For DMI-compatible devices, update options include:

- Lock configuration during update.
- Update to candidate config first before commit to running config.
- Use confirmed commit.

- Rollback candidate config to running config in error.
- Discard uncommitted changes when exclusive lock is available.

Tracking Device Updates

Use Job Manager to track device updates in real time. You can view the status of a running update and the status of completed updates in the Job Manager module.

When you send a command to a device or group of devices using NSM, the management system creates a *job* for that command and displays information about that job in the Job Information window. The command you send is called a *directive*. Job Manager tracks the progress of the directive as it travels to the device and back to the management system. Each job contains:

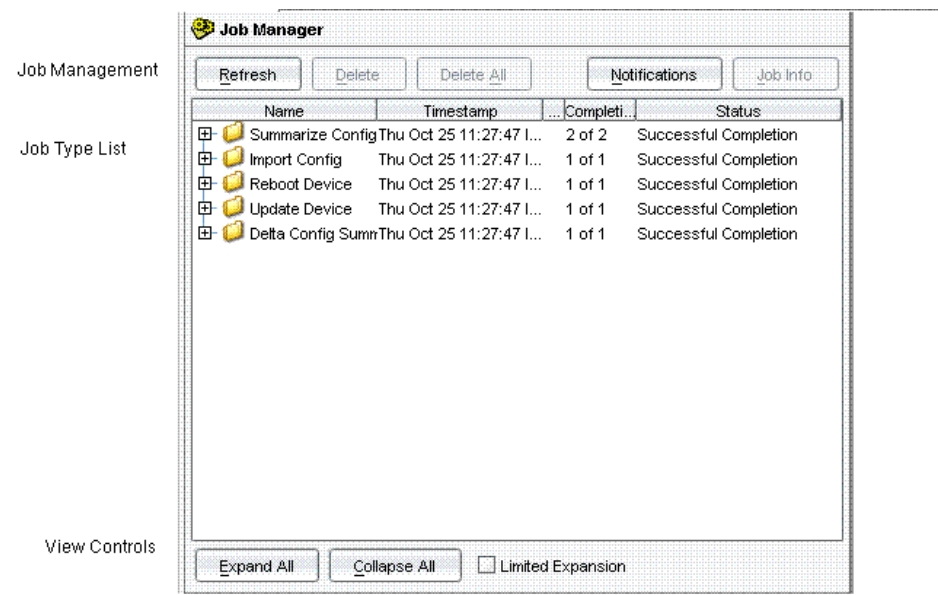
- Name of the command
- Date and time the command was sent
- Completion status for each device that received the command
- Detailed description of command progress
- Command output, such as a configuration list or CLI changes on the device



NOTE: Job Manager configuration summaries and job information details do not display passwords in the list of CLI commands or XML script for administrators that do not have the assigned activity “View Device Passwords”. By default, only the super administrator has this assigned activity.

You can initiate directives from multiple locations in the NSM UI, including the Devices and Tools menus in the NSM toolbar (to access the Update directive, from the File menu, select **Devices > Configuration > Update Device Configuration**). The Job Manager module is shown in [Figure 69 on page 274](#).

Figure 69: Job Manager Module




Job Manager includes the following utilities and information:

- **View Controls**—Use View controls to set the information level you want displayed in Job Manager:
 - *Expand All* displays all devices associated with a directive type.
 - *Collapse All* displays the directive type.
- **Job Type (Directive) List**—Displays the job type (directives) and associated timestamp completion status information. All current and completed jobs appear, including device updates. However, if you have not yet performed an update using NSM, the Job List does not display an Update Configuration directive.
- **Notification Controls**—Enables you to manually view job completion status.
- **Job Information**—Enables you to view job information, including errors, job completion status, job state, automatic job completion notification setting, and start time of job.

Reviewing Job Information

The Job Information dialog box displays the changing device states as the directive is executed. Device state changes, error messages, and warning messages are displayed in real time. A sample Job Information dialog box is shown in [Figure 70 on page 275](#).

Figure 70: Job Information Dialog Box



Job Manager tracks the overall progress of one or more jobs executed on a single device. For multiple device updates, Job Manager tracks the progress of each job on each device in addition to the overall progress for all devices. To view the Job status for an individual device (including error messages and percent complete), select the device in the Percent Complete pane; the status appears in the Output pane.

The Job Information includes:

- **Job Type**—The type of task being tracked. Job Types include Update Device, Reboot Device, and Config Summary. Job Type is also known as a directive.
- **Time Stamp**—The time at which NSM began executing the directive.
- **Job Status**—The current state of the job.
- **Number of Jobs Completed**—The number of jobs completed out of the total number of jobs.
- **Percent Complete**—The percentage of total jobs successfully executed. When performing multiple jobs on multiple devices, this field displays the percentage complete for each device. When the job has completed, successfully or unsuccessfully, this field displays 100 percent.
- **Device Name**—The name of the device on which the job is executed.
- **State Description**—The current state of the job.
- **Completion Level**—The percentage of a job that has executed successfully.
- **Output**—Displays the content of the update, including commands that have been interpreted from the NSM data model into device-specific commands, error messages, and existing commands deleted from the device. The Output Display Region displays

all errors, warnings, device verification output, and device state information associated with the job.



NOTE: If the Job Information dialog box might contain Chinese, Japanese, or Korean characters, you must uncheck the Fixed Font box to display them.



NOTE: Job Manager configuration summaries and job information details do not display passwords in the list of CLI commands for administrators that do not have the assigned activity “View Device Passwords”. By default, only the super administrator has this assigned activity.

Device States During Update

During an update, the managed device changes device state. You can view the current device state in real time in the State Description field of the Job Information dialog box. [Table 29 on page 276](#) lists the states that a device can have.

Table 29: Device States During Update

Device State	Description
None	No update activity has occurred on the device.
Loading in Progress	NSM is sending the update image to the flash memory of the device.
Pending	Device is accepting the parameters from the update configuration that has been sent to the device flash memory.
Converting Data Model to Device Data Model	The parameters that have been set in the NSM configuration are being changed to corresponding device-specific CLI commands that execute on the device.
Completion	Device has successfully been updated with the modeled configuration.
Failed	Device has not been successfully updated with the modeled configuration. The Output pane of the Job Manager dialog box displays error messages and error codes.

Understanding Updating Errors

When an update fails for any reason, Job Manager displays error codes and error messages that can help you identify and locate the problem. Typical errors include:

- The modeled configuration contained invalid values that the device could not process.
- During the update process, the connection between the managed device and the Device Server was lost.

- The modeled configuration caused the managed device to lose its connection to NSM.

For these update errors, the Job Information dialog box displays the Job Status as “Failed.”

You can also check the Connection Status and Configuration Status columns for the device in the Realtime Monitor to determine whether the device is running.

After a device is updated, you can run a delta configuration summary to determine any remaining differences between the modeled configuration and the running configuration; the output of this summary appears in the Job Manager information window. For successful updates, no discrepancies are found or displayed. For failed updates, the output area lists remaining discrepancies.

For example, a failed update job is shown in [Figure 71 on page 277](#).

Figure 71: Failed Update Job Dialog Box

```

Name:      Sgt
Description: failure
Completion: 100%

Error Code:
Error Text: Exception caught during update device: Verification failed. Configuration tree is not a sub-tree
Error Details:
  Detail of the diff:
  pppoe$name$untrust$clear-on-disconnect: .true. -> empty

logs:
  Getting config tree from device ...
  Generating removing CLI commands ...
  Generated 1 removing CLI commands.
  1 unset policy 1
  Sending removing CLI commands to device ...
  Getting config tree from device (second time)...
  Generating configuration CLI commands ...
  Generated 5 delta config CLI commands.
  1 set pppoe name untrust clear-on-disconnect
  2 set policy id 700029 from trust to untrust Any Any ANY permit
  3 set policy id 700029
  4 exit
  5 save
  Sending configuration cli commands to device ...
  Device Warning on command:
  2 set policy id 700029 from trust to untrust Any Any ANY permit
  policy id = 700029
  Verifying configuration ...
  Verification failed. Configuration tree is not a sub-tree of device tree.
  Detail of the diff:
  pppoe$name$untrust$clear-on-disconnect: .true. -> empty

```

In the Output area of this job, the update:

- Successfully removed existing commands on the device (**Generating Removing CLI Commands**)
- Unsuccessfully added new commands that were not present in the running configuration (**Generated 5 Delta Config CLI Commands**). Specifically, the update could not set the command: **pppoe name untrust clear-on-disconnect**

The delta configuration summary correctly detected a difference between settings on the managed device and settings in NSM. This error might be the result of a command that was disabled by another NSM administrator or a local device administrator.

CHAPTER 7

Managing Devices

This chapter describes device management tasks you might need to perform in specific situations, such as upgrading the software version on your devices, obtaining and activating a deep inspection subscription, and handling an RMA device in the Network and Security Manager (NSM) UI.

This chapter also provides details for ScreenOS and IDP about how the components of the NSM management system handle device capabilities, and how device configuration settings are imported and updated. This material is provided for reference only, and does not contain specific configuration tasks.

This chapter also describes the Juniper Update mechanism that enables new operating system versions and new device types to be added to NSM without the need to upgrade NSM. This feature applies only to devices with XML-based schemas.

This chapter contains the following sections:

- [Managing Device Software Versions on page 280](#)
- [Viewing and Reconciling Device Inventory on page 285](#)
- [Managing Large Binary Data Files \(Secure Access and Infranet Controller Devices Only\) on page 289](#)
- [Backing up and Restoring SA and IC Devices on page 301](#)
- [Managing User Sessions for SA and IC Devices on page 304](#)
- [Activating Subscription Services on page 305](#)
- [Managing the Attack Object Database on page 306](#)
- [Updating AV Pattern Files on page 317](#)
- [Updating the Web Category List on page 317](#)
- [Miscellaneous Device Operations on page 318](#)
- [Managing ScreenOS Device Capabilities on page 331](#)
- [Archiving and Restoring on page 337](#)
- [Managing Device Schemas Through the Juniper Update Mechanism on page 338](#)

Managing Device Software Versions

You can use NSM to upgrade or adjust the software on any managed device running ScreenOS, IDP 4.1 or later, Junos, SA or IC..

NSM does not support the upgrade of NetScreen-500 and ISG2000 security devices from ScreenOS 5.1 to ScreenOS 5.2. This migration requires a boot-ROM upgrade; for more details, refer to the *ScreenOS 5.2 Migration Guide*.

When a software upgrade is applied to a Junos device with dual Routing Engines, the upgraded software is applied to both Routing Engines. The backup is upgraded first. The router then reboots and the backup becomes the master. Then the former master is upgraded, as is the standard procedure for upgrading Junos devices with dual Routing Engines.

Upgrading the Device Software Version

Upgrading the operating system is a three-step process:

1. Download the new software image file from the Juniper Networks website to your computer running the client UI.
2. Copy the image file to a repository on the GUI server using the NSM Software Manager, which you access from the Device Manager launchpad by selecting **Manage Device Software** (Select Tools > Software Manager from the menu bar).

The Software Manager lists all software image files in the repository. To add the one you just downloaded, click the Add icon, navigate to the software image file you just downloaded, and then click **Open**.

3. Select the software image file you want to install and the devices you want to install it on using the Install Device Software wizard, which you access from the Device Manager launchpad by selecting **Install Device Software** (Select Devices > Software > Install Device Software).

When you use the Install Device Software wizard, you need to specify the device family by selecting the OS name. All software image files available for the selected OS are then listed for you to choose from. Next, select the devices you want to install the software image file on. When installing software on Junos devices, ensure that the devices you select are in the Managed, In Sync state. The wizard responds by informing you that it is ready to perform the upgrade, or with an error message if the upgrade is not allowed.

You can also choose from a list of additional options:



NOTE: The additional options are available only for the devices running Junos OS and not for devices running ScreenOS.

- Reboot the device after successful installation.
- Remove the software image file from the device after successful installation in order to save space on the device.
- Check the compatibility of the new version with the current configuration on the device.
- Back up the currently running and active file system, which takes a snapshot of the file system of the device in case you encounter problems with the new software version and you need to revert to the current version of device software.

Alternatively, instead of accessing the Install Device Software wizard from the Device Manager launchpad, do the following if you need to install device software on one selected device:

1. From the device tree, right-click on the device in which you want to perform a software upgrade and select **Software > Install Device Software**. The Select Software Image to install dialog box appears.
2. In the Select Software Image box, select the software image that you want to install and click **Next**. The View Selected Details dialog box appears.
3. Verify the details and click **Finish**.



NOTE: Do not change the name of the image file. The name of the image file must be exactly the same as the filename that you download from Juniper Networks, for example, `ns5xp.4.0.3r2.0` or `sensor_4_1r1.sh`.

When upgrading multiple device types, ensure that you have loaded the same version of the image file for each type of device on the Device Server. For example, you can upgrade the firmware on a NetScreen-208, a NetScreen-50, and a NetScreen-5XP at the same time, but the image files for each device type must exist on the Device Server and must be the same OS version.

When a new version of Junos is installed on a device, it internally transforms its configuration to conform to the new version. At the end of the upgrade process, these devices display their Config Status as **Managed, Device Changed**. You need to import the transformed configuration into NSM and then their Config Status becomes **Managed, In Sync**.

When upgrading software on your ScreenOS or IDP devices, you can use different methods depending on the OS version running on the devices:

- For devices running ScreenOS 5.x, the Device Server automatically uses Secure Server Protocol (SSP) to load firmware onto your managed devices. SSP is the protocol used for the management connection between the physical device and the NSM Device Server.
- For devices running IDP Sensor software, the Device Server uses SSH/SFTP to upload and run the upgrade software. The Management Port on the IDP Sensor must be active and reachable over the network via SSH.

Select the Automate ADM Transformation option to automatically update the Abstract Data Model (ADM) when the firmware is loaded onto the managed device. If you deselect this option, the firmware is loaded onto the device, but you cannot manage the device from the UI until the ADM is updated. For example, you might want to deselect this option to first verify that the device is properly operating with the uploaded firmware before managing it from the NSM UI. To enable management, you must reconcile the firmware that you uploaded on the device with the ADM, as described in [“Adjusting the Device OS Version” on page 283](#). For more information about the ADM and NSM components, see [“Managing ScreenOS Device Capabilities” on page 331](#).

For step-by-step instructions on upgrading a device, refer to the *Network and Security Manager Online Help* topic, “Upgrading Firmware on Devices.”

Upgrading a Device Software Version from NSM

If upgrading device software from NSM, you can only upgrade to published versions of the software and not to unpublished versions. You can upgrade the software of NSM-managed devices to unpublished versions only through the device CLI or Web UI. You can continue to manage these devices in NSM after the upgrade.

To add a new device running a software version that is either published or unpublished in the Juniper Update Server, you must first add the device to NSM and then import and update it as described in [“Performing an Update” on page 270](#).

Upgrading a Device Software Version outside NSM

If the software version of a device is upgraded outside NSM, through the device CLI or Web UI, NSM behaves differently depending on whether the upgraded software version is published and whether it is a major release.

To illustrate NSM's behavior in these various scenarios, let us assume that 6.3R1, 6.3R2, 6.3R3, 6.3R4 and 6.4R1 are officially published releases; and 6.3R5, 6.3R6 and 6.4R2 are maintenance releases but are not published in the Juniper Update Server. You see the following behavior:

In a Minor Upgrade

When you upgrade from one version to another version in the same release, you can upgrade from:

- One published version to another published version (6.3R2 > 6.3R3)
- One published version to an unpublished version (6.3R2 > 6.3R5)
- One unpublished version to a published version (6.3R5 > 6.3R2)
- One unpublished version to another unpublished version (6.3R5 > 6.3R6)

After any one of the upgrades mentioned above, the device reboots. When the device reconnects with NSM, its status is **sw inventory out of sync**. NSM does not allow you to import or update the device. You must:

1. Run **reconcile inventory** to synchronize the device software inventory.
2. Import and update the device.

In a Major Upgrade

When you upgrade from one major release version to another, you can upgrade from:

- One published version to another published version (6.3R1 > 6.4R1)
- One published version to an unpublished version (6.3R1 > 6.4R2)
- One unpublished version to a published version (6.3R5 > 6.4R1)
- One unpublished version to another unpublished version (6.3R5 > 6.4R2)

After any one of the upgrades mentioned above, the device reboots. When the device reconnects with NSM, its status is **adjust os version needed** and **sw inventory out of sync**. NSM does not allow you to import or update the device. You must then:

1. Run **adjust os version** which also causes the system to automatically synchronize the inventory.
2. Import and update the device.

Adjusting the Device OS Version

When importing or updating devices, NSM alerts you if it detects a mismatch between the OS running on the managed device and the OS that NSM has recorded for the device.

OS mismatches can occur when:

- A device administrator changes the OS on the device using the Web UI or CLI commands (through a console, Telnet, or SSH session).
- The Automate ADM Transformation option in the Firmware Update Availability dialog box was deselected during a firmware upgrade by NSM. (See [“Upgrading the Device Software Version” on page 280.](#))

To reconcile the OS versions, right-click a device and select **Adjust OS Version** to display the Adjust OS Version Wizard. Follow the directions in the wizard. For step-by-step instructions on how to upgrade a device, refer to the *Network and Security Manager Online Help* topic, “Adjusting the OS Version on Devices.”

If a Junos device was upgraded using the Web UI or CLI, after the Adjust OS Version Wizard, you need to import the device configuration into NSM. This step is necessary because the upgrade process causes the Junos device to internally transform its configuration to conform to the new version. After you import this transformed configuration into NSM, the Config Status of the device becomes Managed, In Sync.

During a firmware upgrade, NSM does not recognize service releases but instead recognizes the last major release running on the device. This results in a version difference in the firmware. [Table 30 on page 284](#) lists the actions Adjust OS Version performs to resolve version differences.

Table 30: Adjust OS Version Directive Actions for Major and Service Releases

Adjust OS Version	Action	Example
Major release to a service release	<p>Performs an Adjust OS version from a major release to a later service release.</p> <p>However, the Adjust OS version does not support changing from a major release to a service release, where the major release when the major release for the service release is earlier than the last running major release of the firmware.</p>	Downgrading from the 11.1R4.4 major release to the 10.4S4.1 service release is not supported, because the major release of the service release is 10.4, which is earlier than the major release of the last running major release of the firmware, which is 11.1.
Service release to a major release	<p>Performs an Adjust OS Version from a service release to a later major release.</p> <p>The action applies only if the service release is upgraded to the next major release.</p> <p>The Adjust OS Version does not support changing from a major release when the major release for the service release that is earlier than the last running major release of the firmware.</p>	Downgrading from 11.1S4.1 service release to the 10.4R4.1 major release is not supported, because the service release is later than the major release.
Service release to a service release	Performs an Adjust OS Version from a later service release to a later service release. However, the directive does not support downgrading to an earlier service release.	Downgrading from the 11.1S4.3 service release to the 11.1S4.1 service release is not supported, because S4.1 is an earlier release.
Major release to a major release	Performs an Adjust OS Version from a later major release to a major later release.	Downgrading from the 11.1R4.1 major release to the 11.2R5.1 major release.

Downgrading the Device OS Version

NSM does not support OS downgrades; you cannot use NSM to install an earlier version of Juniper Networks OS than is currently running on the device. You must use the Web UI or CLI commands to downgrade a managed device, and then add the device to NSM again.

Rolling back the Device OS Version

NSM allows administrators to rollback an IC or SA device to the previous software version that was installed on it.

To rollback to the previous software version:

1. From the device tree, right-click on the device in which you want to perform a software rollback and select **Software > Rollback Device Software**. The Rollback Software dialog box appears listing the OS on the device, the device rollback OS version installed on the backup partition of the device, the device rollback build, and the current NSM managed OS version.

2. Click **Rollback**.

Deleting the Device OS Version

To delete a software image:

1. From the menu bar, select **Tools > Software Manager**. The Software Manager dialog box appears.
2. Select the software image you want to delete and click the delete button (-).

Upgrading Device Support

Use the Upgrade Device Support directive to automatically upgrade all existing devices that are eligible to provide forward or full management support for features in a future release of ScreenOS.



NOTE: You must install a schema patch for the future version of ScreenOS before upgrading the device support.

The directive performs the following actions:

- Performs an Adjust OS Version from the previously known ScreenOS version to the new version of ScreenOS running on the selected devices.
- Optionally performs an import on the selected devices.

Viewing and Reconciling Device Inventory

Device inventory management in NSM allows you to display information about the hardware, software, and license components of each device. It also provides features to update the NSM database with the most current inventory information from the device. In addition, you can use Device Monitor, Device List, and the device tooltip to view the status of inventory synchronization.

These inventory management features are available for all Junos devices, Secure Access devices, and Infranet Controller devices. You cannot use these features with ScreenOS security devices or IDP sensors. You can use these features to make the NSM database match the device inventory, but you cannot write new inventory information to the device.

Initially, the device inventory in the NSM database is generated when the device is first imported into NSM. Immediately after import, the device inventory in the NSM database matches exactly the inventory on the device itself.

If the hardware on the device is changed, the software is upgraded through the Web UI or CLI, new software packages are installed, or a new license key is installed on the device, then the inventory on the device is no longer synchronized with the NSM database.

The Device Monitor, Device List, and tool tip shows the hardware inventory status, the software inventory status, and the license inventory status for each device. Possible states include:

- In Sync
Inventory in the NSM database matches the device.
- Out of Sync
Inventory in the NSM database does not match the device.
- N/A
Either the device is not yet connected and managed by NSM, or the device is a ScreenOS security device or IDP sensor.

The following sections provide details:

- [Viewing the Device Inventory on page 286](#)
- [Comparing and Reconciling Device Inventory on page 287](#)

Viewing the Device Inventory

NSM displays the hardware, software, and license inventory for each device according to the information it has in its database. For a device with dual Routing Engines, NSM collects the inventory data from the master Routing Engine.

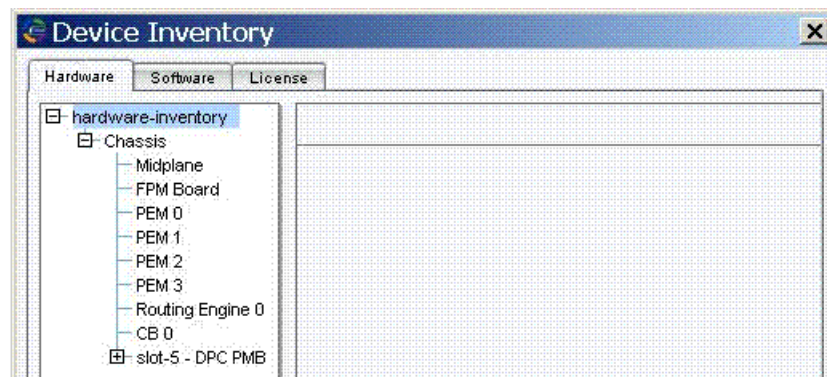
To view the device inventory, the device must be in the “Managed” state.

To view the device inventory, follow these steps:

1. In the navigation tree, select **Device Manager > Devices**.
2. Right-click the device whose inventory you want to view.
3. Select **View/Reconcile Inventory**.

The Device Inventory window opens, similar to the example shown in [Figure 72 on page 286](#).

Figure 72: Viewing the Device Inventory



4. Select the **Hardware** tab to display information about hardware modules in the device, including the I/O module, the Routing Engine, and so on.
5. Select the **Software** tab to display information about the software packages installed in the device, including the installed OS and its version, and any other installed packages.
6. Select the **License** tab to display the license usage summary and details of each installed license.

The usage summary lists the features that are licensed, the capacity of each license (for example, how many VPNs a license supports), how many licensed units are already in use, and how many more are needed.

The license details include the key, name, or ID of the license, the date the license was created, and the validity status of the license.

Comparing and Reconciling Device Inventory

Changes to the device inventory are not automatically updated in the NSM database. The device does not notify NSM of such changes. If a hardware module is added, removed, or replaced, the software is upgraded through the device CLI or Web UI, a new software package is installed, or a new license key is added, the NSM database becomes out of sync with the device inventory.

Run the Inventory Diff tool to check for differences between the NSM database and the device inventory. To run this tool, follow these steps:

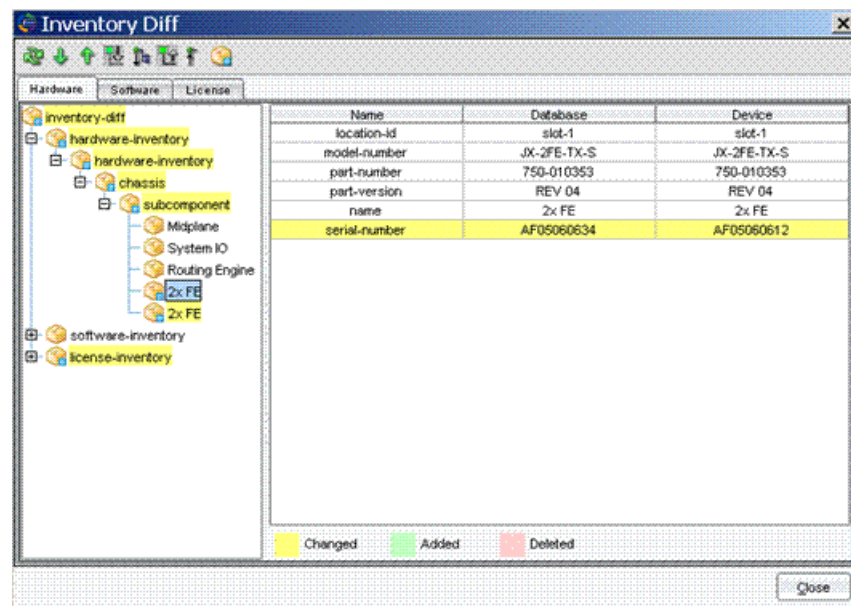
1. In the Device Manager, select **Devices**, and then right-click on the device you want to compare.
2. Select **View/Reconcile Inventory**.

The Device Inventory window appears.

3. Click **Refresh** to start the Inventory Diff tool.

The Inventory Diff window appears, similar to the example shown in [Figure 73 on page 288](#).

Figure 73: Comparing the Device Inventory with the NSM Database



Select an item in the left pane to display information about its objects in the right pane. The right pane shows the name of the object, its value in the database, and its value in the device itself. Differences between the database and the device are highlighted:

- Yellow shows inventory objects that have changed.
- Green shows inventory objects added in the device.
- Red shows objects deleted in the device.

If differences are detected, the inventory status of the corresponding category changes to Out of Sync in the Device List, the Device Monitor, and the device tooltip, and the Reconcile button in the Device Inventory window becomes active.

4. When you have finished viewing the differences, click **Close**.
5. To resolve any differences and update the NSM database with the current inventory, click **Reconcile** in the Device Inventory window.

NSM reimports the updated inventory, first applying changes to the software inventory, then to the license inventory, and finally to the hardware inventory. On completion, NSM displays the updated Device Inventory window.

The inventory status also changes to “Out of Sync” if differences exist between the NSM database and the device inventories when the device reboots and reconnects, or when an Update Device directive is issued to the device. In either case, you can reconcile the inventory immediately by clicking the Reconcile button in the Device Inventory window. You do not need to run the Inventory Diff tool, although you might choose to do so if you want to know what the differences are.

You might also need to reconcile the device inventory after activating a modeled device. The Reconcile button becomes active for this purpose after you issue an Activate Device directive on a device with connection status “Update Needed”.

If you issue an Update Device directive while the inventories are not synchronized, NSM will return an error message telling you that it cannot complete the request until you have reconciled the inventories.

If a member device of a Junos cluster reboots and connects back to NSM, the hardware inventory might show as “Out-of-Sync” in the Device list table because the device takes some time to fully initialize the chassis modules. You can reconcile the inventories to get NSM back in sync with the device.

If the operating system is upgraded using the device CLI or Web UI, the Software Inventory Status will change to “Out of Sync” when the device reboots and reconnects to NSM. For this special case, you must reconcile the inventory by right-clicking the device in the Device Manager and selecting **Adjust OS Version**. It is the only option available from the drop-down list if the operating system versions are not synchronized. See [“Adjusting the Device OS Version” on page 283](#) for details.

Managing Large Binary Data Files (Secure Access and Infranet Controller Devices Only)

Large binary data files that form a part of the configuration of Secure Access and Infranet Controller devices are handled differently from the remainder of the configuration in NSM. The size of some of these binary files could make configurations large enough to overload resources on the NSM server. Consequently, only the large binary files you specify are imported into NSM, and those files are configured as shared objects, which avoids duplication if they are applied to multiple devices.



NOTE: NSM supports binary data files up to 20 MB.

Large binary data files are not imported with the rest of the configuration during a normal device import operation. Instead, the file is represented in the device configuration tree by a stub containing an MD5 hash and file length designation. If you need to manage such a file in NSM, you upload the file separately, and configure it as a shared object. To include the file as part of the device object in NSM, you must then establish a link between the node in the device configuration tree and the shared binary data object. When you establish the link, a pointer to the shared binary data object replaces the MD5 hash and length.

After you have established the link, an Update Device directive will push all linked binary data files to the device along with the rest of the device configuration. No binary data is pushed for nodes that still contain the MD5 hash and length designators.

If you do not need to manage a large binary data file from NSM, then you do not need to include it in the device object configuration. For example, suppose you have a hosted Java applet that resides on a Secure Access device, and you have no intention of updating this applet. In this case, no shared object creation or file upload is necessary. NSM device objects will contain only the MD5 hash stub for these endpoints. Any delta configuration

operation between NSM and the device will indicate identical configurations because the MD5 hash in NSM will match the file on the device. For the same reasons, an Update Device directive will have no effect on the device.

The following sections provide detailed instructions for managing large binary data files in NSM, and specific instructions about how to upload each type of file and link it to the device configuration object.

- [Uploading and Linking Large Binary Data Files on page 290](#)
- [Importing Custom Sign-In Pages on page 294](#)
- [Importing Antivirus Live Update Settings on page 295](#)
- [Importing Endpoint Security Assessment Plug-in \(ESAP\) Packages on page 296](#)
- [Importing Third-Party Host Checker Policies on page 297](#)
- [Importing a Secure Virtual Workspace Wallpaper Image \(Secure Access Devices Only\) on page 298](#)
- [Importing Hosted Java Applets \(Secure Access Devices Only\) on page 299](#)
- [Importing a Custom Citrix Client .cab File \(Secure Access Devices Only\) on page 300](#)

Uploading and Linking Large Binary Data Files

This topic describes the complete procedure for downloading a large binary data file and linking that file into the Secure Access or Infranet Controller device configuration tree. Subsequent sections provide details about each type of large binary data file.



NOTE: NSM supports binary data files up to 20 MB.

To upload and link a large binary data file, follow these steps:

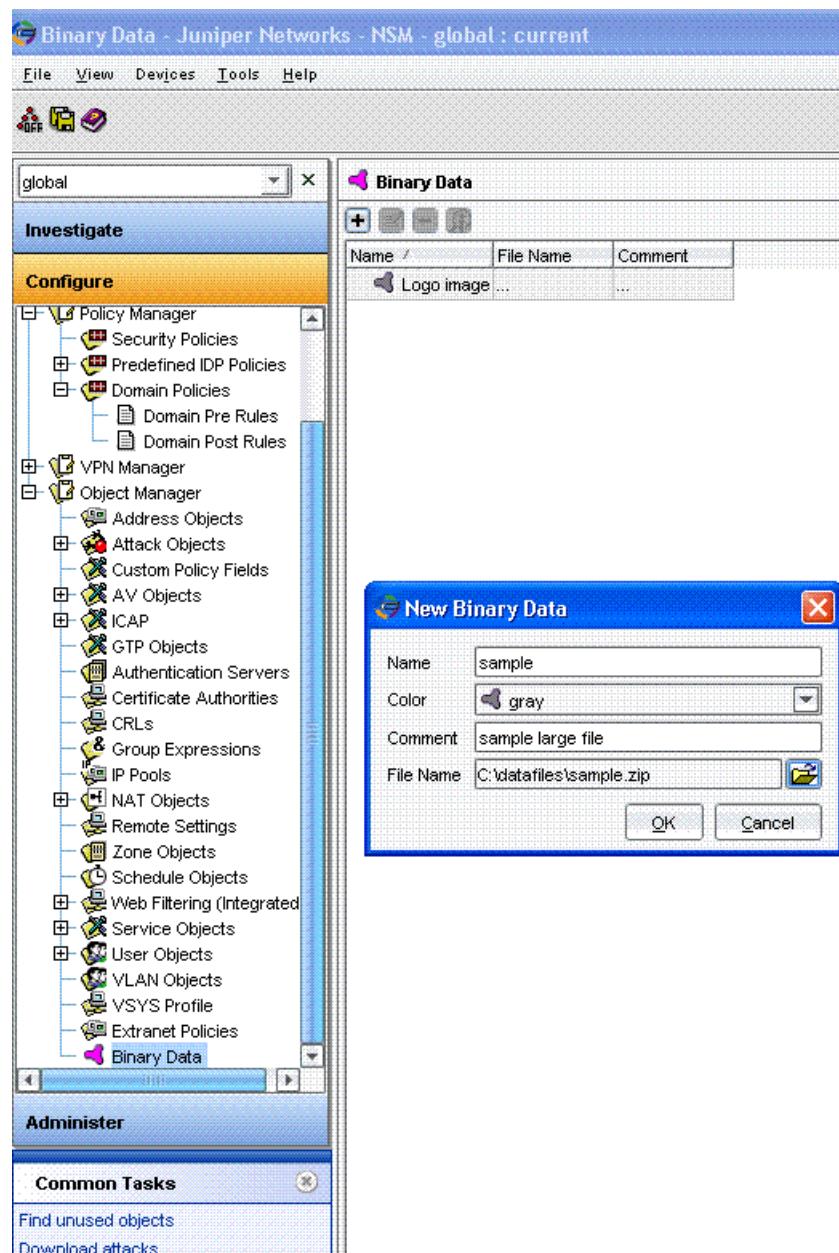
1. In the Device Manager, right-click the device icon and select **Import Device** from the list to import the Secure Access or Infranet Controller device configuration.

When the import job is finished, the device object configuration contains the MD5 stubs for each of the large binary data files.
2. Upload each required large binary data file onto the NSM client workstation.

Use the device Web UI to upload binary files from the Secure Access or Infranet Controller device. Other files, such as ESAP configuration files, should be downloaded from the site of origin.
3. To create a shared object in the NSM Object Manager for the binary file:
 - a. In the Configure panel of the NSM navigation tree, select **Object Manager > Binary data**, and then click the Add icon.

- b. In the Binary Data dialog box, enter a name for the object, select a color for the object icon, add a comment if desired, and select the file you uploaded in Step 2. See [Figure 74 on page 292](#). Click **OK**.

Figure 74: Adding a Shared Binary Data Object



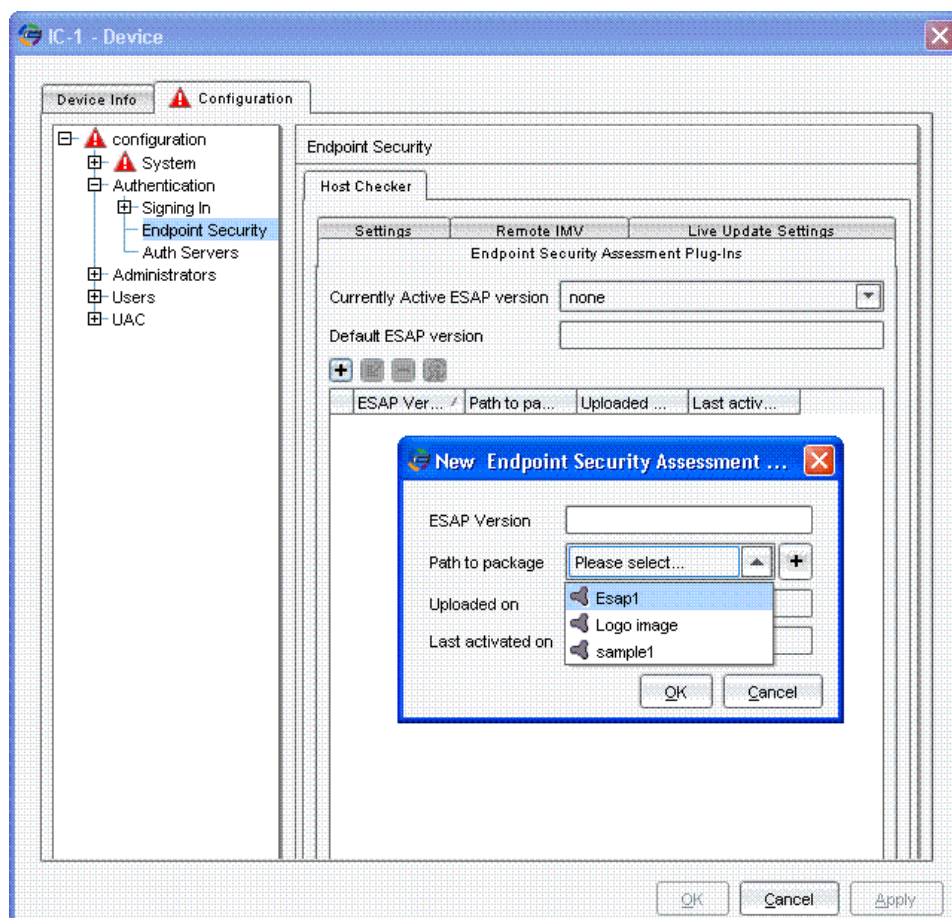
4. Link the shared object to the corresponding node in the device configuration tree:
 - a. In the Device Manager, double-click the Secure Access or Infranet Controller device to open the device editor, and then select the **Configuration** tab.
 - b. Navigate to the node in the configuration where you want to load the binary file.

For example, to load an ESAP package, expand **Authentication** and then select **Endpoint Security**. In the Host Checker tab, select **Endpoint Security Assessment Plug-Ins**, and then click the Add icon.

- c. Select the shared object.

To continue the ESAP example, in the New Endpoint Security Assessment Plug-Ins dialog box, enter a version number, and select a shared binary data object from the **Path to Package** list, as shown in [Figure 75 on page 293](#). This list includes all shared binary data objects. Click **OK**.

Figure 75: Linking to a Shared Binary Data Object



If the object you want is not in the list, you can add it to the shared binary data list by clicking the Add icon. The Binary Data dialog box appears as in step 3.

- d. Click **OK** to save the newly configured links.

Importing Custom Sign-In Pages

The customized sign-in pages feature is a licensed feature that enables you to use your own access pages, rather than modifying the sign-in page included with the Secure Access or Infranet Controller device. Infranet Controller devices can use customized sign-in access pages. Secure Access devices can use customized sign-in access pages and customized sign-in *meeting* pages.

Creating a Custom Sign-In Page

Customized sign-in pages are THTML pages that you produce using the Template Toolkit language and upload to a shared object in NSM in the form of an archived ZIP file.

For more information on customized sign-in pages, including details about how to create them using the Template Toolkit language, see the *Secure Access Custom Sign-In Pages Solution Guide* or the *Unified Access Control Custom Sign-In Pages Solution Guide*.

Linking to a Custom Sign-In Page Shared Object

To create a link from a Secure Access or Infranet Controller configuration tree to a shared object containing a custom sign-in access page, follow these steps:

1. In the Device Manager, double-click the Secure Access or Infranet Controller device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Expand **Signing-In**.
4. Expand **Sign-in Pages**.
5. Select **Users/Administrator Sign-in Pages**, and then click the Add icon in the right pane.
6. Enter a name for the access page.
7. Select **Custom Sign-in Pages**.
8. Select a shared binary data object from the **Custom Pages Zip File** list.
9. Click **OK** once to save the link, and again to save the configuration.

To create a link from a Secure Access configuration tree to a shared object containing a custom sign-in meeting page, follow these steps:

1. In the Device Manager, double-click the Secure Access device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Expand **Signing-In**.
4. Expand **Sign-in Pages**.
5. Select **Meeting Sign-in Pages**, and then click the Add icon in the right pane.
6. Enter a name for the sign-in meeting page.
7. Select **Custom Sign-in Page**.
8. Select a shared binary data object from the **Template File** list.
9. Click **OK** once to save the link, and again to save the configuration.

Importing Antivirus Live Update Settings

Uploading Live Update Settings

Retrieve the latest AV live update file from the Juniper Networks Downloads Web site:

https://download.juniper.net/software/av/uac/epupdate_hist.xml

Retrieve the latest patch file also:

<https://download.juniper.net/software/hc/patchdata/patchupdate.dat>

Linking to a Live Update File Shared Object

To create a link from a Secure Access or Infranet Controller device configuration tree to a shared object containing an antivirus (AV) live update file, follow these steps:

1. In the Device Manager, double-click the Secure Access or Infranet Controller device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Select **Endpoint Security**.
4. From the Host Checker tab, select **Live Update Settings**.

5. Select a shared binary data object from the **Manually import virus signature list**.
6. Click **OK** to save the configuration.

To create a link from a Secure Access or Infranet Controller configuration tree to a shared object containing an AV patch live update file, follow these steps:

1. In the Device Manager, double-click the Secure Access or Infranet Controller device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Select **Endpoint Security**.
4. From the Host Checker tab, select **Live Update Settings**.
5. Select a shared binary data object from the **Manually import patch management data** list.
6. Click **OK** to save the configuration.

Importing Endpoint Security Assessment Plug-in (ESAP) Packages

The Endpoint Security Assessment Plug-in (ESAP) on the Secure Access or Infranet Controller device checks third-party applications on endpoints for compliance with the predefined rules you configure in a Host Checker policy.

Uploading ESAP Packages

To upload the Endpoint Security Assessment Plug-in from the Juniper Networks Customer Support Center to your NSM client computer, follow these steps:

1. Open the following page:
<https://www.juniper.net/customers/csc/software/ive/>
2. To access the Customer Support Center, enter a user name and password for a Juniper Networks Support account.
3. Click the **ESAP** link.
4. Click the **ESAP Download Page** link.
5. Navigate to the ESAP release you want.
6. Upload the plug-in zip file to your computer.

Linking to an ESAP Package Shared Object

To create a link from a Secure Access or Infranet Controller configuration tree to a shared object containing an ESAP package, follow these steps:

1. In the Device Manager, double-click the Secure Access or Infranet Controller device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Select **Endpoint Security**.
4. From the Host Checker tab, select **Endpoint Security Assessment Plug-Ins**, and then click the Add icon.
5. In the New Endpoint Security Assessment Plug-Ins dialog box, enter an ESAP version number.
6. Select a shared binary object from the **Path to Package** list.
7. Click **OK** once to save the link, and again to save the configuration.

Importing Third-Party Host Checker Policies

For Windows clients, you can create global Host Checker policies that take a third-party J.E.D.I. DLL that you upload to the IVE and run on client machines.

Uploading a Third-Party Host Checker Policy

Follow these steps to upload a package definition file that the device will recognize:

1. Name the package definition file **MANIFEST.HCIF** and include it in a folder named **META-INF**.
2. Create a Host Checker policy package by creating a **zip** archive. Include the **META-INF** folder that contains the **MANIFEST.HCIF** file along with the interface DLL and any initialization files. For example, a Host Checker policy package might contain:

META-INF/MANIFEST.HCIF hcif-myPestPatrol.dll hcif-myPestPatrol.ini
3. Upload the Host Checker package to the NSM shared object. You can upload multiple policy packages to NSM shared objects, each containing a different **MANIFEST.HCIF** file.



NOTE: After you upload a Host Checker policy package to the NSM shared object, you cannot modify the package contents. Instead, you must modify the package on your local system and then upload the modified version to NSM.

4. Implement the policy at the realm, role, or resource policy level using the options described in the *Secure Access Administration Guide* or the *Unified Access Control Administration Guide* section about configuring host checker restrictions.

To verify that the package itself is installed and running on the client computer, you can use the name you specified when you uploaded the policy package (for example, **myPestPatrol**). To enforce a particular policy in the package, use the syntax **package-name.policy-name**. For example, to enforce the **FileCheck** policy in the **myPestPatrol** package, use **myPestPatrol.FileCheck**.

Linking to a Third-Party Host Checker Policy Shared Object

To create a link from a Secure Access or Infranet Controller device configuration tree to a shared object containing a third-party host checker policy, follow these steps:

1. In the Device Manager, double-click the Secure Access or Infranet Controller device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Select **Endpoint Security**.
4. From the Host Checker tab, select the **Settings** tab, and then click the Add icon in the Policies box.
5. From the Policy type list, select **3rd Party Policy**.
6. Give the policy a name.
7. Select a shared binary data object from the **Package** list.
8. Click **OK** to save the configuration.

Importing a Secure Virtual Workspace Wallpaper Image (Secure Access Devices Only)

These files are generated by the system administrator.

Uploading a Secure Virtual Workspace Wallpaper Image

Upload the original generated file, and create a shared object for it.

Linking to a Secure Virtual Workspace Wallpaper Image Shared Object

To create a link from a Secure Access device configuration tree to a shared object containing a secure virtual workspace wallpaper image, follow these steps:

1. In the Device Manager, double-click the Secure Access device to open the device editor, and then select the **Configuration** tab.
2. Expand **Authentication**.
3. Select **Endpoint Security**.
4. From the Host Checker tab, select the **Settings** tab, and then click the Add icon in the Policies box.
5. From the Policy type list, select **Secure Virtual Workspace Policy**.
6. Select the **Options** tab.
7. Select a shared binary data object from the **Desktop wallpaper image** list.
8. Click **OK** to save the configuration.

Importing Hosted Java Applets (Secure Access Devices Only)

You can store Java applets of your choice as shared objects in NSM without using a separate Web server to host them. You can then use these applets to intermediate traffic to various types of applications through the Secure Access device. For example, you can upload the 3270 applet, 5250 applet, or Citrix Java applet to shared NSM objects. These applets enable users to establish sessions to IBM mainframes, AS/400s, and Citrix MetaFrame servers through terminal emulators. To enable the Citrix Java ICA client through an IVE session, you must upload multiple Citrix **.jar** and **.cab** files or configure a Citrix Terminal Services resource profile to host the Java applets.

You can upload individual **.jar** and **.cab** files or **.zip**, **.cab**, or **.tar** archive files to NSM shared objects. Archive files can contain Java applets and files referenced by the applets. Within the **.zip**, **.cab**, or **.tar** file, the Java applet must reside at the top level of the archive.

To ensure compatibility with both Sun and Microsoft Java Virtual Machines (JVMs), you must upload both **.jar** and **.cab** files. The Sun JVM uses **.jar** files. The Microsoft JVM uses **.cab** files.



NOTE: Uploading Java applets requires signed ActiveX or signed Java applets to be enabled within the browser to download, install, and launch the client applications.

Uploading a Java Applet

The source of the Java Applet is implementation dependent.

Linking to a Hosted Java Applet Shared Object

To create a link from a Secure Access device configuration tree to a shared object containing a Java applet, follow these steps:

1. In the Device Manager, double-click the Secure Access device to open the device editor, and then select the **Configuration** tab.
2. Expand **Users**.
3. Expand **Resource Profiles**.
4. Select **Hosted Java Applets**, and then click the Add icon in the right pane.
5. Give the applet and file each a name.
6. Select a shared binary data object from the **Applet file to be uploaded** list.
7. Click **OK** once to save the link, and then again to save the configuration.

Importing a Custom Citrix Client .cab File (Secure Access Devices Only)

The custom Citrix client file enables you to provision the Citrix client from the Secure Access device instead of preinstalling it on end-user machines or downloading from another Web server.

Uploading a Custom Citrix Client .cab File

Obtain a Citrix client **.cab** file from an existing Citrix server installation or from the Citrix corporate Web site. Upload it to your NSM client computer and create an NSM shared object for it.

Linking to a Custom Citrix .cab File Shared Object

To create a link from a Secure Access device configuration tree to a shared object containing a Custom Citrix **.cab** file, follow these steps:

1. In the Device Manager, double-click the Secure Access device to open the device editor, and then select the **Configuration** tab.
2. Expand **Users**.
3. Select **User Roles**.

4. Select the **Global Role Options** tab.
5. In the Global Terminal Services Role Options tab, select a shared binary data object from the **Citrix Client CAB File** list.
6. Click **OK** to save the configuration.

Backing up and Restoring SA and IC Devices

NSM allows you to create multiple backup versions of the data in IC and SA devices and store the backup versions in the NSM database. You can create backup versions for multiple devices at one go, or select a backup version of a device and restore this data onto one or more devices.

You can view the backup versions of a selected device, specify the maximum number of backup versions for a device, purge the older backup versions of a device, delete backup versions of a device and edit the comment or note associated with the backup version of a device. You can also set the RMA state on an SA or IC device, activate the device set to the RMA state and do a full restore after activation.



NOTE: NSM users who have the privileges to import a device can perform a backup operation. NSM users who have the privileges to update a device can perform a restore operation.

Backing up an SA or IC Device

To create backup versions of the data in an IC or SA device:

1. From the device tree or device list view, right-click on the device on which you want to perform a backup and select **Backup Manager**. The Backup Manager dialog box appears listing all existing backup versions of the device.
2. Click the **Backup** button in the Backup Manager dialog box. The Backup Comment dialog box appears.
3. Enter the comment that you want to associate with this backup and click **OK**.

Restoring SA or IC Devices

To restore a backed up version of the data to one or more IC or SA devices:

1. From the device tree or device list view, right-click on the device on which you want to perform a backup and select **Backup Manager**. The Backup Manager dialog box appears listing all existing backup versions of the device:

2. Select the backup version you want to restore and click on the **Restore** button. The Restore to Device(s) dialog box appears.
3. Select the device or devices to which you want to restore the backup version and click **OK**.

Backing up multiple SA or IC Devices

To create backup versions of the data in multiple IC or SA devices:

1. Select **Devices > Configuration > Backup Device(s)** from the menu bar. The Backup Device(s) dialog box appears.
2. Enter comments for the backup in the Comments text-box and select the devices which you want to backup and click **OK**.

Configuring Preferences for Backing up and Restoring SA or IC Devices

To configure preferences for backup and restore of SA or IC devices:

1. Select **Tools > Preference** from the menu bar. The New Preferences dialog box appears.
2. In the preference navigation tree, select **Config Backup Settings**.
3. In the Config Backup Settings pane:
 1. Specify the maximum number of backup versions that should be saved for a device. The default value is 3 while the maximum allowed is 5 versions.
 2. Select the **Purge Config File versions** checkbox to automatically purge older backed up versions of the device after the maximum limit of backup versions has been exceeded. If this option is disabled, once the maximum limit for number of backup versions for that device is reached, subsequent backup operations will fail.
4. Click **OK**.

Viewing Backed up Versions for an SA or IC Device

To view backed up versions of an SA or IC device:

1. From the device tree or device list view, right-click on the device and select Backup Manager. The Backup Manager dialog box appears listing the following information for all existing backup versions of the device:
 - Version number of the backup
 - Date and time when the backup was taken
 - OS version running on the device when the backup was taken.
 - Size(in MB) of the data backup.

- Name of the NSM admin who took the backup.
 - Comments that the NSM admin entered while backing up.
2. Click the **Backup** button to backup a device. Select a backed up version and click **Restore** to restore a backed up version to devices. Select a backed up version and click **Edit Comments** to edit the comments associated with a backup. Select a backed up version and click **Delete** to delete the backed up version from the NSM database.



NOTE: The backup and restore feature is available in the NSM UI on root clusters but not on cluster members. However when the backup/restore operation is performed NSM automatically chooses one of the cluster members.

Setting the RMA State on an SA/IC Device

If you need to send a device back to the factory and replace it with a new device, you can set the device to the RMA state. This state allows NSM to retain the device configuration without a serial number or connection statistics. When you install the replacement device, all you need to do is activate the device with the serial number of the replacement unit and do a full restore for the device.

To set the RMA state on an SA/IC device:

1. From the device tree or device list view, right-click on the device on which you want to set the RMA state and select **RMA Device**. The Confirm RMA Device dialog box appears.
2. Click **OK**. The Latest Backup Details dialog box appears. (If you do not have a backup you will be prompted to take a backup before proceeding with).
3. Click **Yes** if you want to take a new backup. Click **No** if the current backup is enough to set the RMA state on the device.

Activating an SA/IC Device Set to the RMA State

1. From the device tree or device list view, right-click on the device that you want to activate and select **Activate Device**. The Activate Device wizard appears.
2. Select either **Device Deployed** or **IP is reachable** or **Device deployed, but IP is not reachable**.
3. If you selected **Device Deployed and IP is reachable**:
 1. Click **Next** and enter the connection information
 2. Enter the IP Address of the security device.

3. Enter the device administrator name and password.
4. After NSM autodetects the device, click **Next** to activate the device in NSM.
4. If you selected Device deployed, but IP is not reachable:
 1. Click **Next**. The Specify the connections settings dialog box opens.
 2. Specify the First Connection One Time Password (OTP) that authenticates the device.
 3. Edit the Device Server Connection parameters, if desired.
 4. Click **Next**. The Specify device connections characteristics dialog box opens.
 5. Click **Show Device Commands** to display a list of CLI commands. The commands enable management and set the management IP address to the Device Server IP address, enable the Management Agent, set the Unique External ID, and set the device OTP. Copy and paste these commands into a text file, and then send the commands to the device administrator. The device administrator must make a Telnet connection to the physical device, paste the commands, and execute them to enable NSM management of the device
 6. Click **OK** to dismiss the Commands window and complete the Activate Device wizard.

Performing a Full Restore of an SA or IC Device

To perform a full restore of an SA or IC device:

- From the device tree or device list view, right-click on the device that you want to restore and select **Full Restore**. All device related data including device configuration, certificate details, network settings and licenses is restored.

Managing User Sessions for SA and IC Devices

NSM allows you to manage user sessions on SA and IC devices. You can query and view user sessions active in a device, refresh the roles of all active user sessions, and delete user sessions from the device.

From the device tree or device list view, right-click on the device on which you want to manage user sessions and select **Active Users**. The Active Users View dialog box appears. The dialog box has the following GUI elements:

The top half of the dialog box has the following GUI elements for user inputs:

1. **Show** text-box to enter maximum number of user sessions to be displayed.
2. **User Name** text-box to enter user name search string. By default, this will be *. You can specify any regular expression string here.
3. **Sort on** drop-down list box to select the name of the field to sort on.
4. **Ordering** drop-down list box to select sort in ascending or descending order

The total records matched field displays the number of user sessions that match a query in the Active Users View dialog box while the total records returned field displays the number of user sessions whose details have been returned as a result of the query.

The bottom half of the dialog box is a table with the following columns detailing all active user sessions. For each user session, the following details will be displayed:

- User name
- Authentication Realm
- User Roles (comma separated list of role-names)
- Sign in time
- Node from which the user signed in (shown only for SA devices)
- IP address assigned for the user's network connect session (shown only for SA devices)

If you have not queried active user sessions using this dialog box, the bottom half of the dialog box will be empty.

- Click on the **Update** button after specifying values for maximum number of user sessions to be displayed and the user name search string, to update details about user sessions in the Active Users View dialog box.
- Click on **Refresh** to refresh the role of all active user sessions.
- Click **Delete** after selecting the user session you want to delete and then click Yes in the confirmation dialog box that appears to delete an active user session.
- Click **Delete All Sessions** and then click **Yes** in the confirmation dialog box that appears to delete all active user sessions on the device. Admin user sessions will not be deleted.



NOTE: The user session management feature for SA and IC devices is available in the NSM UI on root clusters but not on cluster members. However when the operation is performed NSM automatically chooses one of the cluster members.

Activating Subscription Services

To use some Juniper Networks services, such as internal AV or Deep Inspection Signature Service, you must activate the service on the device by first registering the device, and then obtaining the subscription for the service. Even though devices with bundled AV services come with a temporary, preinstalled subscription, you must register your product and retrieve the subscription to receive your fully paid subscription.

To register your product, go to www.juniper.net/support. After you have registered your product, you can retrieve the service subscription.

To obtain the subscription for a service:

1. From the Device Manager launchpad, select **Other**, and then select **Get Entitlement from Entitlement Server**. The Get Entitlement dialog box appears.
2. Select the devices or group of devices for which you want to retrieve a subscription.
3. Click **OK**. The Job Information window displays the status of the subscription retrieval.

Managing the Attack Object Database

The attack object database stored on the device contains predefined attack objects and groups designed to detect known attack patterns and protocol anomalies within network traffic. You use attack objects when using Deep Inspection (DI) or Intrusion Detection and Prevention (IDP) as attack detection mechanisms in a security policy rule.

Keep your attack object database and IDP detector engine firmware current. You can do this interactively from the UI, or you can schedule updates automatically. Juniper Networks provides frequent attack database updates, available for download from the Juniper Networks Web site. New attacks are discovered daily, so it is important to keep your attack object database up-to-date. Verify that the attack database version on the managed device matches the one on the NSM GUI server. The IDP engine is dynamically changeable firmware that runs on IDP Sensors, optional security modules for the ISG Series Integrated Security Gateways and IDP-capable devices.

The following sections explain how to manage the attack object database:

- [Updating the Attack Object Database on page 306](#)
- [Verifying the Attack Object Database Version on page 309](#)
- [Updating the IDP Detector Engine on page 311](#)
- [Example: Confirm IDP Engine Version on page 312](#)
- [Scheduling Security Updates on page 313](#)

Updating the Attack Object Database

You can update the attack object database for managed devices that have deep inspection or IDP capabilities.

- For devices running ScreenOS version 5.0.0-IDP1, ScreenOS 5.1 and later, or standalone IDP, or supported versions of Junos, you must download new attack objects from the attack object database server to the GUI Server, and then download the new objects to your managed devices. IDP attack objects are loaded automatically when an IDP rulebase is loaded; DI attack objects must be loaded manually.
- For devices running ScreenOS version 5.0, you must configure the devices to contact the attack object database server, and then prompt the devices to download new attack objects from the server.

To update a managed device with new DI attack objects, you must first obtain a DI subscription for your device. For details, see [“Activating Subscription Services” on page 305](#).

Updating Attack Objects for IDP-Enabled devices

You can update attack objects by downloading new attack objects and a new detector engine from the attack object database server to the GUI Server, then downloading the new objects to your managed devices.

You can perform a network update if the NSM GUI Server has an Internet connection, either directly or through a proxy. During a network update, the GUI Server contacts the Attack Object Database server (managed by Juniper Networks) and automatically downloads the necessary attack object files.

You can perform a local update if the GUI Server does not have Internet connectivity or you do not want to perform a network update. To prepare for a local update, you manually download the attack objects files from the Attack Object Database server (managed by Juniper Networks), then copy these files to a local directory on the GUI Server. Then, during the local update, you specify the path to these files.

Preparing for a Local Update

Complete the following steps before you perform a local update:

1. Obtain the attack update data file from the Juniper Networks Web site.

Browse to

<https://services.netscreen.com/restricted/sigupdates/nsm-updates/NSM-SecurityUpdateInfo.dat>.

Copy and paste the content from the URL into a text file called **NSM-SecurityUpdateInfo.dat**.

Make sure the file has no HTML tags, RTF tags, or control characters. Use a text editor to make sure there are no control characters in the file. The file should begin with

(updateInfo

and end with a closing parenthesis

)

2. Open the .dat file and locate the "url" line. For example:

:url ("NSMFP6-DI-IDP.zip")

The zip filename is the name of the attack database zip file.

3. Download the attack database zip file from

<https://services.netscreen.com/restricted/sigupdates/nsm-updates/<zipFileName>>.

For example:

<https://services.netscreen.com/restricted/sigupdates/nsm-updates/NSMFP6-DI-IDP.zip>

Download the file to your local disk. Do not change the filename.

4. Put both files in a local directory on the NSM GUI Server or on an internal Web server that is reachable by the NSM GUI Server.
5. Change the permissions on both files to make them readable by all users, but do not change the filenames.

Running the Attack Object Update (Local and Network)

To update the attack object database on the NSM GUI Server:

1. Navigate to the **global** domain.

You can only update the attack object update settings and download a new attack object database from the global domain.

2. Select **Tools > Preferences** to open the New Preferences dialog box.
3. In the preference navigation tree, select **Attack Object**.
4. In the Download URL box for the appropriate device family (ScreenOS or Junos), configure the URL for the attack update file. When you update the attack object database, the management system contacts this server and downloads the latest database version to the GUI Server
 - To perform a network update, enter the URL of the Attack Object Database web server in the Download URL box. To restore the default server, select **Restore Defaults**.
 - To perform a local update, specify the local directory path to the **.dat** file you previously downloaded in the Download URL box. Example:
file:///tmp/NSM-SecurityUpdateInfo.dat
 - To use a proxy server for attack object download, select the **Enable Proxy** check box. Then, enter the proxy server IP address, port, user name, and password.
 - For DI devices, click the **+** button to enter a DI license key and specify Deep Inspection Packs.
5. Select **Tools > Update NSM Attack Database**. The Update NSM Attack Database dialog box appears.
6. Follow the instructions in the Attack Update Manager to download the new Signature and Protocol Anomaly Attack Objects to the NSM GUI Server.

After you have updated the attack object database on the GUI Server, you can use that database to update the attack object database on your managed devices.

IDP attack objects are loaded automatically when you load an IDP rulebase. DI attack objects must be loaded manually.

To load the attack object database update to your managed devices:

1. From the Device Manager launchpad, select **Security Updates > Update Device Attack Database**, or from Devices in the menu bar, select **Deep Inspection/IDP > Update Device Attack Database**. The Update Device Attack Database dialog box appears.
2. Click **Next**, then select the managed devices on which you want to install the attack object update.
3. Follow the directions in the Change Device Sigpack wizard to update the attack object database on the selected managed devices.

Updating DI Attacks on ScreenOS 5.0 Devices

You can update attacks for ScreenOS 5.0 and earlier devices (not 5.0.0 IDP1) by configuring your managed devices to contact the attack object database server, then prompting the devices to download new attack objects from the server.

To configure the device to contact the attack object database server:

1. In the main navigation tree, select **Device Manager > Devices**, and then double-click the device for which you want to configure the database.
2. In the device navigation tree, select **Security > AttackDB > Settings**.
3. For Attack Database Server, enter
`https://services.netscreen.com/restricted/sigupdates`
4. For Mode, select **Update**.
5. Click **OK** to save your changes

To prompt your managed devices to contact the server for updates:

1. From the Device Manager launchpad, select **Security Updates > Update Device Attack Database**, or from Devices in the menu bar, select **Deep Inspection/IDP > Update Device Attack Database**. The Update Device Attack Database dialog box appears.
2. Click **Next**, then select the managed devices that you want to update their attack object database.
3. Follow the directions in the Change Device Sigpack wizard.

Using Updated Attack Objects

After you download updated attack objects and groups to the GUI Server (or to the device), any new attack objects in the update are available for selection in NSM Object Manager. Updated IDP attack objects are also available for selection within an IDP rulebase in a security policy.

You can use new and updated DI attack objects immediately within a DI profile (in a firewall rule), or use the new and updated IDP attack object within an IDP rulebase. When you install the security policy on your managed devices:

- For a security policy that uses IDP attack objects, NSM pushes only the attack objects that are used in IDP rules for the device from the GUI Server to the device.
- For a security policy that uses DI attack objects, NSM pushes all updated signatures from the GUI Server to the device.

Verifying the Attack Object Database Version

New attack objects are added to the attack object database server frequently; downloading these updates and installing them on your managed devices regularly ensures that your network is protected against the latest threats. As new attack objects

are added to the attack object database server, the version number of the database increments by 1. When you download a version of the attack object database from the server, NSM stores the version number of that database.

Automatic Verification

The management system uses the database version number to detect and notify you when the stored attack object database on the GUI server is:

- Older than the most recent database available from the attack object database server
- Newer than the attack object database currently installed on your ScreenOS 5.1 and later managed devices

When NSM detects that managed device contains an older attack object database version than the one stored on the GUI Server, the UI displays a warning for that device, indicating that you should update the attack object database on the device.

Manual Verification

You can also manually check to see if the attack object database on the server is more recent than the one on the security device.

To manually check the attack object database version:

1. From the Device Manager launchpad, select **Security Updates > Check Attack Database Server Version**, or from Devices in the menu bar, select **Deep Inspection/IDP > Check Attack Database Server Version**. The Check Attack Database Server Version dialog box appears.
2. Select the devices or group of devices to be checked.
3. Click **OK**. The Job Information window displays the status of the version check.



NOTE: To view the attack object database version installed on Firewall/IDP devices, place your mouse cursor over the device name in the device list or device tree view. The tooltip displays the Attack DB version. The device list view also includes an Attack Database version column displaying attack object database versions installed on Firewall/IDP devices.

Managing Different Attack Database Versions

Each managed device can contain a different attack object database version. However, the NSM GUI Server can contain only one version of the attack object database at one time. Therefore, when you update the device configuration on a device, you must also update the database on the managed device to match the version of the database on the GUI Server (if the version on the GUI Server is more recent). If the version on the managed device is identical to or more recent than the version on the GUI Server, the device ignores the attack object updates.



NOTE: Although each managed device can contain a different attack object database version, we recommend that you use the most recent version of the attack object database available to ensure that your network is protected against the latest threats.

Although devices running 5.0 update their attack object database independently of the GUI Server, they also must remain synchronized with the attack object database version on the management system if you intend to disable attacks at the device level:

- When the databases are in sync, you can disable attacks at the device level.
- When the databases are out of sync, you cannot disable attacks at the device level. You must update the attack object database on the device using the procedure detailed in [“Updating DI Attacks on ScreenOS 5.0 Devices” on page 309](#).

For details on disabling attacks, see the *Network and Security Manager Online Help* topic, *Configuring Firewall/VPN Devices*.

Example: Updating Devices with Different Attack Object Database Versions

On Monday, you update the attack object database to version 2.0 on the GUI Server, then update two managed devices running ScreenOS 5.2, Device A and Device B. Both devices (and the GUI Server) have the same version of the attack object database.

On Wednesday, in response to a security alert, you update the attack object database to version 2.1 on the GUI server, but install the update on only one of your managed devices, Device A. Device A (and the GUI Server) is now running a different version of the attack object database from Device B.

On Friday, you make miscellaneous configuration changes to Device A and B, then attempt to update both devices with the modeled configuration. During the update, the UI warns you that Device B is running an older version of the attack object database than the GUI Server contains.

Updating the IDP Detector Engine

The IDP engine is dynamically changeable firmware that runs on ISG security devices running ScreenOS 5.0.0-IDP1, standalone IDP appliances, J Series devices, SRX Series devices, and MX Series devices. Automatic updates to the IDP engine occur when you:

- Upgrade security device firmware—The upgraded firmware includes the most recent version of the IDP engine as well as a new version of ScreenOS.
- Manually load a new detector engine—New detector engines may be downloaded with normal attack object updates. You must load the new detector engine onto the device manually.



NOTE: You cannot downgrade the IDP engine version on the device.

To update the IDP engine manually for a ScreenOS or IDP sensor device:

1. From the Device Manager launchpad, select **Security Updates > Update ScreenOS Device Detector**. The Load IDP Detector Engine wizard starts.
2. Click **Next**, and then follow the instructions in the wizard to update the IDP engine on the selected device.

To update the IDP engine for a Junos device:

- From the Device Manager launchpad, select **Security Updates > Update Junos Device Detector**. The Load Junos IDP Detector Engine wizard starts.
- Click **Next**, and then follow the instructions in the wizard to update the IDP engine on the selected device.



NOTE: Updating the IDP engine on a device does not require a reboot of the device.

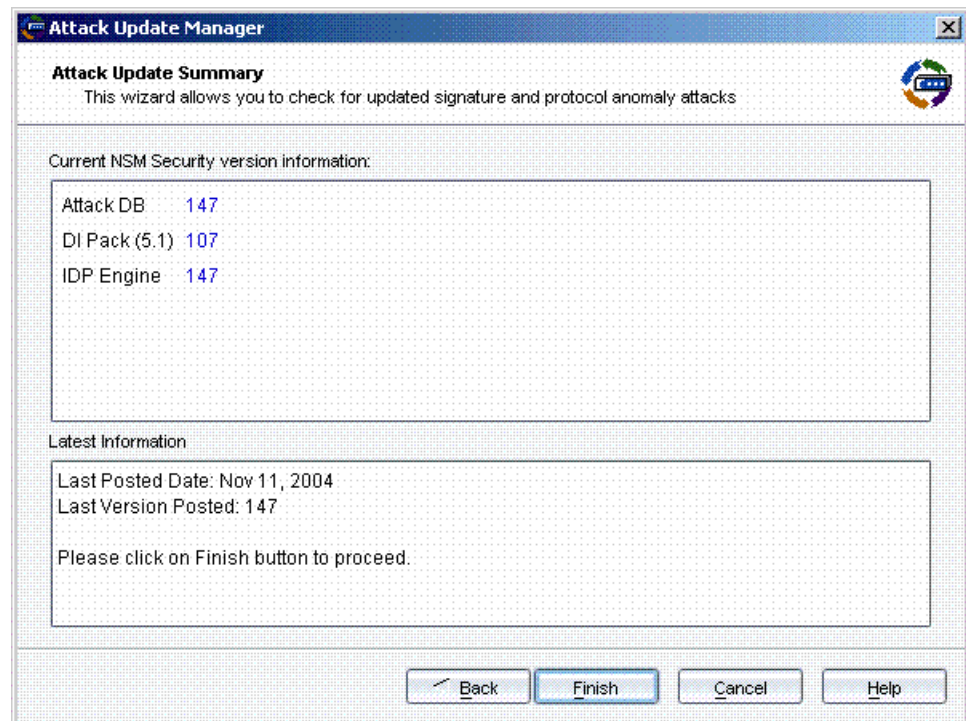
You can also download the new detector engine automatically. See [“Scheduling Security Updates” on page 313](#).

Example: Confirm IDP Engine Version

To see the version of the IDP engine that is currently running on an ISG2000 or ISG1000 device:

1. Select **Tools > View/Update NSM Attack Database**. The Attack Update Manager wizard appears.
2. Click **Next**. The Attack Update Summary displays information about the current version downloaded on the GUI Server and the latest version available from Juniper Networks. See [Figure 76 on page 313](#).

Figure 76: Attack Update Summary



3. Click **Cancel** to exit the Attack Update Manager.

Scheduling Security Updates

For security devices running ScreenOS 5.0.0-IDP1, 5.1 and later, and IDP 4.0 and later, J Series devices, SRX Series devices, and for MX Series devices, you can configure the NSM system to automatically update the attack object database on the GUI Server and on those devices.

For ScreenOS devices running ScreenOS 5.0 (except ScreenOS 5.0.0-IDP1), NSM does not automatically install new attack objects on the device but instead flags the device for manual updating using the UI.



NOTE: Unlike the GUI-based attack object updates, Scheduled Security Updates automatically pushes any new IDP detector engine that has been downloaded.

Using the command-line utility `/usr/netscreen/GuiSvr/utlils/guiSvrCli.sh`, direct the management system to obtain the latest attack objects from the attack database server (managed by Juniper Networks), then specify the action you want the server to take.

For a successful update, the device configuration must be “In-Sync”, meaning that the device is connected and that no configuration differences exist between the configuration on the physical device and the modeled configuration in NSM, or “Sync Pending”, meaning that the device is unconnected and that the physical device will be updated with the

modeled configuration when the device reconnects to the management system. If a device is connected but its configuration is not “In-Sync”, the update process skips that device to avoid installing unexpected changes.

To handle unconnected devices during the update, you must also specify additional post-action options, shown in [Table 31 on page 314](#).

Table 31: Scheduled Security Update (SSU) Command Line Parameters

Parameter	Definition
--dmi	Directs the system to download attacks for devices running Junos OS.
--help	Lists command-line options for <code>guiSvrCli.sh</code> .
--update-attacks	Directs the system to update its attack database by connecting to and downloading the latest attack database, if newer. Requires post-action parameter.
--post-action	Indicates that a post-action instruction will follow (none or update-devices). Requires none or update-devices parameter.
--none	No post-action. SSU updates the attack database, but does not push the new attacks to devices. No other parameters needed.
--update-devices	Updates managed security devices with newly updated attack objects. Requires an unconnected devices handling option (skip or retry).
--skip	Directs the server to skip any unconnected device (server does not try to update attack objects on that device.) No other parameters needed.
--retry	Directs the server to update the device the next time it connects. If the device has changed while offline, the server will take action based on the next parameter. Requires modified device parameter (abort or override).
--abort	Directs the server to abort the update attempt if the device has changed while offline. The device configuration state is set to “Both Changed”, indicating that both the device and NSM have pending changes.
--override	Directs the server to update the device with the new attack objects, overwriting any out-of-band changes made to the device.

Example: Update Attack Objects and Push to Connected Devices

To download a new attack database and push it to connected devices only (ignore unconnected devices), use the following command line.

```
/usr/netscreen/GuiSvr/Utils/guiSvrCli.sh --update-attacks --post-action
--update-devices --skip
```

Scheduling the Update

You can perform a one-time security update using **guiSvrCli.sh** directly, or you can use **crontab** (or another scheduling utility) to configure the update to run at the intervals you desire.



NOTE: Before performing or scheduling a security update, we recommend that you disable the **autoupdate** setting for all managed devices. To disable this setting in the device configuration, from the device navigation tree, select **Security > Attack DB > Settings**, then set the **Schedule Mode** to **Disable**.

To perform a one-time security update:

1. Log in to the NSM GUI Server as root.
2. Change to the utility directory by typing: **cd /usr/netscreen/GuiSvr/Utils.**
3. Type the following to update attacks, including specifying the post-action options for the update:

```
guiSvr.sh --update-attacks --post-action post-action options
```

4. Enter your domain/username and password when prompted.

To configure a scheduled security update using **crontab**:

1. Log into the GUI server.
2. Change to the utility directory by typing: **cd /usr/netscreen/GuiSvr/Utils.**
3. Create a shell script with the following elements:
 - Set the **NSMUSER** environment variable with an NSM domain/user pair. The command for setting environment variables depends on your OS.

```
Example: export NSMUSER=domain/user
```

- Set the **NSMPASSWD** environment variable with an NSM password. The command for setting environment variables depends on your OS and shell.

```
Example: export NSMPASSWD=password
```

- Specify a **guiSvrCli** command string.

```
Example: /usr/netscreen/GuiSvr/utls/guiSvrCli.sh --update-attacks --post-action  
post_action_options
```

4. Make the script executable. Make sure the person who will create the cron job has permission to run the script.
5. Run the crontab editor.

```
crontab -e
```

6. Add an entry for the shell script

```
<minutes after hour> <hour> * * * path/attack_update_shell_script
```

During the update, the **guiSvrCli** utility updates its the attack object database, then performs the post actions. After updating and executing actions, the system generates an exit status code of 0 (no errors) or 1 (errors).

Example: Using Crontab to Schedule Attack Updates

In this example, you use crontab to update attack objects for online managed security devices every day at 5:00 am. This example assumes you are running Linux, that you have a domain called **idp**, and that there is an NSM user called **idpadmin**.

1. Log into the GUI server.
2. Change to the utility directory by typing: **cd /usr/netscreen/GuiSvr/utls.**
3. Create a shell script called **attackupdates.sh** with the following contents.

```
export NSMUSER=idp/idpadminexport  
NSMPASSWD=idpadminpassword/usr/netscreen/GuiSvr/utls/guiSvrCli.sh  
--update-attacks --post-action --update-devices --skip
```

4. Make the script executable.

```
chmod 700 attackupdates.sh
```

5. Run the crontab editor.

```
crontab -e
```

6. Add the script to the crontab.

```
0 5 * * * /usr/netscreen/GuiSvr/utls/attackupdates.sh
```

You can view expanded update results using the Job Manager and Audit Log Viewer in the NSM UI, as detailed in the following sections.

Viewing Scheduled Security Updates in the Job Manager

Each scheduled security update generates a Job Manager entry, entitled Scheduled Attack and Device Update. The entry contains job status information, such as “connected to server” or “no new security update available”.

If the post-action was update-attacks, the job information also includes:

- A list of devices that the server attempted to update with new attack objects.
- For each device, the status of the update, such as “update successful”, “device skipped due to pending changes”, or “update aborted”.

To view a Job Manager entry, in the main navigation tree of the NSM UI, select Job Manager, then double-click the entry you want to view.

Viewing Scheduled Security Updates in the Audit Log Viewer

Each scheduled security update generates an entry in the Audit Log Viewer. The entry contains the following information:

- Time Generated—Specifies the time at which the update began.
- Admin Name/Domain—The administrator name for security update is guiSvrCli and the domain is Global (entry appears as guiSvrCli/Global).
- Action—The action appears as “ Scheduled Attack and Device Update”.

To view an audit log entry, in the main navigation tree of the NSM UI, select **Audit Log Viewer**.

Updating AV Pattern Files

Some security devices provide antivirus (AV) scanning for specific application-layer transactions using an internal AV scanner developed by Trend Micro. The internal AV scanner references a virus pattern file to identify virus signatures. As new viruses emerge, the pattern file on the device needs to be updated.

To update the AV pattern file for a device:

1. From the Device Manager launchpad, select **Security Updates > Update Pattern**. The Update Pattern dialog box appears.
2. Select the devices or group of devices to be updated.
3. Click **OK**. The Job Information window displays the status of the update.

Updating the Web Category List

Web categories (predefined by SurfControl) are used to create the default Web Filtering Profile object, which you can use in a firewall rule to permit or deny specific URL requests to or from your protected network.

The SurfControl CPA server periodically updates its predefined category list, but does not notify its clients when the list is updated. To ensure that the security device and NSM use most up-to-date predefined categories, you must update the list manually, first on the device, then for the NSM system.



NOTE: The security device periodically polls the CPA server for category updates. The default interval is every two weeks; for details on changing this settings, see the *Network and Security Manager Online Help* topic, “Configuring Firewall/VPN Devices”.

You must perform both steps listed below, in the following order:

1. In the Device Manager launchpad, select **Security Updates > Update Web Categories**. This option updates the security device predefined categories from the SurfControl CPA server.

You must perform this step before updating the categories on the NSM management system. When the Select Devices dialog box appears, select the security device you want to contact SurfControl.

2. In the Device Manager launchpad, select **Security Updates > Update System Categories**. This option updates the NSM management system predefined categories from a security device.

You must perform this step after updating the predefined categories on the security device.

Miscellaneous Device Operations

This section describes other device management tasks that you can perform using the NSM UI.

The following sections describe each management task:

- [Launching a Web UI for a Device on page 319](#)
- [Launching a Telnet CLI Window on page 319](#)
- [Rebooting Devices on page 319](#)
- [Refreshing DNS Entries on page 320](#)
- [Updating the Device Clock with an NTP Server on page 320](#)
- [Setting the Root Administrator on a Device on page 321](#)
- [Failing Over or Reverting Interfaces on page 322](#)
- [Setting the RMA State on a Device on page 322](#)
- [Managing Existing and Adding New Devices Using MIP IPv6 Addresses on page 323](#)
- [Upgrading the OS Version During an RMA-Activate Device Workflow on page 330](#)
- [Troubleshooting a BGP Peer Session on a Device on page 330](#)
- [Reactivating Wireless Connections on page 331](#)
- [Finding Usages on page 331](#)

Launching a Web UI for a Device

You can launch a web UI for any device listed in the Device Manager if the device IP address is available. If the SSH connection between the device and NSM is down, you might want to start a web UI on the device to troubleshoot the device.

To launch a web UI on a device, follow these steps:

1. In the Device Manager, select **Devices**.
2. Right-click the device for which you want to start a web UI.
3. Select **Launch Web UI**.

Launching a Telnet CLI Window

You can launch a Telnet CLI window from NSM for all connected Junos OS-based devices (their connected status is up).

1. Right-click on a device to open a Device Manager menu.
2. Select **Launch Telnet** to open the Telnet login window. Log in and issue device commands as desired. User credentials for Telnet are verified by the device and not by NSM.

Invoking the **Launch Telnet** menu item causes the Telnet window to appear even if the Telnet service is not enabled in the device. The **Launch Telnet** menu is disabled if:

- The device is not connected (connected status is down).
- The device does not have an IP address.



NOTE: NSM invokes the default Telnet client provided in the Windows and Linux operating systems.

Rebooting Devices

To reboot a device:

1. Select **Devices > Reboot Device** from the NSM UI.
2. Select the operating system (OS). NSM displays a list of connected devices (connectivity status is up) belonging to the chosen OS.

3. If you chose any OS family other than Junos, select the devices to be rebooted.
4. Select **OK**.

Using the existing scheduled reboot functionality of the Junos devices, NSM allows you to choose one of the following options in the **Reboot Device(s)** window.

- **Reboot now:** This causes an immediate reboot.
- **Reboot after delay of x minutes:** This reboots the selected devices after the specified delay period. Minimum and maximum time limits are based on the Device CLI.
- **Reboot at device local time:** This allows you to choose the date and time for a device reboot. Each device interprets this as its local time. A tool tip in the UI reminds you that the time selected is the device local time.

You can cancel a scheduled reboot. NSM disables the cancel option if you set reboot options and vice versa.



NOTE: NSM also provides a reboot option in the popup menu for EX Series switches, SA and IC devices. You can right-click on the device you want to reboot from the device tree, and select **Reboot Device**.

Refreshing DNS Entries

To enable a security device to use Domain Name System (DNS) to resolve domain names to IP addresses, you configure the IP addresses of the primary and secondary DNS servers on the device. The device can automatically refresh entries in its DNS table by checking them with the specified DNS server at regularly scheduled times or intervals, or after an HA failover.

You can also manually direct the device to refresh its DNS table entries. When you direct the device to refresh its DNS entries, it connects to the previously configured DNS server to perform a lookup of each entry in its table.

To direct one or more devices to refresh their DNS table entries:

1. From the Device Manager launchpad, select **Others > Refresh DNS Entries** from the Devices menu. The Refresh DNS Entries dialog box appears.
2. Select the devices or the group of devices on which DNS tables should be refreshed.
3. Click **OK**. The Job Information window displays the status of the refresh.

Updating the Device Clock with an NTP Server

The security device can use the Network Time Protocol (NTP) to synchronize its system clock with a configured NTP server over the Internet. You can configure the device to perform this synchronization automatically at specific time intervals (see the *Network and Security Manager Online Help* topic, “Configuring Firewall/VPN Devices”), or you can

direct the device to synchronize its clock immediately to a previously-configured NTP server, as described in the following steps.

To direct one or more devices to synchronize their clocks:

1. From the Device Manager launchpad, select **Others > Perform NTP Time Update**. The Perform NTP Time Update dialog box appears.
2. Select the devices or group of devices that should be synchronized with NTP servers.
3. Click **OK**. The Job Information window displays the status of the synchronization.

Setting the Root Administrator on a Device

All security devices ship with the same default login and password for the root administrator. Because these default settings are known, you should change the login and password for the root administrator as soon as possible and as often as necessary.



NOTE: All passwords handled by NSM are case-sensitive.

Each security device can have only one root administrator, who has the following privileges:

- Manages the root system of the security device
- Adds, removes, and manages all other administrators
- Establishes and manages virtual systems, and assigns physical or logical interfaces to them
- Creates, removes, and manages virtual routers
- Adds, removes, and manages security zones
- Assigns interfaces to security zones
- Performs asset recovery
- Sets the device to FIPS mode
- Resets the device to its default settings
- Updates the OS
- Loads configuration files

After you change the root administrator login and password, only persons who know the new login and password can log into the device and perform the tasks listed above.

To configure the login and password for the root administrator for a security device:

1. In Device Manager, right-click a device icon and select **Admin > Set Root Admin**. The Set Root Admin dialog box appears for the device.
2. Enter the new name in the Administrator Name field.

3. Enter the new password in the Password field and then reenter the password in the Confirm Password field.
4. Click **OK**.

For more details on managing device administrators, including the root administrator, see the *Network and Security Manager Online Help* topic, "Configuring Firewall/VPN Devices".

Failing Over or Reverting Interfaces

Some security devices support port modes that bind a second backup interface to the untrust zone. For these port modes, the backup interface is used only when there is a failure on the connection through the primary interface or when you manually force traffic from the primary interface to the backup.

To force a security device to fail over to the backup interface:

1. Right-click a device from the security device Tree or the security device List tab in the Device Manager and select **Admin > Failover**. The Failover Action dialog box appears.
2. Click **Force to Failover**.
3. Click **OK**.

To force a security device to revert to the primary interface:

1. Right-click a device from the security device tree or the security device List tab in the Device Manager and select **Admin > Failover**. The Failover Action dialog box appears.
2. Click **Force to Revert**.
3. Click **OK**.

Setting the RMA State on a Device

If you need to send a device back to the factory and replace it with a new device, you can set the device to the RMA state. This state allows NSM to retain the device configuration without a serial number or connection statistics. When you install the replacement device, all you need to do is activate the device with the serial number of the replacement unit.



NOTE: The replacement device must be the same platform and OS version as the unit that is being replaced. Setting the RMA state cannot be undone.

In the RMA state, the device object is functionally identical to a modeled device, but its status is "RMA" in the Device Monitor.

To set a device to the RMA state:

1. Right-click a device from the security device tree or security device List tab in the Device Manager and select **RMA Device**. The Confirm RMA Device dialog box appears.
2. Click **OK**. In the Device Monitor window, the device status is RMA.

When the replacement device is installed, activate the device with the serial number of the replacement. For information about activating a device, see [“Activating a Device” on page 135](#).



NOTE: If you are placing an HA device in the RMA state, we highly recommended that you perform a flash sync immediately after activating and updating that device. This will ensure that the configuration is synchronized from the device that is not in the RMA state to the device that is in the RMA state. The flash sync ensures that the two HA devices are in sync in NSRP. It also ensures that the case sensitivity of the original zone name is preserved on the device in the RMA state.

Managing Existing and Adding New Devices Using MIP IPv6 Addresses

Beginning in NSM 2012.2R10, NSM supports management of devices with IPv6 addresses. NSM manages devices with IPv6 addresses in the following two scenarios:

- When both the device and the NSM server are configured with IPv6 addresses
- When the device is configured with an IPv6 address and the NSM server is configured with both an IPv4 address and the device server mapped IP (MIP) in the IPv6 address

Configuring a MIP Address in NSM

You can configure the device server to use a MIP address. A MIP maps the destination IP address in an IP packet header to another static IP address, enabling the managed device to receive incoming traffic at one IP address, and automatically forward that traffic to the mapped IP address. MIPs enable inbound traffic to reach private addresses in a zone that contains NAT mode interfaces.

The following sections provide detailed instructions for converting the IP format of devices running ScreenOS and SRX Series high-end devices from IPv4 to IPv6 using an NSM MIP address and adding the devices with IPv6 addresses to the NSM server with a MIP IPv6 address.

- [Converting Device Management IP Format from IPv4 to IPv6 Using an NSM MIP Address on page 323](#)
- [Adding a New Device with an IPv6 Address to the NSM Server with a MIP IPv6 Address on page 326](#)

Converting Device Management IP Format from IPv4 to IPv6 Using an NSM MIP Address

This section provides a detailed description of how to convert the IP format of the following devices from IPv4 to IPv6 using an NSM MIP address:

- [ScreenOS Devices on page 324](#)
- [SRX Series High-End Devices on page 325](#)

ScreenOS Devices

To convert the IP format of a device running ScreenOS from IPv4 to IPv6 using an NSM MIP address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Double-click the device running ScreenOS in Devices.
4. Click **Network>Interface**.
5. Configure the IPv6 address in the management interface.
6. Click **Network>Virtual Router**.
7. Edit the default virtual router.
8. Click **Routing Table** and add a route entry for the DevSvr MIP IPv6 address.
9. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update; when the job status displays successful completion, click **Close**.
10. Double-click the device running ScreenOS again in Devices.
11. Click **Info>Startup**.
12. Click **Use Device Server Through MIP** and select the MIP IPv6 address in the **Device Server IP Address** drop-down menu. Click **OK**.
13. Right-click on the device and select **RMA Device**.
14. Right-click on the device, select **Activate Device**, and click **Next**.
15. Click the **IPv6 address** radio button and configure the management IPv6 address of the device.
16. Click **Next** to add the device to NSM.
17. After the device is added, click **Next** to update the device configuration.

18. Click **Finish** to complete the Add Device wizard.
19. Double-click the device in Device Manager to view the updated configuration.

To check the device configuration status, mouse over the device in Device Manager (you can also check configuration status in Device Monitor). The device status displays as **Managed**, indicating that the device has connected and the management system has successfully imported the device configuration.

SRX Series High-End Devices

To convert the IP format of an SRX Series high-end device from IPv4 to IPv6 using an NSM MIP address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Double-click the SRX Series high-end device in **Devices**.
4. Click **Configuration>Interface**.
5. Configure the IPv6 address in the management interface.
6. Click **Routing-Options** and add a route entry for the DevSvr MIP IPv6 address.
7. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update; when the job status displays successful completion, click **Close**.
8. Double-click the SRX Series high-end device again in **Devices**.
9. Click **Info>Startup**.
10. Click **Use Device Server Through MIP** and select the MIP IPv6 address in the **Device Server IP Address** drop-down menu.
11. Click **Configuration>System>Services>Outbound ssh>Client** to edit the client.
12. Click **Servers** to delete the IPv4 address and configure an IPv6 address for DevSvr.
13. Click **Ok**.
14. Click **Tools>Preferences>Device Update>Netconf**; and uncheck the **Use confirmed commit** and **Rollback candidate config to running config in error** check boxes.

This ensures that the configuration updates will not roll back when the device connection is lost because the IP format changed from IPv4 to IPv6.

15. Update the device configuration by right-clicking the device and selecting **Update Device**. The Job Information box displays the job type and status for the update; when the job status displays successful completion, click **Close**.

Adding a New Device with an IPv6 Address to the NSM Server with a MIP IPv6 Address

This section provides a detailed description of how to add the following devices with IPv6 addresses to the NSM server with a MIP IPv6 address:

- [Adding and Importing Devices with Static IPv6 Addresses on page 326](#)
- [Adding Devices with Dynamic IPv6 Addresses on page 328](#)

Adding and Importing Devices with Static IPv6 Addresses

A static IPv6 address is an IPv6 address that does not change. Not all device families support static IPv6 addresses.

To import a device running ScreenOS and SRX Series high-end devices with a known IPv6 address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device is Reachable** (default).
5. Click **Next**. The Specify Connection Settings dialog box opens.
6. Enter the following connection information:
 - IPv6 address of the security device



NOTE: Select the IPv6 radio button to configure the IPv6 address of the security device.

- Username of the device administrator
- Password for the device administrator



NOTE: All passwords handled by NSM are case-sensitive.

- Select the connection method (Telnet, SSH version 1, SSH version 2) and the port number for the selected service.

If you selected Telnet, click **Next** and go directly to step 7.

If you selected an SSH version, click **Next** and the Verify Device Authenticity dialog box opens. The device wizard displays the RSA Key FingerPrint information; to prevent man-in-the-middle attacks, you should verify the fingerprint using an out-of-band method.

7. After the wizard displays the autodetected device information, verify that the device type, ScreenOS/SRX high-end version, and the device serial number are correct. NSM autodetects the hostname configured on the device and uses it as the device name. You can also change the autodetected hostname.

8. Modify the autodetected device name within the device from its config editor page in NSM, and select **Update device**.

If you modified the device hostname through the Junos OS CLI, SNMP, or J-Web interface, you can modify the device name again in NSM after importing the device, using the edit option. If the device was bulk added, the name you specify in the CSV file is used.



NOTE: If you select the **Device is not reachable** or the **Model Device** workflow, NSM cannot detect the hostname automatically. You need to specify a device name.

9. Select the device server connection parameters: Use a MIP to configure the device to connect to the NSM device server through a mapped IPv6 address and port.

10. Click **Next** to add the device to NSM.

11. After the device is added, click **Next** to import the device configuration.

12. Click **Finish** to complete the Add Device wizard.

13. Double-click the device in Device Manager to view the imported configuration.

To check the device configuration status, mouse over the device in Device Manager (you can also check configuration status in Device Monitor). The device status displays as **Managed**, indicating that the device has connected and the management system has successfully imported the device configuration.

Adding Devices with Dynamic IPv6 Addresses

A dynamic IPv6 address is an IPv6 address that changes. To add a device that uses a dynamic IPv6 address, the device must support NACN.

To import a device running ScreenOS or an SRX Series high-end device with an unknown IP address:

1. From the domain menu, select the domain in which to import the device.
2. In Device Manager, select **Devices**.
3. Click the Add icon and select **Device** to open the Add Device wizard.
4. Select **Device Is Not Reachable**, and then click **Next**.

The Specify Connection Settings dialog box opens.

5. Enter a name for the device and select a color to represent the device in the UI.
6. From the OS Name list, select ScreenOS/IDP or JUNOS. Select the device platform type and the Managed OS version running on the device from the other pull-down menus. If desired, enable Transparent Mode.
7. Select the license key model for the device. Available selections depend on the type of security device and can include baseline, advanced, extended, plus and 10-user.
8. Select the device server connection parameters: Use a MIP to configure the device to connect to the NSM device server through a mapped IPv6 address and port.

9. Click **Next**, and then perform the following tasks on the Specify One-Time Password screen:
 - a. Make a note of the unique external ID for the device. The device administrator will need it to connect the device to NSM. This ID number represents the device within the management system. The wizard automatically provides this value.
 - b. Specify the First Connection One Time Password (OTP) that authenticates the device.



NOTE: All passwords handled by NSM are case-sensitive.

- c. Click **Show Device Commands** to display the list of CLI commands that must be executed on the device to connect to NSM. The commands enable management and set the management IP address to the device server IPv6 address, enable the management agent, set the unique external ID, and set the device OTP.
 - d. Copy and paste these commands into a text file.
 - e. Click **Finish** to complete the Add Device wizard and include the new device in the Device Manager list.
10. Add the commands to the device console. Send the commands to the device administrator. The device administrator must make a Telnet connection to the physical device, paste the commands, and execute them to enable NSM management of the device.
11. To check the device configuration status, mouse over the device in Device Manager or check in Device Monitor.

The status message “Waiting for 1st connect” might appear briefly.

After the device connects, the status displays “Import Needed”, indicating that the device has connected but the management system has not imported the device configuration yet.
12. Import the device configuration by right-clicking the device and selecting **Import Device**. The Job Information box displays the job type and status for the import; when the job status displays successful completion, click **Close**.
13. After the import is completed, double-click the device in Device Manager to view the imported configuration.

To check the device configuration status, mouse over the device in Device Manager or check in Device Monitor. The device status displays as Managed, indicating that the device

has connected and the management system has successfully imported the device configuration.

Upgrading the OS Version During an RMA-Activate Device Workflow

As of Release 2007.3, you can use NSM to upgrade the OS version of your replacement device rather than upgrading it manually.

To upgrade the OS version:

1. Activate the replacement device. For information about activating a device, see [“Activating a Device” on page 135](#). After activation, if NSM detects that the replacement device is running an older OS version than the replaced device, it prompts you to either continue or reject device activation.
2. Click **Next** to put the device in “Firmware upgrade needed” state. If you click Cancel, the activation fails.

The Software Manager allows you to upgrade the firmware version in the physical device before RMA. After upgrading, NSM puts the device in the “Update needed” state.



NOTE: The current OS version of the device is also stored in the device object and is visible in the UI. The directives allowed for this device are “Firmware upgrade,” “Adjust OS version,” and “RMA.”

Considerations for vsys devices and NSRP directives:

- When a device is in the “Firmware upgrade needed” state, directives are not allowed for its vsys devices.
- A peer device cannot perform “Upgrade firmware” and “Adjust OS version” directives on a cluster member device that is in “Firmware upgrade needed” state, but directives such as “Update device” and “Import device” are allowed.
- NSRP directives are not allowed for either of the two cluster member devices.

Troubleshooting a BGP Peer Session on a Device

To troubleshoot BGP peer configurations, you can connect and disconnect BGP connections to a specific neighbor. You can also test the TCP connection to a specific neighbor. To perform these tests, you need to have configured a virtual router and the BGP dynamic routing protocol on the device, and enabled BGP on the virtual router and on the interface to the BGP neighbor.

To connect or disconnect to a BGP peer:

1. In the main navigation tree, select **Device Manager > Devices**. Right-click a device and select **Admin > Modify BGP Peer Session**. The Modify BGP Peer Session dialog box appears.
2. Select the virtual router in which the BGP configuration resides.

3. Select the peer to which you want to connect or disconnect from the list of configured BGP neighbors.
4. Select **Connect** to establish a BGP connection to the selected peer, to terminate the BGP connection to the selected peer, or **TCP Connect** to test the TCP connection to the selected peer.
5. Click **OK**.

Reactivating Wireless Connections

You can deploy a Juniper Networks NetScreen-5GT Wireless security device running ScreenOS 5.0.0-WLAN as a wireless access point (WAP). When you make changes to the wireless settings for the security device, you must update the device with your changes before the new settings take effect. Additionally, the device must reactivate its WLAN subsystem to use the new settings. The NSM automatically reactivates the WLAN subsystem within the wireless security device during the device update process.



NOTE: When using an authentication server for wireless authentication, if you enable 802.1X support on that server, you must also reactive the WLAN subsystem before the change can take effect.

The reactivation process takes approximately 10 seconds. During reactivation of the WLAN subsystem, the device severs all wireless connections and clears all wireless sessions from the session table. Previously connected wireless clients must reconnect to reestablish their disrupted sessions.

For details on configure wireless settings, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

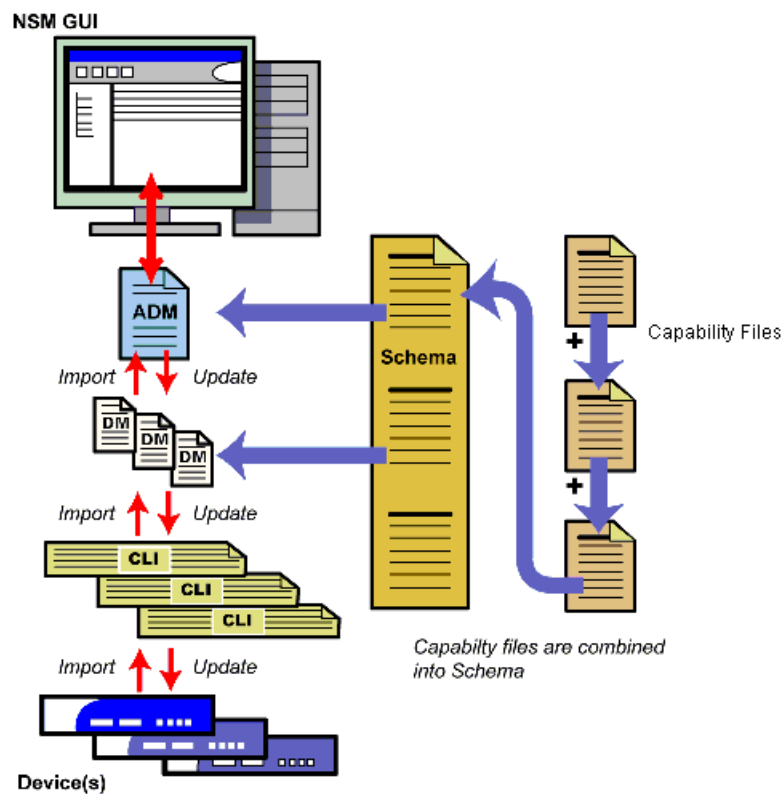
Finding Usages

To locate groups, vsys, policies, and VPNs that reference a specific device, right-click a device and select Find Usages. The Find References box appears.

Managing ScreenOS Device Capabilities

This section presents a detailed description of how NSM components enable you to add, configure, update, and manage ScreenOS security devices. [Figure 77 on page 332](#) is an overview of the components and how they interact with each other. A description of each component follows.

Figure 77: Import/Update Architecture



Abstract Data Model

The Abstract Data Model (ADM) is an XML file that contains configuration data for all objects in a specific domain. The ADM is stored in the GUI Server, but you do not access the ADM directly. When you create, update, or import a device, the GUI Server edits the ADM to reflect the changes. The Management console uses the ADM to determine the current options, fields, screens, and data range to display in the UI for each object.

Data Model

A Data Model (DM) is an XML file that contains configuration data for an individual device. The DM is stored in the Device Server. When you create, update, or import a device, the GUI Server edits the ADM to reflect the changes, then translates that information to the DM.

Data Model Schema

The structure of the ADM and DM is determined by the Data Model (DM) schema. The DM schema reads from a *device capability file* to determine the supported features for the ScreenOS version that is running on the managed devices. A device capability file lists the fields and attributes that a specific ScreenOS version supports.

Your network may contain similar security devices that are running different ScreenOS versions. For example, a NetScreen-5XT may run ScreenOS 5.x, which supports the Routing Information Protocol (RIP), while another NetScreen-5XT runs ScreenOS 4.0.0r2,

which does not support RIP. The DM schema links to the appropriate device capability file for each device.

Device capability files make it easier to integrate devices into NSM and also make upgrading the software on your security devices easier. Each software release includes device capability files that describe the new and changed fields, attributes, and allowable ranges of values.

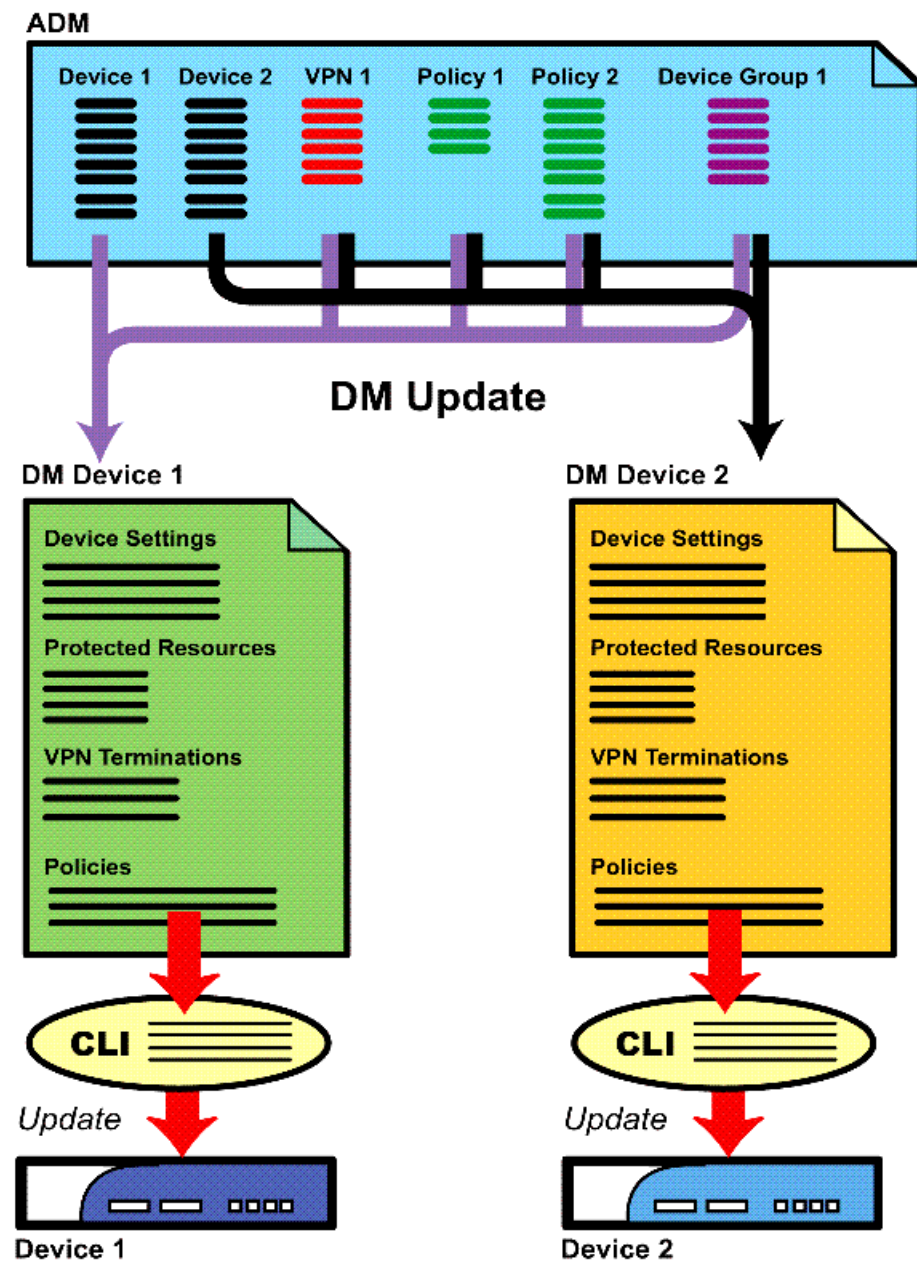
Data Model Updating

Data Model update is the process of translating the objects and object attributes in the ADM domain into individual DMs with device-specific configuration information.

In the ADM, objects are arranged similarly to objects in the management console: each item (VPN, policy, device, device group, and so on) is represented by an object. In the DM, each item is a property of a single device. During the data model update process, the GUI Server identifies the objects that contain properties for a device, and translates those object properties into properties of that device.

When you update a device configuration using the management console, the GUI Server translates the objects and object attributes in the ADM domain into device configuration information in a DM. The Device Server then translates the device configuration information in the DM into CLI commands and sends the commands to the device. See [Figure 78 on page 334](#).

Figure 78: Data Model Update



For example, the ADM contains a VPN with tunnel interfaces, a routing table, and users. When you update a selected device, the DM update identifies the devices that are involved in the VPN and creates interfaces, routing tables, users, and VPN rules in the DM for each device. The DM contains only the VPN information that relates to the specific device, not the entire VPN.

During the device model update process:

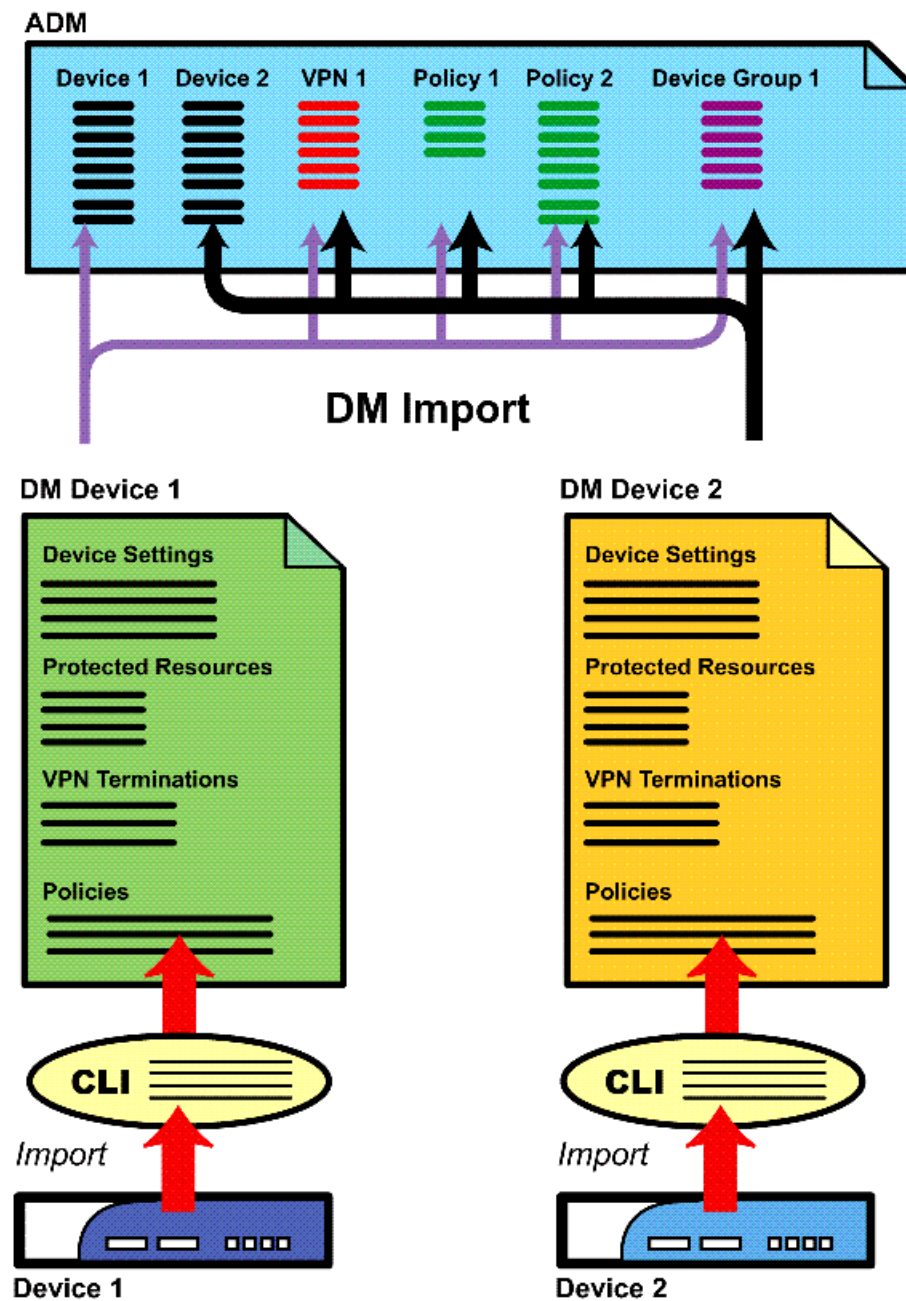
- The GUI Server translates the object and object attributes in the ADM domain into device configuration information in a DM.
- The Device Server translates the device configuration information in the DM into CLI commands.
- The Device Server sends the CLI commands to the device.

Data Model Importing

Data Model Import (DM import) is the process of translating the device-specific configuration information in individual DMs into the objects and object attributes in the ADM domain.

When you import a device configuration using the management console, the device sends CLI commands to the Device Server, which translates the CLI commands into a DM with device configuration information. The GUI Server then translates the device configuration in the DM into objects and object attributes in the ADM, and uses the ADM to display current information in the management console.

Figure 79: Data Model Importing



During the device model import process:

- The device sends CLI commands to the Device Server, which translates the CLI commands into a DM with device configuration information.
- The GUI Server translates the device configuration in the DM into objects and object attributes in the ADM.

The GUI Server then reads the ADM and displays the current information.

Archiving and Restoring

You can archive and restore log and configuration data in NSM using standard Unix commands. Logs reside on the Device Server; all other configuration information, including device configuration data, administrators, policies, audit logs, and job information, resides on the GUI Server.



NOTE: You can also configure the management system to perform local database backups on a regular basis during installation or upgrade. Refer to the *Network and Security Manager Installation Guide* for more information.

Before archiving, you must stop the processes running on both servers, then use the `ls -al` command to identify the actual paths of the GUI Server and Device Server data directories:

- For all information on the GUI Server: `/usr/netscreen/GuiSvr/var`
- For all information on the Device Server: `/usr/netscreen/DevSvr/var`

These directories are links representing the paths that were entered at the time the servers were installed.

After archiving, restart the processes on both servers. For details on stopping, starting, and restarting processes on the management system, refer to the *Network and Security Manager Installation Guide*.

Archiving Logs and Configuration Data

To archive log and configuration data:

1. Stop the Device Server and the GUI Server.
2. Use the `ls -al` command to discover the actual paths of the GUI Server and Device Server data directories. These are the directories you need to back up.

For example:

```
ls -al /usr/netscreen/GuiSvr/var
lrwxrwxrwx 1 root root 21 Feb 25 16:04 /usr/netscreen/GuiSvr var ->
/var/netscreen/GuiSvr
```

The output in the example indicates that the actual location of the GUI Server data is in `/var/netscreen/GuiSvr`. On your own system, verify where your data is stored and which directories should be backed up. Follow the same procedure to determine the location of your data on the Device Server.

3. Run the appropriate backup command on your Solaris or Linux platform to backup the GUI Server data. For example:

```
tar -cvf /netscreen_backup/db-date.tar /var/netscreen/GuiSvr
```

4. Run the appropriate backup command on your Solaris or Linux platform to backup the Device Server data.

For a large amount of log data, using tar may not be appropriate. We recommend using Secure Copy (scp) or File Transfer Protocol (FTP) to backup the Device Server data.

Example using scp:

```
scp -r <local directory> usr@host:<remote-directory>
```

Example using FTP:

```
ftp <host name>  
bi  
hash  
lcd <local directory>  
prompt  
mput
```

5. Start GUI Server and Device Server processes.

Restoring Logs and Configuration Data

These instructions apply only to systems where the var directory links point to a true location outside the prescribed locations (`/usr/netscreen/GuiSvr` or `/usr/netscreen/DevSvr`). We recommend that you do not set these links to point to locations that are inside `/usr/netscreen/GuiSvr` or `/usr/netscreen/DevSvr`; doing so can complicate upgrades to NSM and requires special precautions during backup and restore procedures.

To restore log and configuration data:

1. Stop Device Server and GUI Server processes.
2. Use the **mv** command to transfer data from the var directories to a safe location. This precaution clears the var directory for restoration of the backups.
3. Untar your backups into both of the locations described above.
4. Start GUI Server and the Device Server processes.

Managing Device Schemas Through the Juniper Update Mechanism

You can download and activate XML-based device schemas without upgrading NSM. This feature applies only to devices whose schemas are defined using XML:

- Secure Access devices

- Infranet Controller devices
- Junos devices

This mechanism does not apply to ScreenOS or IDP devices.

The latest device schema is placed by Juniper Networks on the Juniper Update Server, which is a publicly available server. From there, schema upgrade is a two-stage process:

1. Download the schema into the staging repository.
2. Apply the schema.

The most recently downloaded schema is known as the staged schema. The most recently applied schema is known as the current running schema.

The schemas placed on the Juniper Update Server consist of XML statements that define the latest device configuration structure. Overlay schemas are also merged with basic schemas to provide customization of the schemas for specific devices in terms of the configuration branches presented and the manner in which the configuration is presented.

If your version of NSM is not the most current, it might not be compatible with a new schema placed on the Juniper Update server for every device family. In such an instance, the device family is said to be *disabled* for the schema in your version of NSM. For disabled device families, you receive information messages before you download or apply the schema so you can decide whether it is appropriate for you to proceed.

Schema upgrade is split into two stages to allow you to schedule the time at which an upgrade is applied. Applying an upgrade involves restarting the GUI Server and the Device Server, so it is important that this activity is appropriately scheduled. Download can happen at any time.

The system administrator role has all the permissions necessary to manage schemas. Alternatively, you can define a custom role for schema management. Three activities are relevant to defining such a role: View Schema Details, Download Schema, Apply Schema. See [“Configuring Role-Based Administration” on page 63](#) for details on configuring such a role.

Downloading Schemas

You can download the latest schema either through the NSM UI or from the GUI Server CLI. You download the schema directly to the NSM schema staging repository or to an intermediary file. The intermediary file approach is provided for users who do not have Internet access from the GUI Server. Using the CLI, you can set up a cron job to check periodically for new schemas and perform the download. The NSM administrator receives e-mail notification when a new schema has been downloaded.



NOTE: Intermediary files must have 'nsm' user permissions set on them to be downloaded to be retrieved by NSM. To set these permissions, in the NSM server CLI, enter the following command:

```
% chmod 777 filename
```

Access to the Juniper Update server uses your Juniper Networks Download Center credentials—the credentials you use to download software from the www.juniper.net Web site. Use **Tools > Preferences > Juniper Update Settings** to provide these credentials. If your credentials are not present or incorrect, you will not be able to download schemas from the Juniper Update server.

Downloading can be done through a proxy server if your Internet connection is through a proxy server. To configure a proxy server, use **Tools > Preferences** and select **Proxy Settings**.

Downloading Schemas Using the NSM UI

To download a schema using the NSM UI, follow these steps:

1. In the Administer panel, select **Server Manager > Schema Information**.

Schema Information appears in the main display area. It shows information about the schema that is currently staged (the most recently downloaded schema) and the schema that is currently active. Details for each schema version include the version number, the date of last update, a brief description of what changed in that version, and a list of files affected by the change.

The schema information also lists any device families that are disabled in either the staged schema or the running schema, along with the last supported schema version. The display also shows a warning message if the staged schema has disabled families.

2. Click **Download Schema** to look for new schemas to download.

The NSM UI displays a new screen for you to select the source from which to update the schema.

3. From the Update Source list, choose **Juniper Update Server** to get any new schemas directly from the server. Choose **File** to retrieve the schema from an intermediary file.
4. Click **Next** to display information about the latest schema on the source (Juniper Update Server or file) along with current schema information. This information will also indicate if the latest schema on the source contains disabled families. Use this information to determine whether you want to download the schema. The Finish button is activated only if a new schema is available.
5. Click **Finish** to launch a job to download the new schema.

The new schema overwrites the previously staged schema in the NSM repository. When finished, NSM sends an e-mail message to the NSM administrator configured for receiving e-mail notifications.

Downloading Schemas Using the GUI Server CLI

To download a schema using the GUI Server CLI, use the **update-schema** command. If an update is available, this command downloads and stages the update into the NSM repository and notifies the NSM administrator by e-mail of the new download.

The primary advantage of using a CLI command to perform this operation is that you can set up a cron job to periodically look for new schemas and download them.

Downloading a schema using GUI Server CLI will fail if the schema on the Juniper Update Server contains a disabled device family that is not yet disabled in the NSM staged schema.

Selective Schema Loading in NSM

Beginning in NSM Release 2012.2, NSM supports selective schema loading for devices. You can choose this option when you install NSM, or you can set it in the configuration files after the installation or upgrade of NSM. Loading of device schema is based on the parameter **schemaLoader.schemaOption**, which is provided in the configuration file for the GuiSvr (guiSvr.cfg) and DevSvr (devSvr.cfg).

Modify the guiSvr configuration file **guiSvr.cfg** under **/var/netscreen/GuiSvr**. [Figure 80 on page 341](#) and [Figure 81 on page 341](#) show examples of the modified GuiSvr and DevSvr file.

Figure 80: GUI Server Configuration File

```
#Parameter for disabling the Junos Devices Schema Loading
# Options
# 1 --Load All Device Family Schemas, 2 - Load only Screen OS Device Family Schemas, 3 -Load ScreenOs and J/SRX family Schemas.
schemaLoader.schemaOption 1
```

Figure 81: Dev Server Configuration file

```
#Parameter for disabling the Junos Devices Schema Loading
# Options
# 1 --Load All Device Family Schemas, 2 - Load only Screen OS Device Family Schemas, 3 -Load ScreenOs and J/SRX family Schemas.
schemaLoader.schemaOption 1
```

The following schema-loading options are available:

- Load all device family schemas.
- Load Screen OS device schema only (Screen OS).
- Load Screen OS and J/SRX devices schema only (Screen OS + J/SRX Series).

Applying a Schema



NOTE: An applied schema cannot be reverted. The only way to revert to an old schema is through the backup and restore mechanism described in [“Archiving and Restoring” on page 337](#). We recommend that you perform a backup of the GUI Server and the Device Servers before applying a new schema.

To apply a new schema, follow these steps:

1. In the Administer panel, select **Schema Information**.

The Schema Information screen appears in the main display area. It shows information about the schema that is currently staged (the most recently downloaded schema) and the schema that is currently active. Details for each version include a version number, the date of last update, a brief description of what changed in that version, and a list of files affected by the change. Compare the version numbers to tell whether the staged schema is more recent than the currently running schema. Check the information about the schema to determine whether you want to update to the new schema; it might contain only changes that do not apply to your devices. Also, if a more recent schema is staged, then the Apply Schema button is active.

2. Click **Apply Schema** to activate the latest staged schema.

The NSM UI warns that the GUI Server and Device Server will restart as a result of this action, and that all UI users will be logged out.

3. Click **Yes** to proceed with the upgrade.

The GUI Server and Device Server restart. When you log on in the restarted UI, the new schema will be active.

The Job Information screen provides information about the progress of the job, and informs you if any device family is disabled in the new schema.

PART 3

Managing

- [Configuring Objects on page 345](#)
- [Configuring Security Policies on page 473](#)
- [Configuring Voice Policies on page 581](#)
- [Configuring Junos NAT Policies on page 585](#)
- [Configuring VPNs on page 597](#)
- [Central Manager on page 675](#)
- [Topology Manager on page 681](#)
- [Role-based Port Templates on page 691](#)
- [Unified Access Control Manager on page 697](#)

CHAPTER 8

Configuring Objects

Objects represent reusable information, such as network addresses, individual users and user groups, and commonly used configuration data. In Network and Security Manager (NSM), objects are shared between the global domain and all subdomains.

Objects are the building blocks of the NSM management system. You can use an object multiple times in the same domain. For example, you can create an address object to represent a host such as an individual workstation, then use the address object in a VPN protected resource and as the source or destination in a firewall or multicast rule.

This chapter contains the following sections:

- [About Objects on page 346](#)
- [Configuring Address Objects on page 351](#)
- [Configuring Application Objects on page 356](#)
- [Configuring Schedule Objects on page 362](#)
- [Configuring Access Profile Objects on page 363](#)
- [Configuring Quality of Service Profiles on page 363](#)
- [Working with DI Attack Objects on page 365](#)
- [Working with IDP Attack Objects on page 369](#)
- [Configuring Custom DI and IDP Attack Objects on page 371](#)
- [Creating Custom DI Attack Groups on page 395](#)
- [Creating Custom IDP Attack Groups on page 396](#)
- [Configuring Application Identification on page 401](#)
- [Unified Threat Management on page 403](#)
- [Configuring Custom Policy Fields on page 416](#)
- [Configuring GTP Objects on page 418](#)
- [Configuring Service Objects on page 424](#)
- [Configuring SCTP Objects on page 432](#)
- [Configuring Authentication Servers on page 432](#)
- [Configuring User Objects on page 441](#)
- [Configuring VLAN Objects on page 445](#)

- [Configuring IP Pools on page 445](#)
- [Configuring Group Expressions on page 447](#)
- [Configuring Remote Settings on page 450](#)
- [Configuring Routing Instance Objects on page 450](#)
- [Configuring Zone Group Objects on page 451](#)
- [Configuring NAT Objects on page 453](#)
- [Configuring Certificate Authorities on page 461](#)
- [Configuring CRL Objects on page 464](#)
- [Configuring Extranet Policies on page 464](#)
- [Configuring Binary Data Objects on page 465](#)
- [Configuring Protected Resources on page 466](#)
- [Configuring IKE Proposals on page 468](#)
- [Configuring Dial-in Objects on page 470](#)
- [Configuring Border Signaling Gateway Objects on page 471](#)

About Objects

In the NSM UI, most objects appear in the Object Manager; VPN-related objects appear in the VPN Manager. For some object types, such as service objects, attack objects, and IKE proposal objects, predefined objects exist. For most object types, however, you must configure an object before you can use it in your device configuration or security policies.



NOTE: If you import an existing device configuration, NSM automatically imports all objects defined in that configuration.

The Object Manager displays objects created in the current domain only. When you work in the global domain, all custom objects are viewable. When you work in a subdomain, only custom objects created in the subdomains are viewable. However, when creating an object group, you can select objects from both the current subdomain and global domain. Any global object that is part of a subdomain object group appears within the subdomain object list.

Use the Object Manager to view and configure the following objects:

- Host and network addresses:
 - Address objects represent individual hosts or subnetworks in your network.
 - NAT objects (DIP, MIP, VIP) represent references to device-specific NAT configurations (dynamic IPs, mapped IPs, and virtual IPs), enabling multiple devices to share a single object.

- IP Pools define ranges of IP addresses used to assign an IP address to a RAS user.
- Remote Settings represent DNS and WINS servers.
- Services and schedules:
 - Schedule objects represent time periods and determine when a rule is in effect.
 - Service objects represent predefined and custom network services, such as HTTP/80.
 - Quality of Service profiles (IP Precedence and DSCP profiles) determine the quality of service for an incoming packet in the network.
 - Shared Stream Control Transmission Protocol (SCTP) objects support protocols such as IUA, SUA, M2UA, M3UA, and so on, and can be applied to policies.
 - Border Signaling Gateway (BSG) objects allow you to specify the egress-service-point in a transaction term's route action. You can specify the device, the gateway in the device, and a service point for every BSG service point object.
 - BSG Admission Controllers control SIP dialogs and transactions. They are defined per gateway and referenced from transaction policies.
- Application layer protection:
 - DI Profiles define the attack signature patterns, protocol anomalies, and the action you want a security device to take against matching traffic.
 - AV Profiles define the server that contains your virus definitions and antivirus software.
 - Web Filtering Profiles define the URLs, the Web categories, and the action you want a security device to take against matching traffic.
- Users and authentication:
 - User objects represent RAS users on your network.
 - Authentication Servers represent the servers in your network used to authenticate NSM administrators, RAS users, and network traffic.
 - Group Expressions define logical expressions used to include or exclude RAS users.
- Certificates:
 - Certificate Authority objects represent the certificate authority's certificate.
 - CRL objects represent the certificate authority's certificate revocation list.
- VoIP protection:
 - GTP objects represent client GTP configurations.
- Extranet policies and custom policy fields:

- Extranet Policy objects define rules and actions that you may apply to certain traffic on an extranet device (third-party router).
- Custom Policy Field objects represent metadata information that you can store and use in a structured manner.

Use VPN Manager to view and configure the following objects:

- Protected Resources represent the network components, a network service, and the security device that protects those components and service.
- IKE Phase1 Proposals represent the phase1 proposals used to establish a secure and authenticated communication channel between two VPN members.
- IKE Phase2 Proposals represent the Security Associations for services (such as IPSec) that require key material or parameters, as exchanged by two VPN members.

Using Objects Across Domains

Objects created in the global domain are available in all subdomains, but objects created in a subdomain are available only in that subdomain.

For example, when creating a VPN:

- You can use a global domain user object in a subdomain VPN.
- You can use a subdomain user object in a subdomain VPN.
- You cannot create VPNs across domains. However, you can use an extranet device to represent the device in the other domain to create a cross-domain VPN.
- You cannot use a subdomain user object in a global domain VPN.

When creating a subdomain protected resource, you can include a subdomain address object and a global domain service object, but you can only select the protected resource when you are logged in to that specific subdomain.

Replacing Objects

You can use Replace With operations to simplify the process of making repeated changes to an object that is referenced in multiple security policies.

The following shared objects support Replace With operations:

- Address Objects
- Service Objects
- Zone Objects
- Routing Instance Objects

To replace an object with another shared object:

1. From the navigation tree, select **Object Manager**.
2. Right-click on the object that you want to replace and select **Replace With** from the menu.

All available objects of the same category from the global domain are displayed, except the selected object that you are replacing.

3. Select an object that will replace all instances of the existing object and click **Next**.

4. Click **Finish**.

The selected object replaces all instances of the original object in your current working domain.

5. Check your security policies for any errors that might result. You can always edit or remove any duplicate objects in the security policy.

Keep the following limitations in mind when replacing objects:

- NSM provides no validation checking when replacing address objects.
- Replace With operations cannot be rolled back.

Working with Unused Shared Objects

Searching for Unused Shared Objects

You can locate unused shared objects in NSM based on selective nodes (address, service, and so on). This feature is available through the NSM GUI.

To locate an unused shared object:

1. Search for the unused object.
 - From the menu bar at the top, select **View > Search Unused Objects**

The Search Unused Objects window appears. By default, the all search categories are preselected. You can narrow the search by unchecking unnecessary categories.

 - Right-click on a shared object node (for example, Address Objects) and select **Search Unused Objects**.
2. Select the search categories and click **Next**.

The Unused Shared Object List appears.

Deleting an Unused Shared Object

You can delete unused shared objects in NSM using the NSM GUI.

To delete a located unused shared object:

1. From the Unused Search Object List, select the shared objects to be deleted and click **Next**.

An Object Deletion Warning appears.

2. Click **Finish** to delete the objects.

The selected objects are deleted and do not appear in the NSM GUI windows.

Working with Object Versions

You can use the NSM GUI to work with multiple versions of NSM objects. You can create a new object version (for example, a database version), search for existing versions with and without filters, edit comments about versions, compare two versions, restore an older version, filter and sort versions, display the differences between versions, and update a device to an older object version.

These operations are described in the [“Automatic Policy Versioning” on page 566](#).

Searching For and Deleting Duplicate Objects

When you create a new object from Object Manager, NSM displays a warning if a similar custom object with same parameters already exists. You can use Object Manager to search for and delete unused duplicate objects from a selected category.

To find and delete a duplicate object from a selected category:

1. From Object Manager, right-click on an object category, for example, **Address Objects**, **Attack Objects** or **Service Objects**, and select **Search Duplicate Objects** from the menu.

The **Search Duplicate Objects** dialog box is displayed.

2. From the Shared Object Category List, select one or more categories that you want to search in the current domain for duplicate objects and click **Next**.

NSM displays the list of Unused Duplicate Objects.

3. From the list of Unused Duplicate Objects, select the objects you want to delete.

NSM displays a message that the selected objects will be deleted and a warning that the operation cannot be reversed.



NOTE: When you select a group of duplicate objects, such as an address group, NSM displays the details of the fields that are common.

4. Click **Next** to delete the selected objects.

NSM deletes the unused duplicate objects that you selected and displays a report of the deleted objects.

5. Click **Finish** to exit.

Configuring Address Objects

An address object is a representation of a component of your network, such as a workstation, router, switch, subnetwork, or any other object that is connected to your network. You use address objects in NSM to specify the network components you want to protect:

- Firewall and IDP Rules—Use address objects or groups to specify the source and destination of network traffic
- Multicast Rules—Use multicast group address objects to specify the destination of multicast traffic.
- VPNs—Use address objects or groups to create Protected Resources for your Policy-Based and Mixed-Mode VPNs.

Viewing Address Objects

In the navigation tree, click **Object Manager > Address Objects** to view all address objects for the current domain. You can display Address objects in a tree or table format:

- The Address Tree tab displays address objects in a tree format. To view the members of an address object group, click the group to display a member list.
- The Address Table tab displays address objects in a table format with the following columns:
 - Name—Name of the address object
 - Type—Type of the address object (Host, Network, Group)
 - IP/Domain Name—The IP address or host name (such as www.juniper.net) of the address object
 - Netmask—Netmask of the address object
 - Comment—A description of the address object

When you initially deploy the NSM system and open the UI for the first time, the Address Object tree and table tabs are empty. Using the Object Manager, you can create address objects that represent network components that are unique to your network. As you add address objects, they appear in the tree and table tabs.

Creating Address Objects

You can create the following address objects:

- **Host**—Represents devices, such as workstations or servers, connected to your network.
- **Network**—Represents divisions or subnetworks in your network.
- **Address Object Group**—Represents multiple address objects.
- **Multicast Group**—Represents the destination of multicast packets.

NSM supports the IPv6 protocol in configuring policy rule bases, IDP, address, and attack objects for devices running ScreenOS 6.3 and later versions, and Junos OS Release 10.2 and later versions.

The following restrictions apply when using IPv6 while configuring objects:

- Address groups cannot contain both IPv4 and IPv6 addresses.
- The wildcard mask option is not supported for IPv6 addresses.
- Rules cannot contain a combination of IPv4 and IPv6 addresses.
- IPv4 addresses cannot be copied to IPv6 rules and vice versa.
- During a device update, an IPv6 policy rule is dropped if the target platform does not support IPv6.

The following sections detail each address object type.

Adding a Host Address Object

To add a host address object:

1. In the navigation tree, select **Object Manager > Address Objects** to open the address object tree. In the main display area, click the Add icon and select **Host**.
2. Enter a unique name for the address object and select a color to represent the address object.
3. Enter a comment about the host (optional).
4. Select the IP version of the address object: IPv4 or IPv6.
5. Enter the address that identifies the host on your network:
 - To identify the host with an IP address, select **IP** and enter the IP address of the host. Click **Resolve** to automatically resolve the domain name for that IP address.
 - To identify the host with a domain name, select **Domain Name** and enter the domain name of the host. Click **Resolve** to automatically resolve the IP address for that domain name.



NOTE: When NSM fails to resolve a name for an IPv6 address, it displays the same address under the domain name. This is an indication that a name is not configured for this address.

6. Click **OK** to add the address object.

The new host address object immediately appears in the Address Tree and Address Table.

Adding a Network Address Object

To add a network address object:

1. In the navigation tree, select **Object Manager > Address Objects** to open the address object tree.
2. In the main display area, click the Add icon and select **Network**.
3. Enter a name for the address object.
4. Select the IP version of the address object: IPv4 or IPv6.
5. Enter the IP address and netmask of the network.
6. Select a color to represent the address object.
7. Enter a comment about the network, then click **OK** to add the address object.

The new network address object immediately appears in the Address Tree and Address Table.

NSM supports the wildcard masking feature policy on all devices running ScreenOS 6.1 and later, except those with IPv6 addresses.

To configure a wildcard mask, follow the procedure for adding a network address object. In the **Network** dialog box, select **Use Wildcard Mask**. You can then add or edit the **Wildcard Mask** field. If you are configuring wildcard masking on a new device, verify that the device update and Delta Config Summary operations are successful.



NOTE: When a firewall policy with network address objects is applied to Junos devices, the device update operation in NSM fails, because DMI devices do not support network address objects.

Editing and Deleting Address Objects

To edit an address object, right-click on the object and select **Edit**. To delete an address object, right-click on the object and select **Delete**. For more information on editing and deleting address objects, refer to the NSM Online Help.

Replacing Address Objects

To replace an address object, right-click on the object to be replaced and select **Replace With**. Replacing address objects simplifies making repeated changes to an address object that is referenced in multiple security policies. If you have permission to view global domain objects for the objects you are replacing, then all objects for the selected category in the current domain and the global domain are displayed in the Replace With wizard, but the object to be replaced is not displayed. When you replace address objects, keep the following in mind:

- There is no validation check when replacing address objects.

- You cannot undo or roll back a Replace With operation.



NOTE: Replacing address objects only affects objects in your current working domain.

After replacing address objects, it is good practice to check your security policies for any errors that may result. You can always edit or remove any duplicate objects in the security policy.

Adding an Address Object Group

To simplify security policies, you can combine multiple address objects in an address object group. An address object group can contain address objects (and other address object groups) from the current subdomain and the global domain.

To add an Address Object Group:

1. In the navigation tree, select **Address Objects**. The address object tree appears.
2. In the main display area, click the Add icon and select **Group**.
3. Enter a unique name for the group.



NOTE: Address object group names must be unique; you cannot give an address object group the same name as an existing address object.

4. Select a color to represent the group.
5. Enter a comment about the group.
6. Select the IP version. The member list is populated with a particular IP version, based on your choice. Each address group is exclusive to the selected IP version. You can only add IPv4 addresses to an IPv4 address group; similarly, you can only add IPv6 addresses to an IPv6 address group.
7. In the Non-members list, select the address objects you want to add to the group (hold Ctrl to select multiple objects), then click **Add**. The selected address objects now appear in the member list.
 - If you are in the global domain, only the global address objects appear in the Non-members list.
 - If you are in a subdomain, both global and subdomain address objects appear in the Non-members list.



NOTE: You can drag address objects into and out of address groups from the main address tree.

8. Click **OK** to add the group.

You can create address object groups with existing users or create empty address object groups and fill them with users later.

Adding a Multicast Group Address Object

To add a multicast group address object:

1. In the navigation tree, select **Address Objects**. The address object tree appears.
2. In the main display area, click the Add icon and select **Multicast Group**. The New Multicast Group dialog box appears.
3. Enter a name for the multicast group address.



NOTE: Multicast Group address object names must be unique; you cannot assign the same name to another existing multicast group address object.

4. Select a color to represent the multicast group address.
5. Enter a comment about the multicast group address.
6. Select an IP version: IPv4 or IPv6.
7. Enter the IP address of the multicast group. All IPv6 multicast addresses should have a prefix of **ff00::/8**. The netmask field is unavailable for IPv6 multicast addresses since the value is fixed and set to 8. NSM validates your entries and prompts a correction in case of an error.
8. Click **OK** to add the multicast group address.

Adding Static DNS Host Addresses

This ScreenOS 5.3 or later feature lets you create a static host name with multiple IP addresses. You can use this feature to create dynamic addressing in NSM.

To add multiple static host addresses:

1. In the navigation tree,
2. Double-click the device you want to configure. The device must be running ScreenOS 5.3 or later.
3. In the navigation tree of the new dialog box, select **Network > DNS**.
4. Click **Settings** to open the Device Settings dialog box.
5. Click the Add icon, enter the host name and host IP address, then click **OK**.
6. Click **OK** to save the changes and close the dialog box.

Example: Using Static Addresses to Share a FW Policy

Static addresses allow two sites with different IP addresses to share a single firewall policy. For example, each site might have a Web server, each with a different IP address. If you define an address object using the hostname “webserver” and then using that

object in the firewall policy, the device will resolve the address object's hostname to the correct IP for that device as defined by its static host entry.

1. In the navigation tree, select **Object Manager > Address Objects**.
2. Click the Add icon, then select **Host** to open the New Host dialog box.
3. Enter the same name in the Name field that you entered for the Device host name in the previous section. These values are case sensitive and must match exactly.
4. Click **OK** to save the name and close the dialog box.
5. Return to the navigation tree and select **Security Policy**.
6. Click the **Add** icon and enter the security policy name, then click **OK**.
7. Double-click the name of the security policy you just created.
8. Right-click the value in the Source column or the Destination column.
9. Select the address object you just created, click **Add**, then click **OK**. When the address object is pushed to a device, the host name resolves dynamically. One policy can be assigned to multiple devices.



NOTE: If an address object is used in multiple zones, NSM pushes the address object into the zones without changing its name. When you import a device, NSM combines address objects with the same name and same content from different zones into a single address object.

Blocked Hosts

NSM can block the IP address of hosts where login attempts fail consecutively for a specified number of times. The default value is 5 times. NSM saves a list of these blocked IP addresses. Select **Tools > Managed Blocked Hosts** to display a list of blocked hosts or to clear the blocked IP addresses. If the local host is blocked, you must use another computer to use this option to unblock the host IP address.

Configuring Application Objects

In the **Application Objects** window (under **Object Manager**), you can:

- View the predefined application objects and the predefined extended application objects.
- View, search, create, edit, and delete the custom application objects and the application group objects.

Viewing Predefined Application Objects

The **Predefined Application Objects** tab in the **Application Objects** window lists all the predefined application objects in NSM. You can view the predefined Application objects in a table format with the following details.

Table 32: Predefined Application Table Tab Information

Field	Description
Name	The name of the application object.
Application Category	The hierarchical category to which the application belongs.
Port Range	The TCP/UDP port ranges to be matched with application signatures. Specifying a small range improves system performance. You must configure either a TCP or UDP field; while optionally, you can configure both.
Application Type	The type of application—predefined or custom type. Port Binding is required for a custom type application while it is not required for a predefined type.
Port Binding	The default TCP/UDP port bindings required for custom application types only. You must configure either TCP or UDP or optionally, both of them.
Match Order	An integer value used to resolve conflict when multiple application signatures are matched for a session. In that case, the application signature with the highest order (smallest value) is taken. It is assumed that no two signatures have the same order value, in which case the first application signature IDP sees (not necessarily the first one in policy) is taken.

You can right-click on an application object and select **Find Usages** to view all the places where this object is referenced.

You can double-click on an application object to view its settings which include the following additional information:

- Supported Platforms
- Application signature definition (including Client-to-Server DFA and PCRE patterns and Server-to-Client DFA and PCRE patterns)
- Minimum data length which is the minimum number of layer-7 data bytes that the first data packet requires to make a successful match. This applies to both Client-to-Server and Server-to-Client packets.

Viewing Predefined Extended Application Objects

The **Predefined Extended Application Objects** tab in the **Application Objects** window lists all the predefined extended application objects in NSM. You can view the predefined extended application objects in a table format with the following details.

Table 33: Predefined Extended Application Table Tab Information

Field	Description
Name	The name of the application object.
Application Category	The hierarchical category to which the application belongs.
Ext ID	A unique identifier. The system uses the unique ID both for logical processing and reporting.

Table 33: Predefined Extended Application Table Tab Information (continued)

Field	Description
Application Type	The type of application: predefined or custom type. Port Binding is required for a custom type application, but not required for a predefined type.
L7 Protocol	Only HTTP layer 7 is supported.
Chain Order	Indicates whether member signatures are ordered or not.

You can right-click on an application object and select **Find Usages** to view all the places where this object is referenced.

You can double-click on an application object to view its settings, including the following additional information:

- **Maximum Transactions**—Maximum number of transactions the device must inspect for each application signature. Once the maximum is reached the inspection stops. For HTTP, a transaction is a complete request-response cycle.
- **Signature Match Order**—An integer value used to resolve conflict when multiple application signatures are matched for a session. In this case, the application signature with the highest order (smallest value) is taken. It is assumed that no two signatures have the same order value, in which case the first application signature that IDP sees (not necessarily the first one in the policy) is taken.
- **Members**—Members are applicable to compound extended application objects (if any).

Only standalone IDP sensors running IDP 5.1 and later support this feature. When you select this action in an APE rule installed on a device running IDP 5.0 or earlier, NSM displays a message warning the user that if unsupported, this APE action might cause a device update failure.

Creating Custom Application Objects

You can create, edit, delete, and search for customized application objects under the **Custom Application Objects** tab in the **Application Objects** window. You can create custom application objects to represent applications that are not predefined. To add a custom application object, in the **Object Manager**:

1. Select **Application Objects > Custom Application Objects**.
2. Click the Add icon (+) to view the **New Custom Application** dialog box.
3. Configure the following parameters in the **General** tab:
 - **Name**—This is a mandatory field.
 - **Application Category**—This is a mandatory field.

- **Supported Platforms**—Use the **Edit** icon to select supported platforms. You must select at least one.
 - **Port Ranges**—Either a TCP or UDP port is mandatory. You can specify hyphen (-) separated port ranges of both types.
4. Configure the following parameters in the **Detector** Tab:
- **Port Binding:**
 - **Application Type**—Select a predefined or custom application type from the drop-down list. This is a mandatory field.
 - **TCP Port Binding**—Specify comma separated ports. A range of ports is not allowed. You must configure either one TCP or UDP port for a custom application.
 - **UDP Port Binding**—Specify comma separated ports. A range of ports is not allowed. You must configure either one TCP or UDP port for a custom application.
 - **Signature:** Specify a DFA and a PCRE pattern under each of the following sections:
 - **Client-to-server**
 - **Server-to-client**

You must specify at least one DFA pattern.
 - **Minimum Data Length:** This is a mandatory field.
 - **Signature Match Order:** This is a mandatory field.
5. Click **OK** to create the application object. Otherwise, click **Cancel**.

Creating Application Group Objects

You can create, edit, delete, and search for application group objects under the **Application Group Objects** tab in the **Application Objects** window.

In releases prior to NSM 2010.4, in order to target multiple applications in an APE policy, you had to add them individually to the policy. An application group simplifies this task as it enables you to group multiple applications (predefined, custom, or both) into a single object, which can then be added to the policy in a single operation.

To create an application group object:

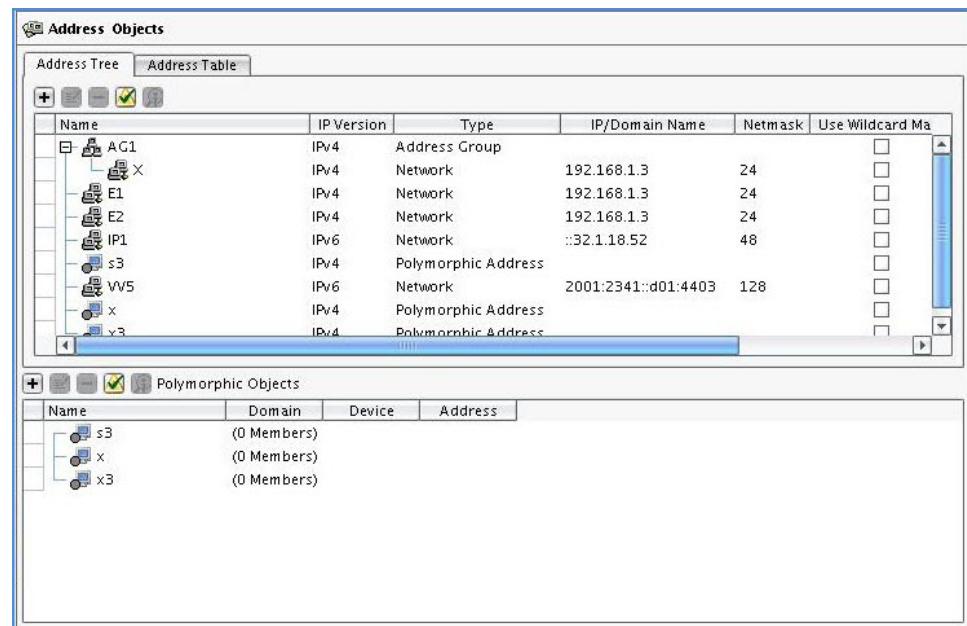
1. Select **Application Objects > Application Group Objects**.
 2. Click the Add icon (+) to view the **New Application Group** dialog box.
 3. Enter a name to uniquely identify the application group object.
 4. Select the predefined application objects, custom application objects, or both from the list.
 5. Click **OK** to create the application group object. Otherwise, click **Cancel**.
- The newly added application group object is displayed in the **Application Group Objects** tab.

Validating Address Objects

Address objects are loaded in the address tree and address table on the NSM client after validating errors and warning messages. Validating all the addresses while the address objects are loading takes a very long time and adversely affects NSM performance. To improve performance, disable validation of address objects during loading by using the option **Disable GUI validation>Shared objects** under **Tools>Preference**. For more information about disabling GUI validation, see [“Disabling GUI Validation” on page 24](#).

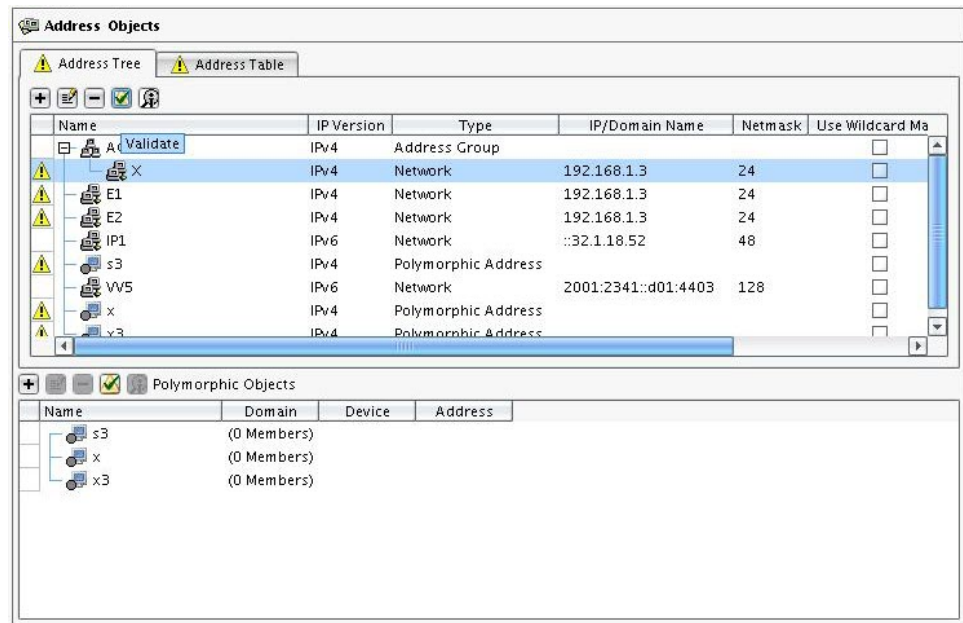
To validate addresses after all the address objects are loaded, click **Validate** at **Address Manager>Address Objects**. [Figure 82 on page 360](#) shows the shared object validation option. When you click **Validate** on the tool bar, all addresses are validated.

Figure 82: Shared Address Object Validation Option



[Figure 83 on page 361](#) shows the warning and error messages for address objects. Both the tree and the table views display validation errors for address objects.

Figure 83: Shared Address Object Validation Warning



NOTE: Validation of address objects are also enhanced on service and custom attack objects.

Editing and Deleting Application Objects

Right-click on an application object and select **Edit** in order to edit it. Similarly, you can right-click on an application object and select **Delete** to delete it. You can only delete a custom application object or an application group object and not a predefined application object. However, before performing an edit or a delete operation, you may want to confirm whether this object is referenced elsewhere, such as within a policy. To verify this:

1. Right-click on an application object.
2. Select **Find Usages**. The **Find References to Application Object <application object name>** dialog box displays.

All referenced areas are displayed as links in this dialog box. Click on a link to navigate to the area where the object is referenced.

If you have performed a number of edits to an application object and want to revert to a specific previous version:

1. Right-click on an application object.
2. Select **View Versions**. The **Version History for appsig** dialog box displays.

3. Select a version and click **Restore**.

The application object is restored to the selected version.

Configuring Schedule Objects

A schedule object defines a time interval that a firewall rule is in effect. You use a schedule object in your firewall rule to determine when a device enforces that rule:

- Use a one-time schedule to control access to a destination for a specific time interval. The schedule object defines a start time, end time, and date during which a rule is enforced. Some examples:
 - Contractor Access Schedule (8:30 AM December 1 to 6:00 PM December 5)
 - Christmas Break Schedule (6:00 PM December 24 to 8:00 AM January 2)
- Use a recurring schedule to control access to a destination for a repeating time interval. The schedule object defines a start time, end time, and days during which a rule is enforced. Some examples:
 - Business Hours Schedule (8:00 AM to 6:00 PM on Monday, Tuesday, Wednesday, Thursday, Friday)
 - After Hours Schedule (6:01 PM to 7:59 AM on Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday)
 - Weekend Schedule (8:00 AM to 6:00 PM on Saturday, Sunday)
- Combine a one-time and recurrent schedule to define a repeated time interval.

Creating Schedule Objects

To add a schedule object:

1. In the NSM GUI navigation tree, **Schedule Objects**. The schedule object tree appears.
2. In the main display area, click the Add icon.
3. Enter a name and comment for the schedule object.
4. Select the frequency of the schedule:
 - To configure a one-time schedule, select **Once**, and enter the start date, start time, stop date, and stop time.
 - To configure a recurrent schedule, select **Recurrent**, and click the Add icon. In the Recurrent Schedule dialog box, select the day of the week and specify the hour and minutes for Start 1 and Stop 1.

To specify a second recurring time interval on the same day, specify the hour and minutes for Start 2 and Stop 2. For example, Business Hours Schedule (8:00 to 12:00 and 13:00 to 17:00 every weekday).

Configuring Access Profile Objects

An access profile consists of a set of attributes that defines access to a device. You can configure multiple profiles and multiple clients for each profile. You can use Object Manager to configure access profile objects. Access profile objects can then be shared across security policies that are assigned to J Series Services Routers and SRX Series Services Gateways managed by NSM. NSM does not provide predefined access profile objects, and you must create an access profile object before you can use it in security policies.

To view all objects, select **Access Profiles** in the navigation tree. The Object Manager displays all the access profiles configured in NSM. Access profiles are listed in a table consisting of the following columns:

- Name—Name of the access profile.
- Comment—Description of the access profile.

You can create, view, edit, or delete access profile objects in the Object Manager. You can also perform a Find Usages operation, and view the version history of an access profile. For details on creating an access profile, see the *Junos System Basics Configuration Guide*.



NOTE: The access profile object is not supported on ScreenOS devices. When a security policy using an access profile is assigned to a ScreenOS device, the access profile settings will be removed before the security policy is updated.

Configuring Quality of Service Profiles

On SSG Series Secure Services Gateways running ScreenOS 6.3 and later, you can define Quality of Service (QoS) profiles as objects under the Object Manager. These profiles can be either IP precedence profiles or Differentiated Services Code Point (DSCP) profiles. For a DSCP QoS profile, the six most significant bits in the ToS field of the IP packet are used to match entries; for an IP precedence QoS profile, only the three most significant bits in the ToS field are used.

You can create, edit, delete and search for a QoS profile.

- [Creating a Quality of Service Profile on page 364](#)
- [Deleting a Quality of Service Profile on page 364](#)
- [Editing a Quality of Service Profile on page 365](#)

Creating a Quality of Service Profile

1. Select **Object Manager > QoS Profiles**. The **QoS Profile** window opens. You can add, edit, delete or search for a QoS profile using the icons at the top of the window.
2. Select **+** to create a new profile. A new QoS profile window appears.
3. Enter a name for the profile in the **Name** field. The name must be unique and between 1 to 31 characters long.
4. Select the QoS profile type from the drop-down list. Your options are: DSCP and IP. The **New/Edit QoS Profile** dialog box opens. In it you can set values for the following fields:
 - **Entry index:** Create or remove an entry based on the DSCP/IP value. Each profile can have multiple entries. Each entry represents a mapping between a special DSCP/IP value and its intended QoS parameters. Each IP profile can have a maximum of 8 entries, and each DSCP profile can have 64 entries. In a QoS profile, an existing entry can be overwritten with the same DSCP/IP Precedence value. If you do not specify values, the values for each entry are set to the lowest by default.
 - **Packet priority:** Set the priority of the packet when sent out. Select 0-63 for DSCP and 0-7 for IP Precedence. Zero has the highest priority.
 - **Policing bandwidth in kbps:** 0-1000000 (in Kbit per sec). The default setting is 0.
 - **Maximum bandwidth in kbps:** 0-1000000 (in Kbit per sec). The default setting is 0.
 - **Guaranteed bandwidth in kbps:** 0-1000000 (in Kbit per sec). The default setting is 0.
5. Click **OK**.

After creating a QoS profile, you can add it to a policy. You cannot, however, delete a QoS profile after it has been added to a policy.



NOTE: QoS profiles cannot co-exist with traffic shaping in the same policy.

Deleting a Quality of Service Profile

1. Select **Object Manager > QoS profiles**. The QoS profile screen opens with a list of QoS profiles.
2. Select a QoS profile to delete.

3. Select the Delete icon (-) at the top of the screen. The **Delete QoS Profile** window opens.
4. Click **OK** to delete the profile.

Editing a Quality of Service Profile

1. Select **Object Manager > QoS profiles**. The QoS profile screen opens with a list of QoS profiles.
2. Select a QoS profile to edit.
3. Select the Edit icon at the top of the screen. The **Edit QoS Profile** window opens.
4. Edit the values of the profile.
5. Click **OK**.

Working with DI Attack Objects

Deep Inspection (DI) attack objects contain attack patterns and protocol anomalies for known attacks and unknown attacks that attackers can use to compromise your network. DI attack objects must be part of an attack object group, and a DI Profile object before you can use them in a firewall rule to prevent malicious traffic from entering your network.



NOTE: Deep Inspection is supported by NS-5GT devices, the NS-HSC, and all devices running ScreenOS 5.3 or later.

To create a Deep Inspection (DI) Profile object, you add predefined attack object groups (created by Juniper Networks) and your own custom attack object groups to the Profile object. After creating the DI Profile, you add the Profile object in the Rule Option column of a firewall rule. If an attack is detected, the device generates an attack log entry that appears in the Log Viewer.

For information about configuring Deep Inspection in a firewall rule, see [“Creating DI Profiles” on page 367](#).

Viewing Predefined DI Attack Objects

NSM contains a database of hundreds of predefined DI attack objects designed to protect networks from multiple attack vectors. Predefined groups contain attack objects, which you can use in a DI Profile to match traffic against known and unknown attacks.



NOTE: NSM displays a superset of all predefined DI attack objects. Based on the platform and ScreenOS firmware version, security devices include a specific subset of DI attack objects. Therefore, the list of predefined DI attack objects displayed in the NSM UI might not match the list of predefined DI attack objects on the physical security device.

To view individual predefined attack objects, select **Attack**. The **Predefined Attacks** tab (default view) displays a table of predefined attack objects that represent known and unknown attack patterns. Use the **Predefined Attacks** tab to quickly view details about an attack object, such as name of the attack object, attack severity, attack category, and attack references. To view the properties for an attack, right-click the attack and select **View**.

To locate all firewall rules that use a predefined attack object or group, right-click the attack object and select **Find Usages**.

Viewing Attack Version Information for Attack Objects

You can view details for predefined attack objects; however, not all details are applicable to all attacks.

The **Pattern** field under the **Details** column in the **General** tab of the attack object dialog box (which appears when you double-click the object) displays the regular expression used to identify the attack. Juniper Networks Security Engineering might choose to hide the exact pattern for specific attack objects. This is done to protect the confidentiality of either the source or target of the specific attack object. In such cases, the field displays Protected instead of the regular expression.

To view attack version information, click one of the supported platform links under the **Platform** column within this attack object dialog box.

Viewing Predefined DI Attack Object Groups

To view predefined attack object groups, in **Object Manager**, select **Attack Objects**, then select the **Predefined Attack Groups** tab. The name of each attack object group indicates the severity, protocol, and attack type of the individual attack objects contained within. For example, the predefined attack object group CRITICAL:DNS:ANOMALY contains predefined protocol anomaly attack objects that detect critical Domain Name Service (DNS) attacks.

To locate all firewall rules that use a predefined attack object or group, right-click the attack object group and select **Find Usages**.

Updating Predefined DI Attack Objects and Groups

You cannot create, edit, or delete predefined DI attack objects or groups, but you can update the attack object database with new attack objects created by Juniper Networks. Updates can include:

- New descriptions or severities for existing attack objects

- New attack objects
- Deletion of obsolete attack objects

Creating DI Profiles

A Deep Inspection (DI) Profile object contains predefined attack object groups (created by Juniper Networks), and your own custom attack object groups. After creating the DI Profile, you add the Profile object in the Rule Option column of a firewall rule.

To create a DI Profile:

1. In the navigation tree, select **Object Manager > Attack Objects > DI Objects**.
2. Select the **Profile** tab.
3. Click the Add icon to add a new Profile object.
4. Configure the name, color, and comments for the profile object.

To add members to the profile object, configure the following:

- **DI Severity**—Select a DI Severity setting for the profile object. The DI Severity setting overrides the severity setting of the attack objects included in each profile member.
- **Signature Category**—Select a category of DI signatures. You can only select categories for which you have a license.

Categories are as follows:

- Server Protection Pack—Designed to protect servers.
- Client Protection Pack—Designed to protect remote and home offices.
- Worm Mitigation Pack—Designed to protect against worms.
- Base (Default) Pack—All signatures. Might be too large for some devices.
- DI Attack Objects and Groups—Add a profile member to the profile object. Each profile member can contain attack object groups, and you can add multiple profile members to the profile object. Within each profile member:
 - Select the attack object groups you want to include in this profile member.
 - Configure the action you want the security device to take when an attack object within the profile member matches traffic. [Table 34 on page 367](#) lists DI profile actions.

Table 34: Deep Inspection Profile Actions

Action	Description
None	The security device takes no action against the connection.
Ignore	The security device ignores the remainder of a connection after an attack object is matched.

Table 34: Deep Inspection Profile Actions (continued)

Action	Description
Drop Packet	<p>The security device drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source IP address.</p> <p>For TCP connections, dropping a single packet will result in the same packet being resent. So, Drop Packet settings are translated to Drop Connection settings for TCP connections.</p>
Drop Connection	The security device drops the connection without sending a RST packet to the sender, preventing the traffic from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
Close Client and Server	The security device closes the connection and sends a RST packet to both the client and the server.
Close Client	The security device closes the connection to the client but not to the server.
Close Server	The security device closes the connection to the server but not to the client.



NOTE: Network security is an ongoing process of defining normal traffic for your network. Eliminating malicious traffic is important, but identifying ambiguous traffic can be equally important. You do not always want to drop traffic that appears abnormal; you might want to reset the connection, block the attacker, set an alert for the event, or all three.

- Configure Deep Inspection Alerts. Enable this option to create an event log entry for matching traffic. If the security device matches network traffic to an attack object in the rule, NSM creates an event log entry that describes that attack (direction, service, and Attack object) and displays an alert in the **Log Viewer**.
- Configure IP Action. Enable this option to direct the device to take action against a brute force attack. When enabled, configure the following IP controls action:
 - **Action.** Select the action you want the device to take when it detects a brute force attack. [Table 35 on page 368](#) lists DI IP actions.

Table 35: Deep Inspection IP Actions

Action	Description
IP Block	The security device logs the event and drops all further traffic matching the target definition for the period of time specified in the timeout setting.

Table 35: Deep Inspection IP Actions (continued)

Action	Description
IP Close	The security device logs the event and drops all further traffic matching the target definition for the period of time specified in the timeout setting and sends a Reset (RST) for TCP traffic to the source and destination addresses.
IP Notify	The security device logs the event but does not take any action against further traffic matching the target definition for the period of time specified in the timeout setting.

- **Target.** Specifies a set of elements that must match for the security device to consider a packet part of a brute force attack. The specified set of elements in an IP packet arriving during a specified timeout period must match that in the packet that the security detected as part of a brute force attack for the subsequent packet to be considered part of the same attack. Possible values are Source, Destination, Destination Port, and Protocol; Source; Destination; From Zone, Destination, Destination Port, and Protocol; and From Zone.
- **Timeout (sec).** A period of time following brute force attack detection during which the security device performs an IP action on packets matching specified target parameters. The default is 60 seconds.

After you have created the DI Profile object, you can use the object in your firewall rules.

Working with IDP Attack Objects

NSM contains a database of predefined IDP attack objects and IDP attack object groups that you can use in security policies to match traffic against known and unknown attacks. Juniper Networks updates the predefined attack objects and groups on a regular basis with newly-discovered attack patterns.

Viewing Predefined IDP Attacks

The **Predefined Attacks** tab displays all attacks in a table format and includes the following information:

- Name of the attack object
- Severity of the attack: **Critical**, **Major**, **Minor**, **Warning**, or **Info**
- Category, displaying the type of application
- Keywords for the attack
- CVE number, identifying the number of the attack in the Common Vulnerabilities and Exposures database
- Bugtraq number, identifying the equivalent attack in the Security Focus Bugtraq database

By default, attack objects are listed alphabetically by Category name. To view attacks in a different order, click on a column heading. To display a detailed description of an attack object, you can do one of the following:

- Double-click the attack.
- Right-click the attack object and select **View** to display the attack viewer.

Viewing Predefined IDP Attack Groups

The **Predefined Attack Groups** tab displays the following predefined attack groups:

- **All Attacks**—A list of all attack objects, organized in the categories described below.
- **Recommended Attacks**—A list of all attack object objects that Juniper Networks considers to be serious threats, organized into categories.
- **Attack Type** groups attack objects by type (anomaly or signature). Within each type, attack objects are grouped by severity.
- **Category** groups attack objects by predefined categories. You can view the IDP version of predefined IDP attack groups. Within each category, attack objects are grouped by severity.
- **Operating System** groups attack objects by the operating system to which they apply: UNIX or Windows. Within each operating system, attack objects are grouped by services and severity.
- **Severity** groups attack objects by the severity assigned to the attack. IDP has five severity levels: **Info**, **Warning**, **Minor**, **Major**, and **Critical**. Within each severity, attack objects are grouped by category.

To locate all rules that use a predefined attack object group, right-click the attack object group and select **Find Usages**.

A predefined static group can include the following members:

- Predefined attack objects
- Predefined static groups
- Predefined dynamic groups

To display a detailed description of an attack object group, right-click the attack and select **View**.

Viewing Attack Version Information for Attack Objects and Groups

NSM lets you look at the details of predefined attack objects and groups. Not all details are applicable to all attacks.

The **Pattern** field under the **Details** column in the **General** tab of the attack object dialog box (which appears when you double-click the object) displays the regular expression used to identify the attack. Juniper Networks Security Engineering may choose to hide

the exact pattern for specific attack objects. This is done to protect the confidentiality of either the source or target of the specific attack object. In such cases, the field displays Protected instead of the regular expression.

To view attack version information, click one of the supported platform links under the **Platform** column within this attack object dialog box.

Updating Predefined IDP Attack Objects and Groups

Juniper Networks updates the predefined attack objects and groups on a regular basis with newly-discovered attack patterns. You can update the attack object database on your security devices by downloading the new attacks and groups to the NSM GUI Server, then installing the new database on your devices.



NOTE: You cannot create, edit, or delete predefined attack object or groups.

Updates to the attack object database can include:

- New descriptions or severities for existing attack objects
- New attack objects
- Deletion of obsolete attack objects

Configuring Custom DI and IDP Attack Objects

You can create custom DI and IDP attack objects to detect new attacks or customize copies of existing attack objects to meet the unique needs of your network. For example, you might want to edit the context of a custom attack object that is producing too many false positives on your network, or you might want to create a new custom attack object to detect the latest virus or Trojan that is sweeping the Internet.

The attack object creation process is similar for custom DI and IDP attack objects. To create both object types, you use the Attack Object Wizard to enter attack object information, attack pattern, and other important information. After you have configured the object however, you use each object differently:

- To use a custom DI attack object to protect your network, you must add the object to a custom attack object group and then a DI Profile object, which you then select within the Rule Options of a firewall rule. For information about creating a custom attack object group, see [“Creating Custom IDP Attack Groups” on page 396](#). For information about creating a DI Profile object, see [“Creating DI Profiles” on page 367](#).
- To use a custom IDP attack object to protect your network, you can add the attack object in an IDP rule.

NSM enables you to import custom attacks and custom attack groups from SRX Series devices and display them as shared objects in **Object Manager**. You can also edit custom attacks and custom attack groups using **Object Manager** and update the device with these changes.

Using the Attack Object Wizard

To help you create custom attack objects, NSM UI uses a Custom Attack Object wizard to guide you through each step. During the creation process, the wizard prompts you for:

- Attack object information—You must supply an attack object name and configure the target platforms that support the attack object. You can also create an attack description, enter attack references, and set a severity for the attack object, if desired. The following sections detail the general attack object information fields.
- Attack Version information—After you have selected the target platforms, you must supply information about the attack version, including the protocol and context used to perpetrate the attack, when the attack is considered malicious, the direction and flow of the attack, the signature pattern of the attack, and the values found in the header section of the attack traffic.

To create a custom attack object, from the main navigation tree, select **Object Manager > Attack Objects > DI Objects** or **IDP Objects**, then select the **Custom Attacks** tab. Click the Add icon to display the custom attack object wizard.

Copying and Editing Predefined Attack Objects to Create Custom Attack Objects

You can also make a copy of a predefined attack object. This copy is a custom attack object, which you can modify like any other custom object. The copy must have a different name than the original, predefined attack object.

To create a custom version of a predefined attack object, open an existing predefined attack object, and click the **Edit** button in the Attack Viewer. A new attack object with the same parameters as the existing predefined attack object appears. The new object has the same name as the previous object, but with “**-Copy**” appended. After editing the parameters that you want, click **OK**.

For more information on IDP Series custom attack objects, see the *IDP Series Custom Attack Object Reference and Examples Guide*.

The following sections explain the attack object creation process; for instructions on creating a custom attack object, see the *NSM Online Help* topic, “Creating Custom Attack Objects.” The fields that can be modified are described below.

Configuring Attack Name and Description

In the **General** tab, enter the basic information about the attack, such as the attack object name and attack severity. You can also enter additional information, such as a general description and keywords, which can make it easier for you to locate and maintain the attack object as you use it in your firewall rules. Specifically, the attack object wizard prompts you for the following:

- **Name**—Enter an alphanumeric name for the object, which appears in the NSM GUI.



TIP: You might want to include the protocol the attack uses in the attack name.

- **Description**—This information is optional. Enter information about the attack, such as why you created the attack object, how the attack or exploit works, and what specific systems on your network the attack object is intended to protect. For example, you might want to include the following information:

- Attack type (buffer overflow, password exploit, format string attack, denial-of-service)
- Affected system (hardware, operating system, software application, or protocol the attack targets)
- Attack mechanism (how the attack works)
- Attack lethality (the consequences of a successful attack)

You are not required to include all this information when creating a new custom attack object, but it is a good idea. If you ever need to edit this attack object, the description can help you remember important information about the attack.

- **Severity**—Select the severity that matches the lethality of this attack on your network. Severity categories, in order of increasing lethality, are: **Info**, **Warning**, **Minor**, **Major**, and **Critical**. Critical attacks are the most dangerous—typically these attacks attempt to crash your server or gain control of your network. Informational attacks are the least dangerous, and typically are used by network administrators to discover holes in their own security system.
- **Category**—Enter the category to which the attack object belongs.
- **Keywords**—Enter descriptive words or numbers associated with the attack. Later, after you have added the custom attack object to the database, you can search using these keywords to quickly locate the attack.
- **Recommended**—Specifies that this attack object is among your highest risk set of attack objects. Later, when you add this attack object to dynamic groups, you can specify whether only recommended attack objects will be included.
- **Recommended Action**—This field only exists in predefined attack objects. When you use an attack object in a policy, you can specify what action the IDP device should take when it detects the attack. However, for IDP-capable devices running IDP 4.1 and later or ScreenOS 6.0 or later, you can tell the device to use the action recommended by Juniper Networks for that attack.
- **Detection Performance**—Select **High**, **Medium**, **Low**, or **Not defined**.

When you have completed entering the basic attack information, you can configure the extended attack information.

Configuring Extended Information

In the **Extended** tab, enter specific information about the attack. Specifically, the attack object wizard prompts you for the following:

- **Impact**—Enter details about the impact of a successful attack, including information about system crashes and access granted to the attacker.
- **Description**—Enter details about how the attack works. You might also consider adding information on the attack history (such as how it attacked your network and what steps you took to neutralize the threat).
- **Tech Info**—Enter information about the vulnerability, the commands used to execute the attack, which files are attacked, registry edits, and other low-level information.
- **Patches**—List any patches available from the product vendor, as well as information on how to prevent the attack. You might find this information in a network security advisory or from the product vendor.



NOTE: Use HTML tags to include a hyperlink within the text.

When you have completed entering the extended attack information, you can configure the external references.

Configuring External References

In the **Extended** tab, enter the external references, such as links to the security community's official descriptions of an attack, you used when researching the attack.

External references, in conjunction with standard network security references, can help other administrators get more information about how an attack works or help you research and compare the attack in relation to a suspected new attack.

Specifically, the attack object wizard prompts you for the following:

- **URLs**—Enter up to three URLs (primary, secondary, and tertiary) for external references you used when researching the attack.
- **Standard References**—Enter the standardized network security organizations' attack designations for the attack:
 - **CVE** (Common Vulnerabilities and Exposures) is a standardized list of vulnerabilities and other information security exposures. The CVE number is an alphanumeric code, such as CVE-1999-0003.
 - **BugTraq** is a moderated mailing list that discusses and announces computer security vulnerabilities. The BugTraq ID number is a three-digit code, such as 831 or 120.

When you have completed entering the external references for the attack, you can select the target platforms for the attack object.

Configuring Target Platforms

In the **General** tab, you must select the target platform, configure the attack version, then set a direction filter (described in [“Configuring the Direction Filter” on page 395](#)) for the attack object. To select the target platform and configure the attack version, click the Add icon, under **Attack Versions** to display the **New Attack Version** wizard.

On the **Target Platform and Type** page, you must select the ScreenOS or IDP versions for which the attack object is designed. Because different versions of ScreenOS and IDP support additional functionality than previous versions, you must specify the versions that must support the attack object. After you have made your selection, the attack object wizard automatically removes options from the custom attack object creation process based on the selected target platforms.

To configure the selected target platform, click the Add icon to display the New Supported Platform dialog box. Select the versions of ScreenOS 5.0 or later or IDP (idp4.0.0) that must support the attack object.



NOTE: The string `isp-sos` in a Target Platform label indicates ScreenOS software that also has IDP capability, such as the software that runs on an ISG2000.

The string `idp` (without the `sos`) in the Target Platform label indicates software that runs on a standalone IDP device, such as an IDP 600C.

Next, select the type of attack that the attack object detects. After you have added the supported platform to the custom attack object, you can configure the attack type on that platform. Select from one of the following attack types:

- **Signature Attack Object**—(DI and IDP attack objects) A signature attack object uses a stateful attack *signature* (a pattern that always exists within a specific section of the attack) to detect known attacks. Stateful signature attack objects also include the protocol or service used to perpetrate the attack and the context in which the attack occurs. If you know the exact attack signature, the protocol, and the attack context used for a known attack, select this option. For more information about creating a signature attack object, see [“Creating a Signature Attack Object” on page 376](#).
- **Protocol Anomaly Attack Object**—(IDP attack objects only) A protocol anomaly attack object detects unknown or sophisticated attacks that violate protocol specifications (RFCs and common RFC extensions). You cannot create new protocol anomalies, but you can configure a new attack object that controls how the security device handles a predefined protocol anomaly when detected. If you don’t know that exact attack signature, but you do know the protocol anomaly that detects the attack, select this option. For more information about creating a protocol anomaly attack object, see [“Configuring a Protocol Anomaly Attack Object” on page 390](#).
- **Compound Attack Object**—(IDP attack objects only) A compound attack object detects attacks that use multiple methods to exploit a vulnerability. This object combines multiple signatures and protocol anomalies into a single attack object, forcing traffic

to match a pattern of combined signatures and anomalies within the compound attack object before traffic is identified as an attack. By combining and even specifying the order in which signatures or anomalies must match, you can be very specific about the events that need to take place before the security device identifies traffic as an attack. For more information about creating a compound attack object, see [“Configuring a Compound Attack Object” on page 391](#).

If you need to detect an attack that uses several benign activities to attack your network, or if you want to enforce a specific sequence of events to occur before the attack is considered malicious, select this option.

Click **Next** to configure the attack version information for the signature attack object. You must enter some general information about attack version and specific details about the attack pattern, such as the protocol and context used to perpetrate the attack. When using a packet-related context, you can also define IP settings and protocol header matches for the attack version.

Creating a Signature Attack Object

When you configure a signature attack object, you enter important information about the protocol and context used to perpetrate the attack, when the attack is considered malicious, the direction and flow of the attack, the signature pattern of the attack, and the values found in the header section of the attack traffic.

Configuring General Attack Properties

In the **General Properties** screen, you can define the false positive frequency for the attack version, the service that the attack uses to enter your network, and the time parameters (scope and count) that determine when a traffic abnormality is identified as an attack. The following sections detail the attack version general properties.

Configuring False Positives

Select a false positive setting that indicates the frequency (**Unknown**, **Rarely**, **Occasionally**, or **Frequently**) the attack object produces a false positive on your network. Although you might not have this information when you initially configure the custom attack object, as you fine-tune your system to your network traffic you can change this setting to help you track false positives.

Configuring Service Binding (IDP Attack Objects Only)

For IDP attack objects, select the service that the attack uses to enter your network. You must select a service other than “**Any**” if you want to choose a service context for the attack object.



NOTE: For DI attack objects, you do not select a service binding. However, for custom DI attack objects, you can select a service binding.

- **Any**—If you are unsure of the correct service, select **Any** and DI attempts to match the signature in all services. Because some attacks use multiple services to attack your network, you might want to select the **Any** service binding to detect the attack regardless of which service the attack chooses for a connection.
- **IP**—If you are not sure of the correct service but know the IP protocol type, select **IP** for the service binding. You can specify the name of the protocol type, or the protocol type number. If you select IP as the service type, you should also specify an attack pattern (in the Detection area) and IP settings values (in the IP area). Additionally, if you use a context binding of first packet, you must leave the attack pattern empty. [Table 36 on page 377](#) lists the supported protocol types.

Table 36: IP Protocol Name and Type Numbers

Protocol Name	Protocol Type Number
IGMP	2
IPIP	4
EGP	8
PUP	12
TP	29
IPV6	41
ROUTING	43
FRAGMENT	44
RSVP	46
GRE	47
ESP	50
AH	51
ICMPV6	58
NONE	59
DSTOPTS	60
MTP	92
ENCAP	98
PIM	103

Table 36: IP Protocol Name and Type Numbers (continued)

Protocol Name	Protocol Type Number
COMP	108
RAW	255

- **ICMP, TCP, and UDP**—Attacks that do not use a specific service might use a specific protocol to attack your network. Some TCP and UDP attacks use standard ports to enter your network and establish a connection; to detect these attack, configure the firewall rule that contains this attack object to monitor traffic on the standard service port or ICMP ID.
- **RPC**—The remote procedure call (RPC) protocol is used by distributed processing applications to handle interaction between processes remotely. When a client makes a remote procedure call to an RPC server, the server replies with a remote program; each remote program uses a different program number. To detect attacks that use RPC, configure the service binding as RPC and specify the RPC program ID.
- **Service**—Most attacks use a specific service to attack your network. If you select **Service** as the service binding, you must select the specific service used to perpetrate the attack. Additionally, you are restricted to general attack contexts (packet, first packet, first data packet, stream, stream 256, stream 1K, or stream 8K context). To detect these attacks, configure the service binding to match the attack service. See [Table 37 on page 378](#).

Table 37: Supported Services for Service Bindings

Service	Description	Default Port
AIM	AOL Instant Messenger	
arp	Address Resolution Protocol	None
bgp	Border Gateway Protocol	TCP/179
Chargen	Chargen	TCP/19, UDP/19
DHCP	Dynamic Host Configuration Protocol	
Discard	Discard	TCP/9, UDP/9
DNS	Domain Name Service	TCP/53, UDP/53
Echo	Echo	TCP/7, UDP/7
encrypt	None	None
Finger	Finger Information Protocol	TCP/79, UDP/79
FTP	File Transfer Protocol	TCP/21, UDP/21

Table 37: Supported Services for Service Bindings (continued)

Service	Description	Default Port
Gnutella	Gnutella	
Gopher	Gopher	
h225ras	Protocol between endpoints (terminals and gateways) and gatekeepers	UDP /1718 , UDP/ 1719
h225sgn	VOIP: H.225 SGN	TCP/1720
HTTP	Hypertext Transfer Protocol	TCP/80, UDP/80
ICMP	Internet Control Message Protocol	
icmpv6	Internet Control Message Protocol version 6	
IDENT	IDENT	TCP/113
iec104	IEC 60870-5-104 process control over IP	TCP/2404
ike	Internet Key Exchange (IPSec)	UDP/500
IMAP	Internet Message Access Protocol	TCP/143, UDP/143
ip	Internet Protocol	
ipv6	Internet Protocol version 6	
IRC	Internet Relay Chat	
LDAP	Lightweight Directory Access Protocol	
lpr	Line Printer spooler	
mgcp	Message gateway control protocol	UDP/2427, UDP/2727
modbus	Serial communication protocol	TCP/502
MSN	Microsoft Instant Messenger	
msrpc	Microsoft Remote Procedure Call	UDP/135, TCP/135
mssql	Microsoft sql server	TCP/1433
mysql	Relational database management system	TCP/3306
NBName	NetBios Name Service	UDP/137 (NBName)
NBDS		UDP/138 (NBDS)

Table 37: Supported Services for Service Bindings (continued)

Service	Description	Default Port
NFS	Network File System	
nntp	Network News Transfer Protocol	
none		
NTP	Network Time Protocol	
POP3	Post Office Protocol, Version 3	TCP/110, UDP/110
Portmapper	Portmapper	TCP/111
RADIUS	Remote Authentication Dial In User Service	
rexec	Rexec	
rlogin	rlogin	TCP/513
rsh	rsh	
rtp/rtpvideo	Real time transport protocol	UDP ports (1024 to 65535)
rtsp	rtsp	
rusers	List the users logged into host	
scan		
sip	Session initiation protocol	TCP/5060, UDP/5060
SMB	Server Message Block	
SMTP	Simple Mail Transfer Protocol	TCP/25, UDP/25
SNMP	Simple Network Management Protocol	TCP/161, UDP/161
SNMPTRAP	SNMP trap	TCP/162, UDP/162
sqlmon	SQL Monitoring For Oracle Databases	UDP/1434
SSH	Secure Shell	TCP/22, UDP/22
SSL	Secure Sockets Layer	
syslog	Syslog	UDP/514

Table 37: Supported Services for Service Bindings (continued)

Service	Description	Default Port
tcp	Transmission Control Protocol	
Telnet	Telnet TCP protocol	TCP/23, UDP/23
TFTP	Trivial File Transfer Protocol	
tns	Oracle Transparent Network Substrate (TNS) protocol	TCP/1521/2483 /1525/1527/1529
udp	User Datagram Protocol	
unspecified		
VNC	Virtual Network Computing	
Whois	whois	
YMSG	Yahoo! Messenger	

- **IPv6 or ICMPv6**—Do not select these options for IDP Series devices, as these devices do not support inspection of IPv6.

Configuring Time Binding

Use **Time Binding** to configure the time attributes for the custom attack object. Time attributes control how the attack object identifies attacks that repeat for a certain number of times. By configuring the scope and count of an attack, you can detect a sequence of the same attacks over a period of time (one minute) across sessions.

After you enable **Time Binding**, configure the following time attributes:

- **Scope**—Select the scope within which the count occurs:
 - **Source**—Select this option to detect attacks from the source IP address for the specified number of times, regardless of the destination IP address.
 - **Destination**—Select this option to detect attacks to the destination IP address for the specified number of times, regardless of the source IP address.
 - **Peer**—Select this option to detect attacks between source and destination IP addresses of the sessions for the specified number of times.
- **Count/Min**—Enter the number of times per minute that the attack object must detect an attack within the specified Scope before the device considers the attack object to match the attack. For example, the TCP Protocol Anomaly “Segment Out of Window” is harmless and is occasionally seen on networks. Thousands of these anomalies between given peers, however, is suspicious.

The minute timer starts when the signature first matches the event. If the signature matches the same event for the specific count or higher within 60 seconds, the signature is considered to have matched the attack object.

If you bind the attack object to multiple ports (see [“Configuring Attack Detection Properties” on page 383](#)) and the attack object detects that attack on different ports, each attack on each port is counted as a separate occurrence. For example, when the attack object detects that attack TCP/80 and then on TCP/8080, the count is two.

After you finish entering the general attack properties for the attack type, you can configure the constraints.

Configuring Constraints

Using constraints, you can define more accurate signatures, thereby reducing the number of false positives.

You can configure the following constraints:

- **With-in Bytes Constraint**—Use this constraint when you want the device to inspect for an attack pattern within a specific byte range:
 - **Lower Limit**—Specify the beginning of the range, that is, byte number 1.
 - **Upper Limit**—Specify the end of the range, that is, byte number 10.
 - **Start Point**—Typically, your selection must be consistent with your pattern context setting. For example, if you configured one of the service contexts, Juniper Networks recommends you select **Context**. If you configured one of the packet contexts, Juniper Networks recommends you select **Packet**. If you configured one of the stream contexts, Juniper Networks recommends you select **Stream**.

The IDP Series device monitors all packets in a session that fall within this range limit for an attack pattern.

You can set multiple constraints. If the starting points are the same, the constraints are evaluated as a Boolean OR. If the starting points are different, the constraints are evaluated as a Boolean AND.

- **With-in Packets Constraint**—Use this constraint when you want the device to inspect for an attack pattern within a specific packet range:
 - **Lower Limit**—Specify the beginning of the range, that is, packet number 1 in the stream.
 - **Upper Limit**—Specify the end of the range, that is, packet number 2 in the stream.

The IDP Series device monitors all packets in a session that fall within this range limit for an attack pattern.

- **Context Check**—Use this constraint to require the matching context be of a specified byte length:
 - **Constraint**—Select **length**.
 - **Comparison Operator**—Select **=**, **!**, **>**, or **<**.
 - **Operand**—Select a byte length.

Only standalone IDP sensors running IDP 5.1 and later support this feature. When you select this action in an APE rule installed on a device running IDP 5.0 or earlier, NSM displays a message warning the user that if unsupported, this APE action might cause a device update failure.

Configuring Attack Detection Properties

In the **Attack Pattern** screen, you can define the signature pattern of the attack, the context in which the attack occurs, and the direction and flow of the attack.

Configuring Attack Pattern

The attack pattern is the signature of the attack you want to detect. A signature is a pattern that always exists within an attack; if the attack is present, so is the signature. To create the attack pattern, you must first analyze the attack to detect a pattern (such as a segment of code, a URL, or a value in a packet header), then create a syntactical expression that represents that pattern. [Table 38 on page 383](#) lists the syntax based on regular expressions to match signature patterns for DI and IDP.

Table 38: Attack Pattern Syntax

Pattern	Syntax
Bit-level match for binary protocols. The length of the bitmask must be in multiples of 8.	\B.0.1..00\B
The first \B denotes the start of the bitmask. The last \B denotes the end of the bitmask.	
The decimal (.) indicates the bit can be either 0 or 1.	
A 0 or 1 indicates the bit at that position must be 0 or must be 1.	
Direct binary match (octal).	\0<octal-number>
Direct binary match (hexadecimal).	\X<hexadecimal-number>\X
Case insensitive matches.	\[<character-set\]
Match any symbol.	.
Match 1 or more symbols.	*
Match 0 or 1 symbols.	?
Match 1 or more symbols.	+
Grouping of expressions.	()
Alternation, typically used with ().	
Character range.	[<start>-<end>]

Table 38: Attack Pattern Syntax (continued)

Pattern	Syntax
Character class. Any explicit value within the bracket at the position matches.	[]
Negation of range.	[^<start>-<end>]
Unicode insensitive matches.	\u<string>\u
Whitespace.	\s
Use a backslash to escape special characters so that they are matched and not processed as regular expression operators.	
Character	Escaped
*	*
(\(
)	\)
.	\.
+	\+
\	\\
[\0133
]	\0135
NOTE: Because the combination of the backslash and the open and close square brackets is used in the case-insensitive expression, you should use the backslash with the octal code for the bracket characters, as shown above.	



NOTE: Regular expression support is provided by the PCRE library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England.

Table 39 on page 384 lists some example syntax matches.

Table 39: Attack Pattern Syntax Example Matches

This syntax	Matches	Example
\X01 86 A5 00 00\X	the five specified bytes verbatim	01 86 A5 00 00

Table 39: Attack Pattern Syntax Example Matches (continued)

This syntax	Matches	Example
(hello world)	hello or world	hello world
(hello world) +	hello or world one or more times	helloworld world hello hellohello
\[hello\]	hello in a case insensitive manner	hElLo HElLO heLLO
[c-e]a(d t)	Anything with the first letter of c, d, or e, the middle letter a and ending in d or t	cad cat dad dat ead eat
[^c-d]a(d t)	Expressions that begin with a letter other than c, d, or e, have the second letter a, and end in d or t	fad zad
a*b+c	Any number of “a” characters followed by one or more b characters followed by a c.	abc aaabbbc

To negate the pattern, enable **Negate**.

Configuring Attack Context

Select the context that defines the location of the signature.



NOTE: For IDP attack objects, if you selected “Any” as the Service Binding in the Attack Pattern screen, you cannot select a service context here.

If you know the service and the specific service context, select that service then select the appropriate service contexts. If you know the service, but are unsure of the specific service context, select **Other** then select one of the following general contexts:



NOTE: If you select a stream, stream 256, stream 1K, stream 8K, or a service context, you cannot specify IP header contents (in the Header Match screen).

- Select packet context to match the attack pattern within a packet. When you select this option, you should also specify the Service Binding (in the **General** tab) and define the service header options (in the **TCP/UDP/ICMP Header Matches** menu). Although not required, specifying these additional parameters helps to improve the accuracy of the attack object and can improve performance.
- Select first packet context to detect the attack in only the first packet of a stream. When the flow direction for the attack object is set to any, the security device checks the first packet of both the server-to-client (STC) and client-to-server (CTS) flows. If you know that the attack signature appears in the first packet of a session, choosing first packet instead of packet reduces the amount of traffic the security device needs to monitor, thereby improving performance.
- Select first data packet context to detect the attack in only the first data packet of a stream. If you know that the attack signature appears in the first data packet of a session, choosing first data packet instead of packet reduces the amount of traffic the security device needs to monitor, thereby improving performance.
- Select stream context to reassemble packets and extract the data to search for a pattern match. However, a security device does not recognize packet boundaries for stream contexts, so data for multiple packets is combined. Select this option only when no other context option contains the attack.
- Select stream 256 context to reassemble packets and search for a pattern match within the first 256 bytes of a traffic stream. When the flow direction is set to any, the security device checks the first 256 bytes of both the STC and CTS flows. If you know that the attack signature will appear in the first 256 bytes of a session, choosing stream 256 instead of stream reduces the amount of traffic that the security device must monitor and cache, thereby improving performance.
- Select stream 8K context to reassemble packets and search for a pattern match within the first 8192 bytes of a traffic stream. If you know that the attack signature will appear in the first 8192 bytes of a session, choosing stream 8K instead of stream reduces the amount of traffic that the security device must monitor and cache, thereby improving performance.
- Select stream 1K context to reassemble packets and search for a pattern match within the first 1024 bytes of a traffic stream. If you know that the attack signature will appear in the first 1024 bytes of a session, choosing stream 1K instead of stream reduces the amount of traffic that the security device must monitor and cache, thereby improving performance.

Configuring Attack Direction

Select the connection direction of the attack. Using single direction (instead of Any) improves performance, reduces false positives, and increases detection accuracy:

- **Client to Server**—Detects the attack only in client-to-server traffic

- **Server to Client**—Detects the attack only in server-to-client traffic
- **Any**—Detects the attack in either direction

Configuring Attack Flows

Select the connection flow of the attack. Using a single flow (instead of Both) improves performance and increases detection accuracy.

- **Control**—Detects the attack in the initial connection that is established persistently to issue commands, requests, and so on.
- **Auxiliary**—Detects the attack in the response connection established intermittently to transfer requested data.
- **Both**—Detects the attack in the initial and response connections.

After you finish entering the attack detection properties for the attack type, click **Next** to configure the attack IP settings and protocol headers.

Configuring Header Match Properties

Specify specific values and options that exist within the header of the attack packet.



NOTE: You can configure header values only for attack objects that use a packet, first data packet, or first packet context. If you selected a stream, stream 256, stream 1K, stream 8K, or a service context (in the Detection area) you cannot specify header contents.

If you are unsure of the options or flag settings for the malicious packet, leave all fields blank and the security device attempts to match the signature for all header contents. For each value you enter, you must specify the relational or equality operator. [Table 40 on page 387](#) lists DI attack header match modifiers.

Table 40: DI Attack Header Match Modifiers

Modifier	Meaning
=	equal to
!	not equal to
>	greater than
<	less than

Additionally, for each flag you must specify whether or not a flag is configured (none), the flag is set (set), or the flag is not set (unset).

Configuring IP Header Matches

For attacks that use IP and a packet context, you can go to the **IP** tab to set values for the IP fields and flags listed below. Note that in the IP header match GUI, you can select either **IPv4** or **IPv6**, but if you select IPv6, you can configure a new ICMPv6 header match in the Protocol header along with existing TCP and UDP protocols. IPv4 and IPv6 header matches cannot coexist in a single attack definition. IPv6-enabled attacks are supported only on ISG1000 with SM and ISG2000 with SM devices.

- **Type-of-service**—Enter the service type. Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the service type. Common service types are:
 - 0000 Default
 - 0001 Minimize Cost
 - 0002 Maximize Reliability
 - 0003 Maximize Throughput
 - 0004 Minimize Delay
 - 0005 Maximize Security
- **Packet length**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the number of bytes in the packet, including all header fields and the data payload.
- **Id**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the unique value used by the destination system to reassemble a fragmented packet.
- **Time-to-live**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the time-to-live (TTL) value of the packet. This value represents the number of routers the packet can pass through. Each router that processes the packet decrements the TTL by 1; when the TTL reaches 0, the packet is discarded.
- **Protocol**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the protocol used.



NOTE: The Protocol field does not appear for DI attack objects.

- **Source**—Enter the source IP of the attacking device.
- **Destination**—Enter the destination IP of the attack target.
- **RB**—Reserved Bit. This bit is not used.
- **MF**—More Fragments. When set (1), this option indicates that the packet contains more fragments. When unset (0), it indicates that no more fragments remain.
- **DF**—Don't Fragment. When set (1), this option indicates that the packet cannot be fragmented for transmission.

Configuring TCP Header Matches

For attacks that use TCP and a packet context, in the **Protocols** tab, select **TCP Packet Header Fields** from **TCP/UDP/ICMP Header Matches** menu, then set values for the following TCP fields and flags:

- **Source Port**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the port number on the attacking device.
- **Dest Port**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the port number of the attack target.
- **Seq. Number**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the sequence number of the packet. This number identifies the location of the data in relation to the entire data sequence.
- **ACK. Number**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the ACK number of the packet. This number identifies the next sequence number; the ACK flag must be set to activate this field.
- **Header Length**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the number of bytes in the TCP header.
- **Data Length**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the number of bytes in the data payload. For SYN, ACK, and FIN packets, this field should be empty.
- **Window Size**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the number of bytes in the TCP window size.
- **UrgPtr**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the urgent pointer. The value indicates that the data in the packet is urgent; the URG flag must be set to activate this field.
- **Urgent Bit**—When set, the urgent flag indicates that the packet data is urgent.
- **ACK bit**—When set, the acknowledgment flag acknowledges receipt of a packet.
- **PSH bit**—When set, the push flag indicates that the receiver should push all data in the current sequence to the destination application (identified by the port number) without waiting for the remaining packets in the sequence.
- **RST bit**—When set, the reset flag resets the TCP connection, discarding all packets in an existing sequence.
- **SYN bit**—When set, the SYN flag indicates a request for a new session.
- **FIN bit**—When set, the final flag indicates that the packet transfer is complete and the connection can be closed.
- **R1 bit**—This reserved bit (1 of 2) is not used.
- **R2 bit**—This reserved bit (2 of 2) is not used.

Configuring UDP Header Matches

For attacks that use UDP and a packet context, in the **Protocols** tab, select **UDP Packet Header Fields** from **TCP/UDP/ICMP Header Matches** menu, then set values for the following UDP fields:

- **Source Port**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the port number on the attacking device.
- **Dest. Port**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the port number of the attack target.
- **Data Length**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the number of bytes in the data payload.

Configuring ICMP Header Matches

For attacks that use ICMP and a packet context, in the **Protocols** tab, select **ICMP Packet Header Fields** from **TCP/UDP/ICMP Header Matches** menu, then set values for the following ICMP fields:

- **Type**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the primary code that identifies the function of the request/reply.
- **Code**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the secondary code that identifies the function of the request/reply within a given type.
- **Seq. Number**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the sequence number of the packet. This number identifies the location of the request/reply in relation to the entire sequence.
- **Id**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the identification number, which is a unique value used by the destination system to associate requests and replies.
- **Data Length**—Specify an operand (**none**, **=**, **!**, **>**, **<**) and a decimal value for the number of bytes in the data payload.



NOTE: IDP Series 5.1r1 devices do not support IPV6/ICMP6 configuration.

Configuring a Protocol Anomaly Attack Object

A protocol anomaly attack object locates unknown or sophisticated attacks that violate protocol specifications (RFCs and common RFC extensions). You cannot create new protocol anomalies, but you can configure a custom attack object that controls how the security device handles a predefined protocol anomaly when detected.



NOTE: Protocol anomaly attack objects are supported by IDP-capable security devices only, such as the ISG2000 or ISG1000 running ScreenOS 5.3 or later IDP.

To configure a custom protocol anomaly attack object, you must:

- Configure the false positive setting—For details, see [“Configuring Attack Detection Properties” on page 383](#).
- Select a predefined protocol anomaly—Select the protocol anomaly you want to use for this attack object. The list of available predefined protocol anomalies depends on the protocols supported by the target platform. For details, refer to the NSM Online Help.
- Configure the time-based settings—For details, see [“Configuring Time Binding” on page 381](#).

Configuring a Compound Attack Object

A compound attack object combines multiple signatures and protocol anomalies into a single attack object, forcing traffic to match all combined signatures and anomalies within the compound attack object before traffic is identified as an attack. By combining and even specifying the order in which signatures or anomalies must match, you can be very specific about the events that need to take place before the security device identifies traffic as an attack.

NSM 2006.1 and later releases also support Boolean expressions for standalone IDP signatures.



NOTE: Compound attack objects are supported by IDP-capable security devices only, such as the ISG series with Security Module or any of the standalone IDP Sensors. ISG series devices do not support Boolean expressions.

When configuring a custom compound attack object:

- All members of the compound attack object must use the same service setting or service binding, such as FTP, Telnet, YMSG, or TCP/80.
- You can add protocol anomaly attack objects to a compound attack object.
- You cannot add predefined or custom attack objects to a compound attack object. Instead, you specify the signature directly within the compound attack object, including such details as service (or service binding), service context, attack pattern, and direction.
- You can add between 2 and 32 protocol anomaly attack objects and signatures as members of the compound attack object. However, all members must use the same service setting or service binding.

Configuring General Attack Properties

False positives and time-based attack properties are configured for a compound attack object the same way as they are for a signature attack object.

Because all members of the compound attack object must use the same service binding, the service binding you select determines the service contexts you can use for an attack

pattern, as well as the available predefined protocol anomaly attack objects you can add as members.

- To match all services, select **Any** as the Service Binding.
 - When adding an attack pattern as a member, you are restricted to the contexts packet, first data packet, and first packet.
 - When adding a predefined protocol anomaly attack object as a member, you are restricted to the IP-based protocol anomaly attack objects.

Additionally, because the number of session transactions are not known for the service, you cannot specify a scope (in the Compound Members page).

- To match a specific service, select the service binding and provide the protocol ID, port/port range, program number if necessary.

Next, configure the members of the compound attack object.

Configuring Compound Attack Members

When configuring members, you add the signatures and protocol anomalies to detect an attack that uses multiple methods to exploit a vulnerability. The attack traffic must match all signatures and anomalies within the compound attack object before the device considers the traffic as an attack. To be explicit about the events in an attack, you can also specify the order in which signatures or anomalies must match before the security device identifies traffic as an attack.

Configuring the Attack Object Scope

If the selected service supports multiple transactions within a single session, you can also specify whether the match should occur over a single session or can be made across multiple transactions within a session:

- Select **Session** to allow multiple matches for the object within the same session.
- Select **Transaction** to match the object across multiple transactions that occur within the same session.

Select **Reset** to detect multiple occurrences of the attack object in the same session; disable it to log multiple occurrences as one.

Configuring a Boolean Expression

The **Boolean Expression** field makes use of the **Member Names** created in the lower part of the dialog.

NSM supports three Boolean operators: **or**, **and**, and **oand** (ordered and). NSM also supports the use of parenthesis to determine precedence.

Boolean operators:

- **or**—If either of the member name patterns match, the expression matches.
- **and**—If both of the member name patterns match, the expression matches. It does not matter in which order the members appear.

- **oand**—If both of the member name patterns match, and if they appear in the same order as in the Boolean Expression, the expression matches.

Example: Boolean Expression

Suppose you have created six signature members, labeled s1 through s5.

Suppose you know that the attack always contains the pattern s1, followed by either s2 or s3. Further, you know that the attack always contains s4 and s5, but their positions in the attack can vary.

You might create the following Boolean expression:

((s1 oand s2) or (s1 oand s3)) and (s4 and s5)

Configuring an Attack Pattern

You configure the attack pattern as a member of a compound attack object as you would an attack pattern in a signature attack object. For details, see [“Configuring Attack Detection Properties” on page 383](#).

To add an attack pattern to the compound attack object, click the Add icon and select **Signature**. In the **Signature Parameters** dialog box, configure the following parameters:

- **Member Name**—Specify a member name.
Juniper Networks signature team uses the following naming convention for members: m01, m02, m03, and so forth. Juniper Networks recommends that you use this same naming convention.
- **Pattern**—Specify the pattern to match. You construct the attack pattern just as you would when creating a new signature attack object. To negate the pattern, enable **Negate**.
- **Context**—Specify the context in which to locate the pattern. The context displays only contexts that are appropriate for the specified Service. If you selected a service binding of **Any**, you are restricted to the service contexts packet, first data packet, and first packet.
- **Direction**—Specify whether the security device should match the pattern in traffic flowing in any direction, from client to server, or from server to client.

Adding a Predefined Protocol Anomaly Attack Object

To add a protocol anomaly to the compound attack object, click the Add icon and select protocol anomaly.

If you selected a service binding of *any*, you are restricted to the IP-based protocol anomaly attack objects.

To add a protocol anomaly to the compound attack object, click the Add icon and select **Anomaly**. In the **Anomaly Parameters** dialog box, configure the following parameters:

- **Member Name**—Specify a member name.

Juniper Networks signature team uses the following naming convention for members: m01, m02, m03, and so forth. Juniper Networks recommends that you use this same naming convention.

- **Anomaly**—Select an anomaly from the drop-down list. The menu only displays protocol anomalies appropriate for the service you selected.

If you selected a service binding of *any*, you are restricted to the IP-based protocol anomaly attack objects.

Configuring an Attack Object Ordered Match

Use the oAND Boolean operator to create a compound attack object that must match each member signature or protocol anomaly in the order you specify. If you do not specify an ordered match, the compound attack object still must match all members, but the attack pattern or protocol anomalies can appear in the attack in random order.

To configure an ordered match, use the oAND Boolean operator in the Boolean Expression field to match a more complex arrangement of attack patterns. For example, if you have created signatures m1 and m2, and you know that attack m1 is always followed by m2, the Boolean expression that you create would be **m1 oAND m2**.



NOTE: With 2010.4 release and later, the **Ordered Match** checkbox is deprecated and its functionality replaced with the oAND Boolean operator.

Configuring Constraints

You can configure the following constraints:

- **Context-Check**—Use this constraint to require the matching context be of a specified byte length:
 - **Constraint**—Select **length**.
 - **Comparison Operator**—Select **=**, **!**, **>**, or **<**.
 - **Operand**—Select a byte length.
- **Match within same context**—Use this constraint when you want the selected signature members to be found in the same context instance (in any order).

Protocol anomaly members are not selectable and not a component of this constraint.

- **With-in Bytes Constraint**—Use this constraint when you want the device to inspect for an attack pattern within a specific byte range in a stream:
 - **Lower Limit**—Specify the beginning of the range, that is, byte number 1.
 - **Upper Limit**—Specify the end of the range, that is, byte number 10.
 - **Member**—Select members. You can select members one at a time and set a lower and upper limit for each one. In this case, the program logic begins counting bytes

from start-of-stream for each member. The member must be found within the byte range indicated.

- **With-in Packet Constraint**—Use this constraint when you want the device to inspect for an attack pattern within a specific packet range of a stream:
 - **Lower Limit**—Specify the beginning of the range, that is, packet number 1 in the stream.
 - **Upper Limit**—Specify the end of the range, that is, packet number 2 in the stream.
 - **Member**—Select members. You can select members one at a time and set a lower and upper limit for each one. In this case, the program logic counts packets beginning with the start-of-stream. The member must be found within the packet range indicated.

Only standalone IDP sensors running IDP 5.1 and later support this feature. When you select this action in an APE rule installed on a device running IDP 5.0 or earlier, NSM displays a message warning the user that if unsupported, this APE action might cause a device update failure.

Configuring the Direction Filter

Use the direction filter to specify the direction (Any, Client-to-Server, Server-to-Client) of traffic in which the attack object attempts to match an attack. Each attack version in the attack object retains its own direction; however, you can use the direction filter to change which direction is monitored by the attack object. Only those attack versions that match the direction filter are active in the attack object.

By default, the direction filter is automatically set to the direction of the most recently-created or edited attack version.

Creating Custom DI Attack Groups

You can create custom attack object groups to contain your custom DI attack objects. After you add these custom groups to a DI profile, you can then configure a firewall rule to use that DI Profile.

All DI attack object groups (both predefined and custom) are considered “static” groups, meaning that they do not change. To add or delete an attack object from the group, you must manually edit the group members.

A custom attack object group can contain custom attack objects and other custom attack object groups. You cannot add predefined attack objects or predefined attack object groups to a custom attack object group. To use both predefined and custom attack objects in a firewall rule, create a DI Profile that includes predefined and custom attack object groups, then use this profile object within the Rule Options of a firewall rule. For information about creating a DI Profile, see [“Creating DI Profiles” on page 367](#).



NOTE: Attack group names cannot be the same as attack object names.

Creating Custom IDP Attack Groups

NSM contains a database of hundreds of predefined attack objects designed to protect networks from multiple attack vectors.

For IDP attack objects, you can create static or dynamic groups to contain predefined or custom attack objects. A static group contains only the groups or attack objects you specify, while a dynamic group contains attack objects based on criteria you specify. Although you do not have to create a group to use an attack object within an IDP rule (you can add attack objects individually or by group), organizing attack objects into groups can help keep your security policies organized.

Creating Static Attack Groups

A static group contains a specific, finite set of attack objects or groups. There are two types of static groups: predefined static groups and custom static groups.

A custom static group can include the same members as a predefined static group (predefined attack objects, predefined static groups, and predefined dynamic groups), plus the following members:

- Custom attack objects
- Custom dynamic groups
- Other custom static groups

Use static groups to define a specific set of attacks to which you know your network is vulnerable, or to group custom attack objects. For example, you might want to create a group for a specific set of informational attack objects that keep you aware of what is happening on your network.

Static groups require more maintenance than dynamic groups because you must manually add or remove attack objects in a static group to change the members. However, you can include a dynamic group within a static group to automatically update some attack objects. For example, the predefined attack object group Operating System is a static group that contains four predefined static groups: BSD, Linux, Solaris, and Windows. The BSD group contains the predefined dynamic group BSD-Services-Critical, to which attack objects can be added during an attack database update.

To create a custom static group:

1. In Object Manager, select **Attack Objects > IDP Objects**. The IDP Objects dialog box appears.
2. Click the Custom Attack Groups tab, then click the Add icon and select **Add Static Group**. The New Static Group dialog box appears.
3. Enter a name and description for the static group. Select a color for the group icon.
4. To add an attack or group to the static group, select the attack or group from the Attacks/Group list and click the **Add** button.
5. Click **OK**.

For instructions for creating a static attack object group, see the *NSM Online Help* topic “Adding Static Attack Groups.”

Creating Dynamic Attack Groups (IDP Only)

A dynamic group contains a dynamic set of attack objects that are automatically added or deleted based on specified criteria for the group. For example, an attack database update can add or remove attack objects from a dynamic group based on the group criteria. This eliminates the need to review each new signature to determine if you need to use it in your existing security policy.

A predefined or custom dynamic group can only contain attack objects and not attack groups. Dynamic group members can be either predefined or custom attack objects.

To create a custom dynamic group:

1. In Object Manager, select **Attack Objects > IDP Objects**. The IDP Objects dialog box appears.
2. Click the Custom Attack Groups tab, then click the Add icon and select **Add Dynamic Group**. The New Dynamic Group dialog box appears.
3. Enter a name and description for the static group. Select a color for the group icon.
4. In the Filters tab, click the Add icon and select one of the following:
 - Add Products Filter to add attack objects based on the application that is vulnerable to the attack.
 - Add Severity Filter to add attack objects based on the attack severity.



NOTE: All predefined attack objects are assigned a severity level by Juniper Networks. However, you can edit this setting to match the needs of your network.

- Add Category Filter to add attack objects based on category.
- Add Last Modified Filter to add attack objects based on their last modification date.
- Add Recommended Filter to include only attacks designated to be the most serious threats to the dynamic group. In the future, Juniper Networks will designate only attacks it considers to be serious threats as Recommended. These settings will be updated with new attack object updates. In addition, you can designate custom attack objects as Recommended or not.

You create filters one at a time; each criteria you add is compared to the attributes for each attack object. Attack objects that do not match the criteria are immediately filtered out. If you create a filter with attributes that no attack object can match, a message appears warning you that your dynamic group has no members.

From the resulting list of matching attack objects, you can then exclude any attack objects that produces false positives on your network, or an attack object that detects an attack to which your network is not vulnerable.



NOTE: A dynamic group cannot contain another group (predefined, static, or dynamic). However, you can include a dynamic group as a member of a static group.

Example: Creating a Dynamic Group

To create a dynamic group:

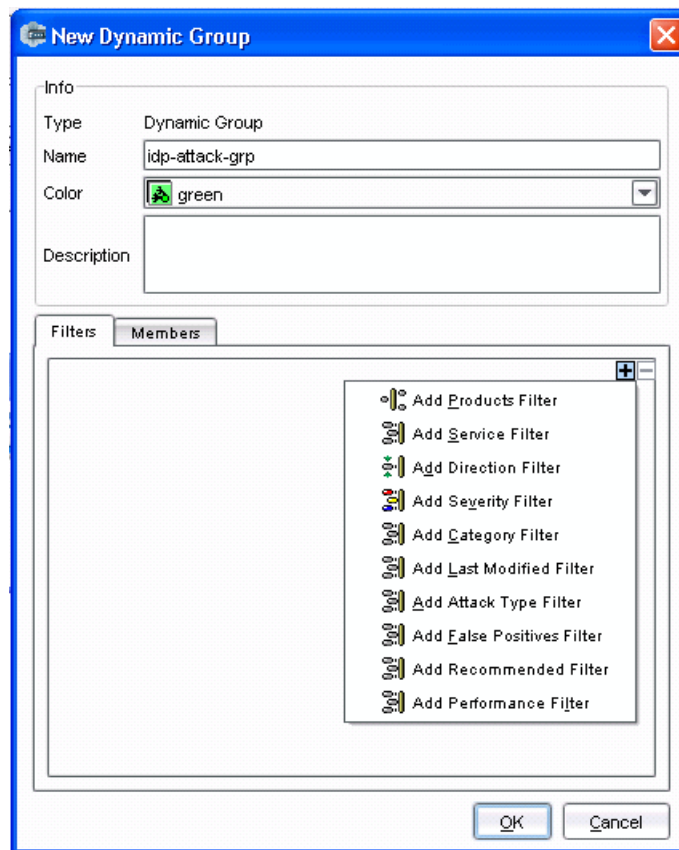
1. In the Custom Attack Groups tab, click the Add icon, and select **Add Dynamic Group**. The New Dynamic Group dialog box appears.
2. Enter a name and description for the group. Select a color for the group icon.

Figure 84: New Dynamic Group

The image shows a 'New Dynamic Group' dialog box with a blue title bar. It has two tabs: 'Info' and 'Filters'. The 'Info' tab is active, showing fields for 'Type' (set to 'Dynamic Group'), 'Name' (set to 'idp-attack-grp'), 'Color' (set to 'green' with a small green icon), and 'Description' (an empty text area). The 'Filters' tab is also visible, showing an empty list area with a '+' icon in the top right corner. At the bottom right are 'OK' and 'Cancel' buttons.

3. In the Filters tab, click the Add icon and add the filters that determine which attack objects should be in the group:

Figure 85: New Dynamic Group Filters

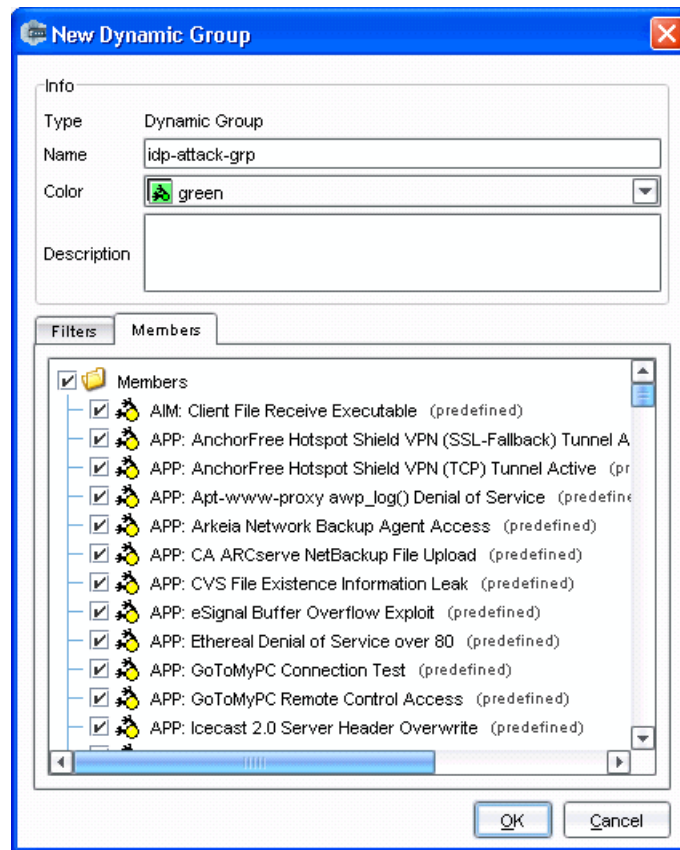


- a. Add a Products filter to add attack objects that detect attacks against all Microsoft Windows operating systems.
- b. Add a Severity filter to add attack objects that have a severity level of critical or major.

IDP automatically applies all filters to the entire attack object database, identifies the attack objects that meet the defined criteria, and adds the matching objects as members of the group.

4. View the members of the group by clicking the Members tab.

Figure 86: New Dynamic Group Members



5. Click **OK** to save the dynamic group.

Updating Dynamic Groups

When you are satisfied with the group criteria and its members, use the group in a security policy. The next time you update your attack objects, the update automatically performs the following:

- For all new attack objects, compares the predefined attributes of each attack object to each dynamic group criteria and adds the attack objects that match.
- For all updated attack objects, removes attack objects that no longer meet their dynamic group criteria. The update also reviews updated attack objects to determine if they now meet any other dynamic group criteria, and adds them to those groups if necessary.
- For all deleted attack objects, removes the attack objects from their dynamic groups.

You can also edit a dynamic group manually, adding new filters or adjusting existing filters to get exactly the type of attack objects your want.



NOTE: You can edit a custom dynamic attack group from within an IDP rule in a security policy. Double-click the group icon in the Attack Objects column of an IDP rule to display the Dynamic Group dialog box, make the desired changes, then click **OK** to save your edits.

Editing a Custom Attack Group

To modify a custom attack group, double-click the group in the Custom Attack Groups tab in the IDP Objects dialog box. The Static Group or Dynamic Group dialog box appears, with the previously-configured information displayed. Enter any changes you want to make and then click **Apply** to continue making changes or click **OK** to close the dialog box.

Deleting a Custom Attack Group

To delete a custom attack group, right-click the group in the Custom Attack Groups tab in the IDP Objects dialog box, and then select **Delete**. A confirmation window asks you to verify that you want to delete the item. Click **OK**.

Configuring Application Identification

Juniper Networks provides predefined application signatures that detect Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) applications running on nonstandard ports. Identifying these applications allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports, thereby ensuring security.

The device uses these signature patterns to identify a TCP or UDP application by matching the first few packets of a session (in both client-to-server and server-to-client sessions). The device then restricts or allows this traffic based on the action specified in the IDP rule.

The steps involved are:

From NSM:

1. Download the IDP signature database, which includes the app-id database.
2. Perform an IDP attack database update, which updates the app-id database.

However, if you are using the CLI, with Junos OS Release 10.2 and later, you can update the app-id database along with the IDP signature attack database or separately.

Updating the NSM App-ID Database

The predefined application identification (app-id) package database is part of the IDP signature database and is available on the Juniper Networks website. The app-id database can be downloaded along with the IDP signature database to the NSM server.

To update the app-id database on the NSM server with the latest version from the Juniper Networks website:

1. From the **Tools** menu, select **View / Update NSM Attack Database**. The Attack Update Manager dialog box appears.
2. Click **Next**. The dialog box displays the following:
 - A list of security packages and their versions installed on the NSM server.
 - The latest version to which these packages can be upgraded to.
3. Click **Finish**.

The Job Information dialog box appears, displaying whether the download was successful and the downloaded app-id database version. If the installation was unsuccessful, an error message is displayed.
4. Click **Close** to exit the Job Information dialog box.

Updating the Device App-ID Database

To update the device with the latest version of the app-id database:

1. Perform the “[Updating the NSM App-ID Database](#)” on page 401 procedure.
2. From the **Devices** menu, select **Deep Inspection/IDP > Update Device Attack Database**. A dialog box appears, listing the SRX Series devices on which an IDP attack update is permissible.
3. Select one or more devices and click **OK**.

The Job Information dialog box appears, displaying whether the installation was successful. If the installation was unsuccessful, an error message is displayed.



NOTE: Trying to update a device that already has the latest version of the IDP attack database installed produces an error. You may want to check the IDP attack database version installed on the device before proceeding to update the device.

4. Click **Close** to exit the Job Information dialog box.

Viewing the Device App-ID Database Version

You can view the app-id database version currently installed on the device. You should check to make sure the device does not already have the latest version installed, as any further update will result in an error. To view the app-id database version:

1. From the **Devices** menu, select **View/Update App-id > Check App-id Database Server Version**. A dialog box appears listing the high-end SRX Series devices that have the app-id database installed.
2. Select one or more devices and click **OK**.

The Job Information dialog box appears, displaying the version of the app-id database installed on the selected device(s). This information is fetched from the respective device/(s).

3. Click **Close** to exit the Job Information dialog box.

Uninstalling an App-ID Database from the Device

You may want to uninstall an app-id database from a device when you no longer want to monitor applications that run on nonstandard ports. To uninstall an app-id database:

1. From the **Devices** menu, select **View/Update App-id > Uninstall App-id Database**. A dialog box appears listing the high-end SRX Series devices that have the app-id database installed.
2. Select one or more devices from which you want to uninstall the app-id database and click **OK**.

The Job Information dialog box appears displaying the progress of the uninstallation and whether it was successful. After the uninstallation, the app-id version installed on the device will be zero.

3. Click **Close** to exit the Job Information dialog box.

Unified Threat Management

- [Creating UTM Profiles on page 403](#)
- [Creating an Antivirus Profile on page 404](#)
- [Creating an Antispam Profile on page 405](#)
- [Creating a Content Filtering Profile on page 406](#)
- [Creating a URL Filtering Profile on page 406](#)
- [Miscellaneous UTM Features on page 407](#)
- [ScreenOS Threat Management Features on page 410](#)

Creating UTM Profiles

A UTM profile can define more than one UTM feature. You can have more than one custom feature profile for the different supported protocols. The predefined, profiles are recommended although you can define custom profiles as well.

To create a UTM profile:

1. Select **Object Manager > UTM > Profiles > +**. The **New UTM Profile** window opens.
2. Enter a name for the profile.
3. Enter a comment or description.
4. Select a color from the drop-down list.
5. Select **Profile Properties**:
 - Connections per client: 0-20000

- Behavior for over-limit connections: either drop or permit
6. Click **OK**. The new profile is displayed in the list of UTM profiles.

You can define profiles for antivirus, anti-spam, URL filters, and content filters for the new profile either from the same window or by navigating from their respective nodes in the navigation pane. You can create miscellaneous objects such as Extension lists, URL Patterns and Categories and so on in UTM and use them in your various UTM profiles such as Antivirus or Antispam.

Creating an Antivirus Profile

1. Select **Object Manager > UTM > Anti-Virus Profiles**. NSM displays two tables: Predefined UTM AV Profiles on top and Custom UTM AV Profiles below. You can only view but not edit the listed predefined profiles. You can create and edit custom profiles.
2. Select **+** in the **Custom UTM AV Profiles** table. The **New Anti-Virus Profile** window opens.
3. Enter a name for the profile.
4. Select an Engine type: either Kaspersky Lab Engine or Juniper Express Engine. Properties change according to the chosen engine.
5. Enter a comment or description.
6. Select a color from the drop-down list.
7. Select **Engine Properties**. If you select Kaspersky Lab Engine:
 - Enable Intelligent Pre-screening.
 - Set maximum content size. Mouse over the field to see a tool tip with the allowed values. The allowed range is 20-20000.
 - Set a time-out period. The allowed range is 1-1800.
 - Set the decompression limit in the range of 1-4.
 - Set the HTTP tricking time-out in the range of 0-600.
 - Set the scan mode: None, All, By-extension.
 - Select an extension list from the drop-down list or create a new one by clicking on **+** beside the field.
 - Set notification options for: Virus detection, Fallback Block, and Fallback Non-Block.
 - Set actions for the various situations: Default, Too many requests, Corrupt file, Out of resources, Engine not ready, Password file, Time out, Maximum content size, and Decompression Layer. You can select one of these actions for each situation: None, Block, Log and Permit.

If you select the Juniper Express Engine, you need to also enable the same settings with the following exceptions, which do not appear in the Juniper Express Engine tab:

- Scan mode
- Extension list
- Actions inapplicable to the Juniper Express Engine: Corrupt file, Password file, and Decompression Layer

8. Select **Apply** and then **OK**.

KAV Updater Support: From the NSM UI, you can update a Kaspersky Lab or Juniper Express Engine Pattern DB from a remote server that is pre-configured in the device. To run an update from the NSM UI, navigate from **Devices > AV Scan Manager > Update Pattern** and select IDP devices to be updated.

Creating an Antispam Profile

You can create an Antispam profile to specify the action to be taken with spam. You can bind custom features with security policies.

1. Select **Object Manager > UTM > Anti-Spam Profiles**.
2. NSM displays two tables: Predefined UTM AS Profiles on top and Custom UTM AS Profiles below. You can only view but not edit the listed predefined profiles. You can create and edit custom profiles.
3. Select **+** in the Custom UTM AS Profiles table. The **New Anti-Spam Profile** window opens
4. Enter a name for the profile.
5. Enter a comment or description.
6. Select a color from the drop-down list.
7. Enable Use default SBL.
8. Select an action: Block, Tag header, Tag subject.
9. Enter a tag string.
10. Select **OK**.

Creating a Content Filtering Profile

Content filtering allows you to specify the type of content to block. For example, you could block executable files such as .exe or .bin files that are prone to containing viruses.

1. Select **Object Manager > UTM > Content Filtering**.
2. Select **+** in the Custom UTM CF Profiles table. The **New Content Filtering Profile** window opens.
3. Enter a name for the profile.
4. Enter a comment or description.
5. Select a color from the drop-down list.
6. Set notification options: Notification type, Notify mail sender, and Custom message.
7. Select the type of content to block.
8. Set filters. You can select from existing lists or create new lists for each filter by clicking **+** beside the field.
 - Permitted command list
 - Block command list
 - Block extension list
 - Block mime list
 - Block mime list except
9. Select **OK**.

Creating a URL Filtering Profile

Based on your requirements, you can use URL filtering to prevent malicious or suspicious URLs from transferring their data. For example, you might wish to filter out gaming or entertainment sites.

1. Select **Object Manager > UTM > URL Filtering**. NSM displays two tables: Predefined UTM UF Profiles on top and Custom UTM UF Profiles below. You can only view but not edit the listed predefined profiles. You can create and edit custom profiles.
2. Select **+** in the **Custom UTM UF Profiles** table. The New URL Filtering Profile window opens.
3. Enter a name for the profile.

4. Enter a comment or description.
5. Select a color from the drop-down list.
6. Select the engine type. If you select Surf-control Integrated, set the following:
 - Default action: Block or permit.
 - Timeout period: In the range of 1-1800. Mouse over the field to see a tool tip with the allowed values.
 - Enter a deny message.
 - Set Fallback actions- either deny or permit- for the following: Default, Server Conn, Time out, Too many requests.
 - Select the list of categories to filter. You can edit the existing list or add a new one by clicking +. You can permit or deny each of the categories.

If you select Websense Redirect as the engine, set the following properties:

 - Enter server information: Host name, Port in the range 1024- 65535, Sockets in the range 1-8. Mouse over the field to see a tool tip with the allowed values.
 - Enter account name.
 - Select Timeout period: In the range of 1-1800.
 - Enter Deny message.
 - Set Fallback actions— either deny or permit— for the following: Default, Server Conn, Time out, Too many requests.
7. Select **OK**.

Miscellaneous UTM Features

The UT Manager provides miscellaneous features that support the main activities of threat management. These include:

- [Multipurpose Internet Mail Extension \(MIME\) Lists on page 408](#)
- [Extension Lists on page 408](#)
- [Command Lists on page 409](#)
- [URL Patterns on page 409](#)
- [URL Categories on page 410](#)

Multipurpose Internet Mail Extension (MIME) Lists

MIME lists contain the type of MIMEs that can be permitted or denied; for example, text or html. You can define the extensions that are to be permitted or denied by creating profiles. The maximum number of characters allowed in a MIME name are 29, in a MIME entry 40, and a MIME list 1023. The maximum of user defined MIME lists is system dependent. To create a MIME list:

1. Select **Object Manager > UTM > Misc > Mime List**. NSM displays two tables: Predefined UTM Mime List Profiles on top and Custom UTM Mime List Profiles below. You can only view but not edit the listed predefined profiles. You can create and edit custom profiles.
2. Select **+** in the **Custom UTM Mime List Profiles** table. The New Mime List Profile window opens.
3. Enter a name for the profile.
4. Enter a comment or description.
5. Select a color from the drop-down list.
6. Enter the multipurpose internet mail extensions for the profile.
7. Select **OK**.

Extension Lists

In an extension file list profile, you can specify various file extensions to be allowed or blocked; for example, .txt, .pdf, .exe. You can use these list profiles to create antivirus settings as well. The maximum number of characters allowed in an Extension name are 29, in an Extension entry 15, and an Extension list 255. The maximum number of user defined Extension lists is system dependent. To create a new Extension list profile:

1. Select **Object Manager > UTM > Misc > Extension List**. NSM displays two tables: Predefined UTM Extension List Profiles on top and Custom UTM Extension List Profiles below. You can only view but not edit the listed predefined profiles. You can create and edit custom profiles.
2. Select **+** in the Custom UTM Extension List Profiles table. The **New Extension List Profile** window opens.
3. Enter a name for the profile.
4. Enter a comment or description.

5. Select a color from the drop-down list.
6. Enter the extension types for the profile.
7. Select **OK**.

Command Lists

A command list defines various commands to be executed in the system for different protocols that you can permit or deny. You can create a new command list profile.

1. Select **Object Manager > UTM > Misc > Command List**.
2. Select **+**. The **New Command List Profile** window opens.
3. Enter a name for the profile.
4. Enter a comment or description.
5. Select a color from the drop-down list.
6. Enter the list of commands for the profile.
7. Select **OK**.

URL Patterns

You can create a URL pattern by listing different URLs to be permitted or denied. To create and view URL patterns:

1. Select **Object Manager > UTM > Misc > URL Patterns**. You can view all the URL patterns and create a new URL pattern.
2. Select **+**. The **New URL Pattern** window opens.
3. Enter a name for the profile.
4. Enter a comment or description.
5. Enter the URLs that make up the pattern.
6. Select **OK**.

URL Categories

A URL category is a list of URL Patterns that belong to the same category. To create a URL category:

1. Select **Object Manager > UTM > Misc > URL Category**. NSM displays two tables: Predefined UTM URL Categories on top and Custom UTM URL Categories below. You can only view but not edit the listed predefined categories. You can create and edit your own categories.
2. Select **+** in the Custom UTM URL Categories table. The **New URL Category** window opens.
3. Enter a name for the category.
4. Enter a comment or description.
5. View or edit entries in the **UTM Pattern Entries** box.
6. Select URL patterns from the list or add new patterns by clicking on **+**.
7. Select **OK**.

ScreenOS Threat Management Features

You can continue to use ScreenOS features to configure and manage AV and Web Filtering objects, as well as configure ICAP servers.

- [Configuring Antivirus Objects on page 410](#)
- [Configuring External AV Profiles on page 411](#)
- [Configuring Internal AV Profiles on page 412](#)
- [Configuring ICAP AV Servers and Profiles on page 413](#)
- [Configuring ICAP AV Profiles on page 414](#)
- [Configuring Web Filtering Objects on page 415](#)

Configuring Antivirus Objects

Security devices provide one or more of the following antivirus scanning methods:

- **External AV scanning**—This method forwards traffic to a Trend Micro device for scanning. (This option is not supported by devices running ScreenOS 5.3 or higher.) The security device forwards all traffic to be scanned to the Trend Micro device. To configure external AV scanning, use the AV Scanner settings (detailed below).
- **Internal AV scanning**—This method uses the AV scanner on the security device, and is not supported by all security devices. Internal scanning may be configured on a

per-device basis, or it may be configured via templates. This section describes how to create the templates.

- ICAP AV scanning—This method forwards traffic to an Internet Content Adaptation Protocol (ICAP) server for examination. To forward traffic to an ICAP server, create an ICAP server object, create an ICAP profile, and then specify that profile in a policy.

Configuring External AV Profiles

External AV profiles define the external Trend Micro AV scanner server that a security device uses to detect viruses in specific protocols. This feature describes the external scanner supported by ScreenOS 5.0 — 5.3. For ScreenOS 5.4 and later, use an ICAP AV profile as described in [“Configuring ICAP AV Profiles” on page 414](#)

You must configure an AV profile when using external AV for virus protection on your security device. After you have configured an AV profile, you can use the profile within a firewall rule.



NOTE: You can configure additional settings for external antivirus protection on the security device itself. For details, refer to *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

External AV profiles contain the following information:

- Server Name and Port—You must specify the IP address and port number of the external antivirus server that contains your virus definitions.
- Protocols and Timeouts—You must specify the protocols (HTTP and SMTP) that the external AV server scans for viruses. The default protocol timeout is 180 seconds, but you can edit this default to meet your networking requirements.

You must use the AV profile in a firewall rule and install that rule on a security device before the external scanner can begin inspecting traffic for viruses. For information about using AV profiles in rules.

In this example, you configure an AV profile that sends all HTTP traffic to an external antivirus server at 1.2.2.20 for virus checking. Because you anticipate heavy HTTP loads on the network, you increase the timeout from 180 seconds (the default setting) to 300 seconds.

1. In the main navigation tree, select **Object Manager>UTM>ScreenOS>AV Objects>External**.
2. In the main display area, click the Add icon. The New AntiVirus Profile dialog box appears.
3. Configure the following:

- For Name, scanner1_HTTP
 - For Server Name, enter 1.2.2.20.
 - For Server Port, leave the default port number of 3300.
4. Select HTTP, then configure the timeout as 300 seconds.
 5. Click **OK** to save the new profile.

Configuring Internal AV Profiles

Internal AV profiles allow you to set AV settings for multiple devices via a policy. NSM comes with a predefined AV profile, or you can create your own.

- [Viewing a Predefined Profile on page 412](#)
- [Creating Custom AV Profiles on page 412](#)
- [Configuring Extension Lists on page 413](#)

Viewing a Predefined Profile

To view the predefined AV profile, select **Object Manager > AV Objects > Internal > Predefined Profiles**. Select the predefined profile, then click the **Edit** button. You cannot make changes to the profile.

Creating Custom AV Profiles

To create a custom AV profile, select **Object Manager > UTM > ScreenOS > AV Objects > Internal > Predefined Profiles/Custom Profiles**. Click the **Add** button.

Populate the fields of the New Internal Antivirus Profile dialog:

- General information—Assign a name and color to the profile, and enter a comment describing the purpose of the profile.
- For each protocol type, check the Enable check box to enable scanning for that protocol, then set the following settings for each enabled protocol:
 - Scan Mode: All, Intelligent, or by File Extension. If you select Scan by File Extension, you must populate the Ext List Include field.
 - Scanning Timeout: Scans that take longer than this time out and are not completed.
 - Decompress Layer: The number of levels of decompression to perform before scanning. A decompression setting of 2 would all the scanning of a .zip file within a .zip file.
 - Skip Mime (HTTP only): If checked, causes the scanner to skip any mime types listed in the Mime List field.
 - Ext List Include: A list of file extensions to examine for viruses. Extension lists are created under Object Manager > AV Objects > Extension Lists.
 - Ext List Exclude: A list of file extensions to not examine for viruses. Extension lists are created under Object Manager > AV Objects > Extension Lists.

- **Mime List (HTTP only):** The list of mime types to not scan. NSM ships with a default mime type list, or you can create your own under **Object Manager > AV Objects > Custom Mime Lists**.
- **Email Notify Virus Sender (IMAP, POP3, SMTP only):** Notifies an e-mail sender if a virus was found in the e-mail.
- **Email Notify Scan-Error Sender (IMAP, POP3, SMTP only):** Notifies an e-mail sender if the e-mail was dropped due to a scan error.
- **Email Notify Scan-Error Recipient (IMAP, POP3, SMTP only):** Notifies an e-mail recipient if the e-mail was passed due to a scan error.

Configuring Extension Lists

You can configure AV profiles to scan (or not scan) files based on their file extension. File extension include lists and exclude lists are the same kind of lists. They become include or exclude lists depending on how they are added to a profile.

To create a file extension list object, select **Object Manager > UTM > ScreenOS > AV Objects > Internal > Extension lists**. Click **Add**.

Populate the following fields in the New Internal Antivirus Ext List dialog:

- **Name:** Give the extension list a descriptive name.
- **Color:** Assign a color, if desired.
- **Comment:** Provide a comment describing the list and its use.
- **File Extension:** Enter a list of file extensions, separated by semicolons. Example: `html;htm;jpg`.

Configuring ICAP AV Servers and Profiles

Before a security device can forward traffic to an ICAP AV server, you must create a server object in NSM. You can create multiple server objects and assign some or all of them to server groups. You can then assign this server object or server group to an AV profile, then assign that profile to a security policy.

To specify a server, you will need the following information:

- **Name:** The name of the ICAP server as it will appear in the NSM GUI.
- **Host:** The IP address of the ICAP server.
- **Port:** The ICAP server port. (Default: 1143)
- **Enable:** If selected, indicates that the server should be reachable and usable by a security device. Deselect this check box if the server is unavailable or should not be used by a security device.
- **Probe URL:** The path on the ICAP AV server to probe for availability.
- **Probe Interval (in seconds and multiples of 5):** Indicates how often the security device should check to see that the server is in service and available to process traffic. If this value is set to 0, then the security device will assume that the ICAP service is available

at all times. If it is set to a positive number of seconds, the security device will check the server's status at that interval. If the server returns as in-service, the security device will send it traffic. If it returns as out-of-service, the security device will not send traffic.

- **Maximum Connections:** The maximum number of TCP connections between the security device and the ICAP AV server.

To create a server group, first create the server objects that will go into it. However, you can create an empty group as a place holder.

You can assign a server to more than one server group.

Configuring ICAP AV Profiles

ICAP AV profiles, when assigned to a policy, redirect traffic to an ICAP AV server.

To create an ICAP AV Profile, select **Object Manager > UTM > ScreenOS > AV Objects > ICAP > Custom Profiles**. Click the add icon.

You need the following information:

- General information—Assign a name and color to the profile, and enter a comment describing the purpose of the profile.
- HTTP tab:
 - HTTP Enable: Selecting this check box in each tab turns on scanning for that protocol.
 - Skip Mime: If checked, causes the scanner to skip any mime types listed in the Mime List field.
 - Time out: Scans that take longer than this time out and are not completed.
 - Mime List: If Skip Mime is checked, also specify the MIME list that will be used for comparison. See [“Multipurpose Internet Mail Extension \(MIME\) Lists” on page 408](#) for information on creating MIME lists.
- SMTP tab:
 - SMTP Enable: Selecting this check box in each tab turns on scanning for that protocol.
 - Time out: Scans that take longer than this time out and are not completed.
 - Email Notification for Virus - Notify Sender: Notifies an e-mail sender if a virus was found in the e-mail.
 - Email Notification for Scan - Notify Sender: Notifies an e-mail sender if the e-mail was dropped due to a scan error.
 - Email Notification for Scan - Notify Recipient: Notifies an e-mail recipient if the e-mail was passed due to a scan error.
- ICAP tab:
 - ICAP Server/Server Group: Assign an ICAP AV server or server group to this profile. See [“Configuring ICAP AV Servers and Profiles” on page 413](#) for information on creating ICAP AV servers and server group objects in NSM.

- Request URL: The request URL on the ICAP AV server.
- Response URL: The response URL on the ICAP AV server.

Configuring Web Filtering Objects

Web Filtering (Integrated) enables you to create a Web Filtering profile for all of your security devices by binding the profile to the firewall rule. With a Web Filtering profile, the security device intercepts each HTTP request and determines whether to permit or block access to a requested website by categorizing the URL and matching the Web category to the Web Filtering profile. You can then bind the Web Filtering profile to the firewall rule.

To configure a security device for Web Filtering, you need to:

- Obtain a license key to enable the Web Filtering option on security devices.
- Configure at least one Domain Name Server (DNS) so the security device can resolve the SurfControl CPA server name to an address.
- Configure Web Filtering on the security device. For details, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.
- [Web Categories on page 415](#)
- [Custom Lists on page 415](#)
- [Predefined Categories on page 416](#)

Web Categories

A Web category is a list of URLs organized by content. There are two types of categories: Custom Lists and Predefined Categories.

Custom Lists

You can group URLs and create custom lists specific to your needs. You can include up to 20 URLs in each list. When you create a list, you can add either the URL or the IP address of a website. When you add a URL to a custom list, the security device performs a Domain Name Server (DNS) lookup, resolves the hostname into IP addresses and caches this information.

When a user tries to access a website by typing the IP address of the website, the security device checks the cached list of addresses and tries to resolve the hostname. It is important to enter both the URL and the IP address(es) of a website.



NOTE: When a URL exists in both a custom list and a predefined category, the security device matches the URL to the custom list first.

In this example you create a custom list called Competitors, Gaming.

1. In the main navigation tree, select **Object Manager > UTM > ScreenOS > Web Filtering (Integrated) > Web categories > Custom Lists**.
2. Click the Add icon. The New Web categories dialog box appears.
3. For Name, enter **Competitors, Gaming**.
4. Click the Add icon. The New URL Entries dialog box appears. Enter your configuration changes, then repeat to add a second URL Entry.
 - For the first URL entry, enter **www.games1.com** then click **OK**.
 - For the second URL entry, enter **www.games2.com** then click **OK**.
5. Click **OK** to save the new Custom List.

Predefined Categories

The security devices can use the predefined SurfControl Web categories to determine the category of a URL. SurfControl Content Portal Authority (CPA) servers maintain a large database of web content classified into approximately 40 categories.

To view the predefined SurfControl Web categories, select **Web Filtering (Integrated) > Web categories > Predefined Categories**.

Configuring Custom Policy Fields

Custom Policy Fields objects represent metadata information that you can store and use in a structured manner. Users can add custom objects to the policy table, such as ticket Number, vendor contact, and so on, for each rule in the rulebase. NSM provides a shared object to store these custom detail data while the table contains a column that corresponds to these custom details.

The **Custom Detail** column (visible in Expanded Mode and hidden in Compact Mode) captures the information about the rule, but does not push the information to the device. The column is able to display multiple shared objects in each cell.

This allows for a better filtering mechanism for the information, reduces data redundancy (in the case where all rules need to have the same e-mail address associated with them), and provides multiple properties for user's needs.

The custom detail object is user configurable. The metadata is designed to capture the following information about each object:

- **Name** -- Determines to which definition of the metadata the objects need to comply.
- **Required** -- Indicates if the metadata for a custom detailed object is defined with the "Required" option set to true, all rules in all rulebases that do not have a value selected for this MetaData displays the yellow warning triangle with a warning message. The policy will not be saved if no value is provided.
- **Validation String** -- A shared object definition in the metadata requires the user to select from a list of Custom Detail objects. A String definition in the metadata allows the user to enter a plain test string. Each instance supports no more than a single string value.

- **Field Type** -- If a regular expression is provided in the definition of an object, the custom detail object is validated against the regular expression. This is required and the custom object instance cannot be saved until this expression is satisfied.
- **Comments** -- This column allows the user to input any comments associated with the new object.

This information will be exported using the Policy Export tool, if the user selects Expanded Mode when exporting data.

Policy filtering is supported on individual values in the Custom Details column.

Defining Metadata

The metadata is defined using the Policy Details node located in the navigation tree. Users can see all metadata definitions as well as add, edit, or delete definitions. Existing metadata is displayed in table format and supported at the domain level. Definitions in the global domain are accessible in subdomains for creating objects that comply with the global domain.

Deleting a metadata definition forces all objects to comply with the definition and lists all usages of those objects. When deleting a metadata definition, all the objects complying that metadata are also deleted. In addition, it removes all usages of the changed objects from the security policy rules that referred to them.

Instantiating New Objects

As with metadata definitions, you can also create custom policy objection on the domain level. Objects you create in the global domain will be available for all subdomains, while objects created in the subdomains will only be available within the subdomain in which it was created.

When you delete an object, NSM displays all the usages of that object in the security policy rules, and will ask you for confirmation of the command. Once you confirm that you want to delete the object, NSM will remove all usages of the object you are deleting from the security policy rules that refer to the deleted object.

Adding Custom Detail Object to Rules

You can add custom detail objects to a rule in the policy using the same mechanism as other shared objects, such as service or address objects. You can use multiple selections for objects using the Shared data type. This allows you to add multiple objects complying with the same metadata. For example, you can add multiple e-mail addresses or phone number for each rule.

Once you have added custom objects to the rules, NSM displays the custom object along with the metadata name. For example, after adding an address to a rule, the value displayed in the rule could look like the following:

Email Address: admin@juniper.net

Requisition Number: JN0001

NSM will sort the entries in the Custom Details cell by the metadata name appended to the custom object value. NSM will copy and paste data in the Custom Details column along with other rule data when a rule is copied and pasted.

Objects with a String data type will provide a special edit dialog that allow you to change the string value contained within. The dialog allowing for this information is accessible by right-clicking on the selected value in the Context Menu. Objects with a Shared data type will have a special dialog that allows you to edit the value contained within. After saving the change, it is reflected in all rules using that object.

[Open Log Viewer](#)

You can open the Log Viewer from any rule in the policy. NSM will open the Log Viewer screen to display only those logs that were generated as a result of the selected rule. This option is available only if the policy has not been edited since the last time it was pushed to a device. Otherwise, the action is displayed as disabled. You can right-click the policy to access this option.

[Configuring GTP Objects](#)

To enable a security device to manage GTP traffic, you must create a GTP object and then apply it to a security policy rule. The rule with the GTP object defines how the device handles GTP packets: If a GTP packet matches the rule, the device attempts to further match the packet data with the parameters set in the GTP object.

For detailed information on GTP, refer to the Concepts & Examples ScreenOS Reference Guide, Volume 13: General Packet Radio Service.

Using GTP objects, you can configure multiple rules that enforce different GTP configurations in the same security policy. For example, you can configure a security policy that enables a device to control GTP traffic differently based on source and destination zones and addresses, action, and so on.

You configure GTP objects in the Object Manager. From the main navigation tree, select **Object Manager > GTP Objects**, then click the Add icon to display the New GTP Object configuration screens. For each object, you can configure the following settings:

- [“Configuring Info” on page 418](#)
- [“Configuring Traffic Logging and Counting” on page 420](#)
- [“Configuring IMSI Prefix and APN Filtering” on page 421](#)
- [“Configuring GTP Message Filtering” on page 423](#)
- [“Configuring Subscriber Tracing \(Lawful Interception\)” on page 423](#)

The following sections detail each GTP setting. For an example on creating a GTP object, see [“Example: Creating a GTP Object” on page 424](#).

[Configuring Info](#)

The Info settings define the basic properties of the GTP object, and specify how the security device should handle GTP messages and tunnels.

Limiting GTP Message Length

To limit the length of a GTP message, you can specify the minimum and maximum number of bytes permitted in a message length field. In the GTP header, the message length field indicates the length of the GTP payload. It does not include the length of the GTP header itself, the UDP header, or the IP header.

The default minimum and maximum GTP message lengths are 0 and 65535, respectively.

Limiting GTP Message Rate

To limit the rate of network traffic from a security device to a GPRS Support Node (GSN), you can specify the number of packets per second permitted for GTP-Control (GTP-C) messages.

Because GTP-C messages require processing and replying, they can overwhelm a GSN. Setting a rate limit on GTP-C messages can protect your GSNs from Denial-of-Service (DoS) attacks such as:

- Border Gateway bandwidth saturation—A malicious operator connected to the same GRX as your PLMN can generate enough network traffic directed at your Border Gateway, so that legitimate traffic is starved for bandwidth in or out of your PLMN, thus denying roaming access to or from your network.
- GTP flood—GTP traffic can flood a GSN, forcing it to spend its CPU cycles processing illegitimate data. This can prevent subscribers from roaming, forwarding data to external networks, or prevent a GPRS attach to the network.

To limit the GTP message rate, enable Limit (packets/second) and enter the maximum number of packets per second that a security device can send to a GSN (the default is unlimited).

Limiting GTP Tunnels

GSNs use GTP tunnels to transmit GTP traffic using the GPRS Tunneling Protocol (GTP). Because GSNs have a limited capacity for GTP tunnels, you might want to configure the security device to limit the number of GTP tunnels created.

To limit GTP tunnels, enable Limit (tunnels/GSN) and enter the maximum number of tunnels permitted for each GSN (the default is unlimited).

Removing Inactive GTP Tunnels

To configure a security device to detect and remove inactive GTP tunnels automatically, configure the GTP Tunnel Inactivity Timeout (hours). A GTP tunnel might hang (become inactive) when a “delete pdp context response” message gets lost on a network, or a GSN does not properly shut down.

The security device automatically removes a GTP tunnel that is idle for the specified timeout value. The default timeout value is 24 hours.

Validating Sequence Numbers

When using a security device between the GGSNs, you can configure the device to validate sequence numbers for the GGSN and drop out-of-sequence packets. This helps conserve GGSN resources by preventing the unnecessary processing of invalid packets.

The header of a GTP packet contains a Sequence Number field, which indicates the order of the packets arriving at the GGSN. During the PDP context activation stage:

- The sending GGSN uses zero (0) as the Sequence Number value for the first G-PDU it sends through a tunnel to another GGSN. The sending GGSN then increments the Sequence Number value for each following G-PDU it sends. The value resets to zero when it reaches 65535.
- The receiving GGSN sets its counter to zero. When it receives a valid G-PDU, it increments its counter by one. The counter resets to zero when it reaches 65535. The receiving GGSN compares the Sequence Number in the arriving packet with the sequence number in its counter: If the numbers correspond, the GGSN forwards the packet; if they differ, the GGSN drops the packet.

To enable the device to validate sequence numbers for the GGSN, enable Sequence Number Validation. By default, validation is disabled.

Filtering GTP-in-GTP Packets

To enable a security device to detect and drop a GTP packet that contains another GTP packet in its message body, enable GTP in GTP Denied.

Removing GTP R6 Informational Elements

GTP R6 contains additional Informational Elements (IEs) that support 3GPP networks: RAT, RAI, ULI, IMEI-SV, and APN Restriction. These new IEs are not supported on 2GPP networks. You can tell the firewall to strip out these elements when traffic passes from a 3GPP network to a 2GPP network.

To enable GTP traffic to flow between 3GPP and 2GPP networks, enable Remove r6 IE.

Inspecting Tunnel Endpoint IDs

You can configure the security device to perform Deep Inspection on the tunnel endpoint IDs (TEID) in G-PDU data messages.

To perform Deep Inspection on tunnel endpoint IDs, enable TEID DI.

Configuring Traffic Logging and Counting

When you enable traffic logging and counting for a GTP object, the security device generates log entries for deleted GTP tunnels and GTP traffic events.

Traffic Counting

A security device can count the number of user data and control messages (or bytes of data), received from and forwarded to the GGSNs and SGSNs that the device protects.

The device counts traffic for each GTP tunnel separately, and differentiates GTP-User and GTP-Control messages.

To enable counting, select Count By Message or Count By Byte. When counting is enabled and tunnel is deleted, the device counts and logs the total number of messages or bytes of data that it received from and forwarded to the SGSN or GGSN.

To view log entries for deleted GTP tunnels, use the Log Viewer.

Traffic Logging

A security device creates log entries for GTP events based on the status of the GTP packet. For each event type, you can also specify how much information (basic or extended) you want about each packet.

To configure GTP logging, select basic or extended for each GTP packet status:

- Log Forwarded Packets—When enabled, the device creates a log entry for each GTP packet that was transmitted because it was permitted by the security policy.
- Log Dropped Packet Due to Type/Length/Version—When enabled, the device creates a log entry for each GTP packet that was dropped because it was denied by the security policy.
- Log Dropped Packet Due to Invalid State—When enabled, the device creates a log entry for each GTP packet that was dropped because it failed stateful inspection.
- Log Dropped Packet Due to GSN Tunnel Limit—When enabled, the device creates a log entry for each GTP packet that was dropped because the maximum limit of GTP tunnels for the destination GSN was reached.
- Log Dropped Packet Due to GSN Rate Limit—When enabled, the device creates a log entry for each GTP packet that was dropped because the maximum rate limit of the destination GSN was reached.

You can also specify the frequency that a security device creates log entries for rate-limited messages. Setting a logging frequency conserves resources on the syslog server and security device, and can avoid a logging overflow of messages. By default, the frequency is 2, meaning the security device creates a log entry for every two messages above the set rate limit.

To view GTP traffic log entries, use the Log Viewer.

Configuring IMSI Prefix and APN Filtering

You can use the IMSI Prefix and APN to restrict access to a specific set of mobile subscribers.

Creating an APN Filter

An Access Point Name (APN) is included in the header of a GTP packet, and provides information on how to reach a network. By default, a security device permits all APNs. However, you can configure the device to filter APNs, enabling access only for those APNs you specify, and restricting roaming subscribers' access to external networks.

You can specify up to 2000 permitted APNs. When APN filtering is enabled, it applies only to “create pdp request” messages. For these messages to pass an APN filter, the GTP packet must match both the APN name filter and the Selection Mode filter:

- **APN Domain Name filter**—The device attempts to match the APN in a GTP packet to the APNs set in the GTP object. If the two APNs match, the device passes the packet to the selection mode filter.
- **Selection Mode Filter**—The device attempts to match the Selection Mode for the GTP packet and the GTP object. If the two modes match, the device forwards the GTP packet; if the modes do not match, the device drops the GTP packet.

Additionally, you can filter GTP packets based on the combination of an IMSI prefix and an APN. For details, see [“Creating an IMSI Prefix Filter” on page 422](#).

Setting the Network ID (APN Domain Name)

To set an APN filter, you need to know the network ID, which identifies the name of an external network.



NOTE: Because the APN domain name (network ID) can potentially be very long and contain many characters, you can use the wildcard “*” as the first character of the APN to indicate that the APN also includes all preceding characters. However, because APN filtering is based on perfect matches, using the wildcard “*” can prevent the inadvertent exclusion of APNs that you would otherwise authorize.

Setting a Selection Mode

You must also set a Selection Mode, which indicates the origin of the APN and if the user subscription has been verified by the Home Location Register (HLR). You can set one of the following Selection Modes:

- **Mobile Station**—MS-provided APN, subscription not verified. This Selection Mode indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user’s subscription to the network.
- **Network**—Network-provided APN, subscription not verified. This Selection Mode indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user’s subscription to the network.
- **Verified**—MS or Network-provided APN, subscription verified. This Selection Mode indicates that the MS or the network provided the APN and that the HLR verified the user’s subscription to the network.

Creating an IMSI Prefix Filter

A GSN (GPRS Support Node) identifies a mobile station by its IMSI (International Mobile Station Identity). An IMSI is composed of three elements:

- The MCC (Mobile Country Code)

- The MNC (Mobile Network Code)
- The MSIN (Mobile Subscriber Identification Number)

The MCC and MNC combine to create the IMSI prefix, which identifies the mobile subscriber's home network (PLMN). By default, a security device does not perform IMSI prefix filtering on GTP packets. You can use the IMSI prefix to configure a security device to deny GTP traffic sent from non-roaming partners.

When you set an IMSI prefix in the GTP object, the security device filters “create pdp request” messages and permits only GTP packets with a matching IMSI prefix. If the prefix does not match, the security device drops the GTP packet. You can set up to 1000 IMSI prefixes for each device (one per each filter).

To disable IMSI prefix filtering, remove all MCC-MNC pairs from the GTP object.

Configuring GTP Message Filtering

By default, the security device permits all GTP message types. You can configure a security device to filter GTP packets and drop them based on their message type.

A GTP message type includes one or many messages. When you drop a message type, you automatically drop all messages of the specified type. For example, if you select to drop the **sgsn-context** message type, you also drop “sgsn context request”, “sgsn context response”, and “sgsn context acknowledge” messages.

You drop message types based on the GTP version number, enabling you to drop message types for one version and permit them for another version.

Configuring Subscriber Tracing (Lawful Interception)

You can configure a security device to identify subscribers based on IMSI prefixes or Mobile Station-Integrated Services Data Network (MS-ISDN) identification, then log the contents of their GTP-User Data (GTP-U) or GTP-Control (GTP-C) messages.

To enable subscriber tracing, you must configure the following:

- Set Subscribers—Set the number of subscribers that the security device actively traces concurrently. The default number of simultaneous active traces is three (3).
- Specify Log Bytes—Specify the number of bytes of data to log for a GTP-U packet. The default value is zero, meaning that the device does not log any content from a GTP-U packet. When you enter a number other than zero, the security device sends the logged packets to an external server (such as Syslog) dedicated to Lawful Interception operations.
- Set ID—For each subscriber you want to trace, enter their ID number and select Based on IMSI or Based on MSISDN.

Example: Creating a GTP Object

1. In Object Manager, select **GTP Objects**, then click the Add icon in the main display area. The New GTP Object dialog box appears.
2. In the Info tab, configure the following settings:
 - For Name, enter **GPRS1**, then enter a color and comment for the object.
 - Select **Sequence Number Validation**.
 - Select **GTP in GTP Denied**.
 - Leave all other defaults.
3. In the GTP navigation tree, select **Traffic Logging/Counting**. Configure the following:
 - For Traffic Counters, select **Count by Message**.
 - Select **Basic** for the following message types: Log Forwarded Packets, Log Dropped Packet Due to Type/Length/Version, and Log Dropped Packet Due to Invalid State.
 - Leave all other defaults.
4. In the GTP navigation tree, select **IMSI Prefix and APN Filtering**. Click the Add icon to display a new IMSI Prefix and APN Filter Entry dialog box. Configure the following, then click **OK**:
 - For APN, enter **mobiphone.com.mnc123.mcc456.gprs**.
 - Select MCC-MNC and enter the code **24656**.
 - For Selection Mode, select **Mobile Station, Network, and Verified**.
5. In the GTP navigation tree, select **Subscriber Tracing**.
 - For Maximum Number of Simultaneous Active Traces, enter **2**.
 - For Number of Bytes to Be Saved to Log, enter **1020**.
6. Click the Add icon to display a New Subscriber ID dialog box. Configure the following, then click **OK**:
 - For ID, enter **345678**.
 - For ID Type, select **Based on IMSI**.
7. Click **OK** to save the new Subscriber ID, then click **OK** to save the GPRS1 object.

Configuring Service Objects

Service objects represent the IP traffic types for existing protocol standards. Security devices monitor and manage network traffic using these protocols. NSM includes predefined service objects for most standard services. You can also create custom service objects to represent services that are not included in the list of predefined service objects, or to represent a custom service running on your network.

You use service objects to create protected resources and specify the type of service within a security policy:

- In a protected resource, select a service or group of services to define the types of traffic you are permitting to and from the resource.
- In individual rules within a firewall or IDP rulebase, select one or more services or groups of services to define the types of IP traffic to which the rule applies. The action of the rule applies when the security device detects packets that use the specified service type.

Viewing Predefined Services

You can view predefined services in a tree or table format. The Service Tree displays services in a tree format, with service groups and individual services. The Service [Table 41 on page 425](#) displays services in a table format, and includes the following details:

Table 41: Service Table Tab Information

Name	Name of the service object
Type	Type of the service object: service or group
Timeout	Service timeout—inactivity timeout after a which a session on a security device is removed
Category	Classification based on the purpose the service is designed for: <ul style="list-style-type: none"> • email—used for sending and receiving e-mail (POP3, for example) • info seeking—used to retrieve specific information from a server (DNS, for example) • remote—used for accessing remote servers (Telnet, for example) • security—enable the access of a remote server securely using well known security mechanisms (HTTPS for example) • other— all other services
Non-ICMP Src Port	The TCP and UDP source port for the service. This column displays a list of IP protocols.
Non-ICMP Dst Port	The TCP and UDP destination port for the service. This column displays a list of IP protocols.
Comment	Contains optional comments.

To view service object properties, double-click a service object. In addition to the service name, category, and service timeout value, you can view the following service settings:

- For Non-ICMP services, the service object displays the protocol ID, source port range, and destination port range.
- For ICMP services, the General tab displays the Internet Control Message Protocol (ICMP) type and code.

- For Sun-RPC services, the Sun-RPC tab displays the Sun Microsystems program identifiers. Sun Remote Procedure Call (Sun-RPC), also known as Open Network Computing (ONC) RPC, enables a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.



NOTE: The transport address is comprised of the port number of the server, the program ID, and the version number.

NSM and security devices support 13 Sun-RPC predefined services. To permit or deny all Sun-RPC requests, include the Sun-RPC-Any service in a firewall or IDP rule; to permit or deny a Sun-RPC request by specific program number, include that service (or create a custom service) in the rule.

- For MS-RPC services, the MS-RPC tab displays the Microsoft universal unique identifiers (UUIDs). Microsoft Remote Procedure Call (MS-RPC) is the Microsoft implementation of the Distributed Computing Environment (DCE) RPC. Like the Sun-RPC, MS-RPC enables a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service program's Universal Unique Identifier (UUID).

NSM and security devices support 27 MS-RPC predefined services and 3 MS-RPC predefined service groups. To permit or deny all MS-RPC requests, include the MS-RPC-Any service in a firewall or IDP rule; to permit or deny an MS-RPC request by specific UUID, include that service (or create a custom service) in the rule.

You can view details for a predefined service object, but you cannot edit that service object.

Creating Custom Services

You can create custom service objects to represent protocols that are not included in the predefined services or to meet the unique needs of your network.



NOTE: Sun-RPC protocols and regular TCP/UDP/ICMP protocols cannot be in the same service object. MS-RPC protocols and regular TCP/UDP/ICMP protocols cannot be in the same service object

To add a service object, in the Object Manager, select **Service Objects > Custom Service Objects**. In the main display area, click the Add icon and select **Service** to display the New Service dialog box. Configure the following parameters:

- Name—Enter a name for the service.
- Timeout—Select the session timeout after which an inactive session is removed.

- Never. The session does not time out.
- Default. Use the default timeout for the selected protocol. The default timeout for TCP connections is 30 minutes. The default timeout for UDP connections is 1 minute.
- User-defined. Enter a session timeout value. The maximum timeout value for TCP and UDP connections is 2160 minutes.
- Color—Select a color to represent this service object in the NSM UI.
- Comment—Add a comment, if desired.
- Add the service entry:
 - For ICMP services, in the General tab click the Add icon. Enter the ICMP type and code, then click **OK**. For information about ICMP type, see the *NSM Online Help*.



NOTE: For Junos OS, ICMP-based service objects should not be created under **Non-ICMP Service Entries**.

- For Sun-RPC services, select the Sun-RPC tab, then click the Add icon. Enter high and low program identifiers, then click **OK**. You can add up to eight program ranges; ensure that the Program High value is greater than or equal to the Program Low value.



NOTE: For the complete list of the Sun Microsystems Program IDs and Microsoft UUIDs, refer to the ScreenOS online Help.

- For MS-RPC services, select the MS-RPC tab, then click the Add icon. Enter a UUID, then click **OK**. A UUID is 36 characters.
- For other non-ICMP services, in the NON-ICMP Service Entries area, click the Add icon. Select the protocol type and configure the source and destination ports, then click **OK**. To create a service object that uses multiple ports for the same service, add two service entries with different ports.

Service Object Groups

You can group services together as a service object group, then use that group in security policies and VPNs to simplify administration. Each service object can be referenced by multiple service object groups. Service object groups can contain both predefined and custom service objects, as well as other service object groups.

To add a service object group:

1. In the navigation tree, select **Object Manager > Service Objects**.
2. In the main display area, click the Add icon and select **Group**. The New Service Group dialog box appears.
3. Enter a name, color, and comment for the service object group.



NOTE: Service object group names cannot be the same as service object names.

4. In the Non-members area, select the service objects or service object groups you want to add to the group (hold Ctrl to select multiple objects), then click **Add**.



NOTE: You can drag service objects into and out of service groups from the main service tree.

5. Click **OK**.

The new service object group appears in the Service Tree and Service Table tabs.

Example: Creating a Custom Service and Group

In this example, you create a custom service object to represent the Ident service and a custom service group that includes this service.

To create the custom Ident service:

1. In the main navigation tree, select **Object Manager > Service Objects > Custom Service Objects**.
2. In the main display area, click the Add icon and select **Service**. The New Service dialog box appears.
3. Configure the following:
 - a. For Name, enter **Ident**
 - b. For Timeout, select **Default**.
 - c. For Color, select **blue**.
 - d. Enter a comment, if desired.
4. In the Non-ICMP Services Entries area, click the Add icon and select TCP. The New Service Entry dialog box appears. Configure the following:
 - a. For Source Port, select **Range**.
 - b. For Source Port Range, enter 0 to 65535.
 - c. For Destination Port, select **Specific**.
 - d. For Specific Port, enter **113**.
5. Click **OK** to save the new service entry, then click **OK** again to save the new service object.
6. In the main display area, click the Add icon and select **Group**. The New Service Group dialog box appears. Configure the following:

- a. For Name, enter **Remote Mail**.
- b. For Color, select **pink**.
- c. Enter a comment, if desired.
- d. In the Non-members area, select the following services (press and hold Ctrl to select multiple services):
 - FTP
 - HTTP
 - Ident
 - MAIL
 - POP3
 - TELNET
- e. Click **Add** to add the services as members of the group, then click **OK** to save the new service group.

Example: Creating a Custom Sun-RPC Service

In this example, you create a service object called my-sunrpc-nfs to use the Sun RPC Network File System, which is identified by two Program IDs: 100003 and 100227. Because Sun-RPC services use dynamically negotiated ports, you cannot use regular service objects based on fixed TCP/UDP ports to permit them in security policy. Instead, you must create sun rpc service objects using program numbers. For example, NFS uses two program numbers: 100003 and 100227. The corresponding TCP/UDP ports are dynamic. To permit the program numbers, you create a sun-rpc-nfs service object that contains these two numbers. The ALG maps the program numbers into dynamically negotiated TCP/UDP ports, and permits or denies the service based on a policy you configure.

To create the Sun-RPC service:

1. In the main navigation tree, select **Object Manager > Service Objects > Custom Service Objects**.
2. In the main display area, click the Add icon and select **Service**. The New Service dialog box appears.
3. Configure the following:
 - For Name, enter **my-sunrpc-nfs**
 - For Timeout, select **Default**.
 - For Color, select **blue**.
 - Enter a comment, if desired.
 - Select the Sun-RPC tab.
4. Configure the first service entry. Click the Add icon to display the New Service Entry dialog box, configure the following, then click **OK**:

- For Program Low, enter **100003**.
 - For Program High, enter **100003**.
5. Configure the second service entry. Click the Add icon to display the New Service Entry dialog box, configure the following, then click **OK**:
- For Program Low, enter **100227**.
 - For Program High, enter **100227**.
 - Click **OK** again to save the new service object.

Example: Creating a Custom MS-RPC Service

In this example, you create a service object called `my-ex-info-store` that includes the UUIDs for the MS Exchange Info Store service. Because MS RPC services use dynamically negotiated ports, you can not use regular service objects based on fixed TCP/UDP ports to permit them in a security policy. Instead, you must create MS RPC service objects using UUIDs. The MS Exchange Info Store service, for example, uses the following four UUIDs:

- 0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde
- 1453c42c-0fa6-11d2-a910-00c04f990f3b
- 10f24e8e-0fa6-11d2-a910-00c04f990f3b
- 1544f5e0-613c-11d1-93df-00c04fd7bd09

The corresponding TCP/UDP ports are dynamic. To permit them, you create an `ms-exchange-info-store` service object that contains these four UUIDs. The ALG maps the program numbers into dynamically negotiated TCP/UDP ports based on these four UUIDs, and permits or denies the service based on a rule you configure.

To create the MS-RPC service:

1. In the main navigation tree, select **Object Manager > Service Objects > Custom Service Objects**.
2. In the main display area, click the Add icon and select **Service**. The New Service dialog box appears.
3. Configure the following:
 - a. For Name, enter **my-ex-info-store..**
 - b. For Timeout, select **Default**.
 - c. For Color, select **blue**.
 - d. Enter a comment, if desired.
4. Select the MS-RPC tab. Configure a service entry for each of the following UUIDs:

- 0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde
- 1453c42c-0fa6-11d2-a910-00c04f990f3b
- 10f24e8e-0fa6-11d2-a910-00c04f990f3b
- 1544f5e0-613c-11d1-93df-00c04fd7bd09

5. Click **OK** to save the new service object.

Editing and Deleting Service Objects

You can edit a service object by right-clicking on the object and selecting **Edit**. You can also delete a service object by right-clicking on the object and selecting **Delete**. For more information on editing and deleting service objects, refer to the NSM Online Help.

Replacing Service Objects

You can replace a service object by right-clicking on the object and selecting **Replace With**. Replacing service objects simplifies making redundant changes to a service object that is referenced in multiple security policies. If you have permission to view the global domain objects for the objects you are replacing, then all objects for the selected category in the current domain and the global domain are displayed in the Replace With wizard, but the object to be replaced is not shown. When replacing service objects however, keep the following in mind:

- There is no validation check when replacing service objects; an error appears for any service objects that are not valid for specific policies. For example, you cannot assign a SUN-RPC-ANY service object to an IDP policy.
- You cannot replace a service object with a service group object that contains the replaced service object.
- You cannot undo or roll back a Replace With operation.



NOTE: Replacing service objects only applies to those objects in the domain in which you are working. Custom Services created in the global domain are not available for Replace With operations in subdomains.

After replacing service objects, it is good practice to check your security policies for any errors that may result. You can always edit or remove any duplicate objects in the security policy.

In this example, you want to replace all references to HTTP with HTTPS in your security policies.

To replace HTTP with HTTPS:

1. In the navigation tree, select the Object Manager and click **Service Objects** to open the service object tree.
2. Click on **Predefined Service Objects**.

3. In the Service Tree or Service Table, right-click on the HTTP service object and select **Replace With**. The Replace With wizard appears displaying a list of objects you can replace the selected service object with.
4. Select the HTTPS service object. Click **Next**. The wizard next displays the objects affected by the Replace With operation.

As an optional step, you can delete any replaced custom service objects by clicking on them and then selecting **Delete Replaced Object**.



NOTE: You cannot delete a predefined service object.

5. Click **Finish**.

Configuring SCTP Objects

With Stream Control Transmission Protocol (SCTP), you can transmit data in messages to the SCTP transport layer. Various protocols including IUA, SUA, M2UA, M3UA, H.248, and DIAMETER, can run on SCTP. You can control the protocols used with the SCTP protocol filtering tool. After you configure an SCTP object, you can apply it to various policies.

Configuring an SCTP Object

1. Select **Object Manager > SCTP Objects**. The **SCTP Objects** window opens. You can add, edit, delete or search for an SCTP object using the icons in the task bar at the upper left of the window.
2. Click the **New (+)** icon. A **New SCTP Object window** opens.
3. Enter a name for the object in the **Name** field.
4. Check the appropriate boxes to drop payload-protocols.
5. Check the appropriate boxes to drop m3ua-services.
6. Click **OK**.

Configuring Authentication Servers

An authentication server provides authentication services for NSM administrators and remote access services (RAS) users on your network. The information stored in an authentication server determines the privileges of each administrator.

When the security device receives a connection request that requires authentication verification, the device requests an authentication check from the external auth server specified in the policy, L2TP tunnel configuration, or IKE gateway configuration. The

device then acts as a relay between the user requesting authentication and the authentication server granting authentication.

In NSM, an auth server is an object used in security policies, IKE gateways, and L2TP tunnels. Each security device includes a default authentication server; however, to enable an external RADIUS, SecureID, or LDAP server to provide authentication, you must configure an external authentication server object. You can also configure a RADIUS authentication server object to provide authentication for the global domain and each subdomain. For information about configuring a RADIUS server, see [“Configuring a RADIUS Authentication Server” on page 435](#).



NOTE: You must also define routes that direct authentication requests to the RADIUS, SecurID, and LDAP servers.

To configure general authentication server object properties, in the main navigation tree, select **Object Manager > Authentication Servers** then click the Add icon. The General, Redundancy, and Identity tabs are the same for all server types; in the Server Type tab, select the authentication server type (RADIUS, SecureID, LDAP) to configure specific settings for that server type.

Configuring General Authentication Server Settings

In the General tab, configure a name, color, and comment that uniquely identify the object, then specify the IP address of the main authentication server; this is the IP address of the server that handles authentication requests.

You can also configure an authentication timeout (default is 10 minutes) to control the number of minutes before an authentication check times out. Timeouts affect the following user types differently:

- **Auth user**—The timeout countdown begins after the first authenticated session completes. If users initiate a new session before the countdown reaches the timeout threshold, they do not need to reauthenticate and the timeout countdown resets. The default timeout value is 10 minutes, and the maximum is 255 minutes. You can also set the timeout value at 0 so that the authentication period never times out.
- **Admin user**—If the length of idle time reaches the timeout threshold, the security device terminates the administrator session. To continue managing the device, the administrator must reconnect to the device and re authenticate. The default timeout value is 10 minutes, and the maximum is 1000 minutes. You can also set the timeout value at 0 so that an administrator session never times out.



NOTE: User authentication timeout is not the same as session idle timeout. If no activity occurs in a session for a predefined length of time, the security device automatically removes the session from its session table.

Configuring Authentication Server Redundancy

In the Redundancy tab, you can configure backup server to handle authentication requests if the primary server fails. For RADIUS servers only, you can also configure a secondary backup server (this option is not supported for SecureID servers).

For RADIUS and LDAP servers only, you can configure a Failover Revert Interval that determines how long the device uses a backup server before attempting to use the primary server again. To configure the interval, enter the number of seconds (1 to 86400); to disable the failover revert, set the interval to 0 (the device continues to use the backup server indefinitely). The interval countdown begins when the device fails over from the primary auth server to the backup or secondary backup server (RADIUS only).

Configuring Authentication for User Types

In the Identity tab, configure the user types that the authentication server supports:

- Admin Users
- Firewall Auth Users
- XAuth Users
- 802.1X Users
- L2TP Users

For RADIUS servers, you can also configure the optional domain name checking and domain name stripping settings, as detailed in the following sections.

Domain Name Checking

Use domain name checking to authenticate users from a specific domain. This setting is optional and is not required to configure a RADIUS authentication server.

To configure, for Domain to Check In Username, enter the domain name (up to 45 characters). For each user authenticating to the server, the server compares the domain name in the username to specified domain (the domain is read as a string from right to left to the first @ character).

To authenticate usernames from all domains, leave this option unconfigured (blank).

Domain Name Stripping

Use domain name stripping to remove the domain name from usernames before sending to the authentication server. This setting is optional and is not required to configure a RADIUS authentication server. However, you might need to configure this setting when implementing a new RADIUS server with an existing network and established usernames.

To configure:

- For Separator Character, enter the separator character used in the usernames.
- For Separator Character Occurrence, enter the number of times (0 to 10) the separator character occurs in the username.

When a user attempts to authenticate, the device examines the username from right to left, then strips domain name information for the specified number of separator characters before sending the username onto the authentication server.

For example, when the Separator Character is @ and the Separator Character Occurrence is 2, the device handles the username `user1@mygrp.abc@myco.com` by stripping the characters `@mygrp.abc@myco.com` and sending only the characters `user1` to the authentication server.

If the device does not locate the separator character in the username, it does not strip the domain name from the username (usernames are passed to the authentication server as-is). Conversely, if the number of specified separator characters exceeds the number of separators found in a username, the device strips domain name information to the number of separators found (when reading right to left).

Configuring Authentication Server Types

In the Server Type tab, select the authentication server type (RADIUS, SecureID, LDAP) to configure specific settings for that server type:

- For RADIUS, see [“Configuring a RADIUS Authentication Server” on page 435](#).
- For SecureID, see [“Configuring a SecurID Authentication Server” on page 439](#)
- For LDAP, see [“Configuring a RADIUS Authentication Server” on page 435](#)

Configuring a RADIUS Authentication Server

The Remote Authentication Dial-In User Service (RADIUS) is a protocol for an authentication server that can support up to tens of thousands of users. The security device acts as a RADIUS client that authenticates users. When users log in, the RADIUS client (the security device) prompts them for their user name and password, then compares these values with the values stored in the RADIUS database. If the values match, the RADIUS client authenticates the user and permits access to the appropriate network services.

For a RADIUS authentication server object, configure the following:

- **RADIUS Port**—The port number on the RADIUS server to which a security device sends authentication requests. The default port number is 1645.
- **RADIUS Secret**—The secret (password) shared between a security device and the RADIUS server. The RADIUS server uses the shared secret to generate a key to encrypt traffic between the security device and the RADIUS server. The security device uses the shared secret to encrypt the user’s password that it sends to the RADIUS server.
- **RADIUS Retry Timeout**—The interval (in seconds) that a security device waits before sending another authentication request to the RADIUS server if the previous request does not elicit a response. The default is three seconds.
- **RADIUS Retries**—The number of unanswered requests (access and accounting) that a security device sends before it considers the RADIUS server unreachable and fails over to a backup server. To configure, enter the number of retries (1 to 20); the default is three.

- **RADIUS Compatible with RFC 2138**—When selected, enables the authentication server to comply with RFC 2138, an older RADIUS standard, with the following considerations:
 - For operations where RFC 2865/66 and RFC 2138 are mutually exclusive, the server complies with RFC 2138 only.
 - For operations where RFC 2865/66 and RFC 2138 are both supported, the server complies with all three RFCs.

When unselected (default), the server is compatible only with the current RADIUS standards RFC 2865 and 2866.

- **Enable Sending Calling-Station-ID**—When selected, the security device transmits the calling station ID within the access or accounting request to the RADIUS authentication server. Because the ID identifies the originator of the call (either the IKE IP address for XAuth or the phone number of the user originating the call), you might not want to send this information to the server. By default, this option is disabled; the device does not send the calling station ID to the server.
- **Length of Account Session ID Attribute**—The byte length of the account-session-id, which uniquely identifies the accounting session. By default, the byte length is 11, and follows the format NS-xxxxxxx. Because some RADIUS servers do not properly accept an 11-byte account session ID, you might want to configure a lower byte length that does not include the "NS-" prefix. To configure, enter a byte length from 6 to 10.
- **Separation of Authentication and Accounting Functions** — In the XAUTH and L2TP authentication process, RADIUS Accounting was coupled with RADIUS authentication resulting in a few issues caused by unavailability of the server's accounting service or network topology policy limitations leading to aborted authentication processes even if correct information was provided. You can separate the authentication and accounting functions by specifying different RADIUS Authentication and Accounting servers. In ScreenOS devices running 6.2 and later, you can enable or disable the accounting function, but not the authentication function. You can configure the RADIUS server accounting port as a value in the range of 1024 - 65535. From the NSM UI,
 - From Edit device > VPN Settings > Defaults, configure the following in the XAuth and L2TP sections: Default Accounting Server from the drop-down list, and Disable Default Accounting checkbox.
 - From Edit device > VPN Settings > Gateway Entry, configure the following in the IKE IDs/XAuth tab: Accounting Server Name from the drop-down list, and Disable Accounting checkbox.
 - From Edit device > VPN Settings > L2TP Entry, configure the following in the Auth Server > Use Custom Settings: Accounting Server Name from the drop-down list, and Disable Accounting checkbox.
 - From Edit Device > VPN Settings > L2TP Entry, configure the following in Accounting Settings: Select Accounting server name from the drop down list, and Disable Accounting checkbox.

Supported User Types

A RADIUS server supports the following user types:

- Auth users
- L2TP users (authentication and remote settings)
- XAuth users (authentication and remote settings)
- Admin users (authentication and privilege assignments)
- User groups

A RADIUS server **does not** support IKE users.

RADIUS Access-Challenge

When a user attempts to log in using telnet, a security device can process access-challenge packets from an external RADIUS server. Access-challenge is an additional authentication level. After a username and password has been authenticated, the RADIUS server sends an access-challenge to the security device, which forwards the challenge to the user. When the user replies, the device sends a new access-request with the user's response to the RADIUS server; if the user's response is correct, the authentication process concludes successfully.



NOTE: Juniper Networks does not support access-challenge with L2TP.

Juniper Networks Dictionary File

A dictionary file defines vendor-specific attributes (VSAs) that you load onto a RADIUS server. After you define the VSA values, the security device can query those values when a user logs on to the device.

You must load a Juniper Networks dictionary file to enable the RADIUS server to support NSM-specific attributes as administrator privileges, user groups, and remote L2TP and XAuth IP address, and DNS and WINS server address assignments. You **do not need** to load Juniper Networks dictionary file to enable RADIUS to make IP address assignments (Juniper Networks uses the standard RADIUS attribute for IP address assignments).

Juniper Networks provides two dictionary files: one for Funk Software RADIUS servers and one for Cisco RADIUS servers:

- For Funk Software RADIUS server dictionary file, go to http://www.juniper.net/customers/csc/research/hetscreen_kb/downloads/dictionary/funk_radius.zip
- For Cisco RADIUS server dictionary file, go to http://www.juniper.net/customers/csc/research/hetscreen_kb/downloads/dictionary/cisco_radius.zip

If using a Microsoft RADIUS server, there is no dictionary file. You must configure it as outlined in *Using a Windows NT Domain / Active Directory for User Authentication Security Devices*, which you can download from the Juniper customer support site.

Each Juniper Networks dictionary file contains the following specific information:

- **Vendor ID**—The Juniper Networks vendor ID (VID; also called an “IETF number”) is 3224. The VID identifies a specific vendor for a particular attribute. Some types of RADIUS server require you to enter the VID for each attribute entry, while other types only require you to enter it once and then apply it globally. Refer to your RADIUS server documentation for further information.
- **Attribute Name**—The attribute names describe individual NSM-specific attributes, such as NS-Admin-Privilege, NS-User-Group, and NS-Primary-DNS-Server.
- **Attribute Number**—The attribute number identifies an individual vendor-specific attribute.
- **Attribute Type**—The attribute type identifies the form in which attribute data (or “value”) appears—a string, an IP address, or an integer.

The RADIUS server automatically receives the above information when you load the Juniper Networks dictionary file onto it. To make new data entries, you must manually enter a value in the form indicated by the attribute type.

Example: Configuring a Radius Auth Server

In the following example, you define an auth server object for a RADIUS server. You specify its user account types as auth, L2TP, and XAuth. You name the RADIUS server “radius1” and accept the ID number that the security device automatically assigns it. You enter its IP address, which is 10.20.1.100; and change its port number from the default port number (1645) to 4500. You define its shared secret as “A56htYY97kl”. You change the authentication timeout value from the default (10 minutes) to 30 minutes and the RADIUS retry timeout from 3 seconds to 4 seconds. You also assign its two backup servers the IP addresses 10.20.1.110 and 10.20.1.120.

In addition, you load the Juniper Networks dictionary file on the RADIUS server so that it can support queries for the following vendor-specific attributes (VSAs): user groups, administrator privileges, remote L2TP and XAuth settings.

1. In the main navigation tree, select **Object Manager > Authentication Servers** and click the Add icon. Enter a name, color, and comment for the authentication server.
2. Configure the RADIUS servers:
 - For Main Server, enter the IP **10.20.1.100**
 - For Primary Backup Server, enter IP **10.20.1.110**
 - For Secondary Backup Server, enter IP **10.20.1.120**
3. For timeout, enter **30**.
4. Select the following:
 - **For Firewall Auth Users**
 - **For XAuth Users**
 - **For L2TP Users**
5. For Server Type, select **RADIUS**.

6. Configure the RADIUS server properties:
 - For server port, enter **4500** (default is 1645)
 - For secret, enter **A56hYY97kl**
 - For retry timeout, select **4**.
7. Click **OK** to save the RADIUS authentication server object.
8. Load the Juniper Networks dictionary file on the RADIUS server.

Configuring a SecurID Authentication Server

Security devices also support the RSA SecurID system. The device acts as a SecurID client, forwarding authentication requests to the external server for approval and relaying login information between the user and the server. Each SecurID user has three authentication credentials:

- User Name
- Personal identification number (PIN)
- Authenticator—a SecurID issued device with an LCD screen that displays a token code, a randomly generated string of numbers that changes every minute. The authenticator uses an algorithm known only by RSA to create the token code that appears in LCD screen; when users enter their username, their PIN, and the token code from their authenticator, the RSA ACE server also performs the same algorithm, generating a match between the server and the user.

When users log in, the SecurID client (the security device) prompts them for their user name, their PIN, and the current token code. The device compares the user input against value generated by the RSA ACE server algorithm. If the values match, the authentication is successful.

For a SecurID authentication server object, you must configure the following:

- Authentication Port—The port number on the SecurID ACE server to which the security device sends authentication requests. The default port number is 5500.
- Encryption Type—The algorithm used for encrypting communication between the security device and the SecurID ACE server (SDI or DES).
- Client Retries—The number of times that the SecurID client (the security device) tries to establish communication with the SecurID ACE server before aborting the attempt.
- Client Timeout—The length of time in seconds that the security device waits between authentication retry attempts.
- Use Duress—An option that prevents or allows use of a different PIN number. When this option is enabled, and a user enters a previously determined duress PIN number, the security device sends a signal to the SecurID ACE server, indicating that the user is performing the login against his or her will, possible under duress. The SecurID ACE server permits access that one time, then denies any further login attempts by that user until he or she contacts the SecurID administrator. Duress mode is available only if the SecurID ACE server supports this option.

Supported Users

A SecurID Ace server supports the following types of users and authentication features:

- Auth users
- L2TP users (user authentication; L2TP user receives default L2TP settings from the security device)
- XAuth users (user authentication; no support for remote setting assignments)
- Admin users (user authentication; administrator user receives default privilege assignment of read-only)

A SecurID ACE server can store L2TP, XAuth, and device administrator user accounts for authentication purposes; but it cannot assign L2TP, XAuth remote settings, or device administrator privileges.

Configuring an LDAP Authentication Server

Lightweight Directory Access Protocol (LDAP) a protocol for organizing and accessing information in a hierarchical structure resembling a branching tree. LDAP is used to locate resources, such as organizations, individuals, and files on a network, and helps authenticate users attempting to connect to networks controlled by directory servers.

To create an LDAP authentication server object, configure the following:

- LDAP Server Port: The port number on the LDAP server to which the security device sends authentication requests. The default port number is 389.
- Common Name Identifier: The identifier used by the LDAP server to identify the individual entered in a LDAP server. For example, an entry of "uid" means "user ID" and "cn" for "common name."
- Distinguished Name (dn): The path used by the LDAP server before using the common name identifier to search for a specific entry. (For example, c=us;o=juniper, where "c" stands for "country," and "o" for "organization.")

Supported Users

An LDAP server supports the following types of users and authentication features:

- Auth users
- L2TP users (user authentication; L2TP user receives default L2TP settings from the security device)
- XAuth users (user authentication; no support for remote setting assignments)
- Admin users (user authentication; administrator user receives default privilege assignment of read-only)

LDAP servers cannot assign L2TP or XAuth remote settings.

Configuring a TACACS Authentication Server

Terminal Access Controller Access Control System (TACACS) is a security application. As of Release 2007.3, you can configure TACACS to authenticate administrator users.

To configure the TACACS server:

1. In the NSM main navigation tree, click **Object Manager > Authentication Servers**.
2. Select the TACACS server type from the Authentication Server dialog box.
3. Configure the following parameters and click **OK**.
 - Secret - The secret (password) shared between the security device and the TACACS server. The device uses this secret to encrypt the user's password that it sends to the TACACS server.
 - Port - The port number on the TACACS server to which the security device sends an authentication request. The default port number is 49.

Configuring User Objects

User objects represent the users of your managed devices. You can include user objects or groups in security policies or VPNs to permit or deny access to individuals or groups. NSM supports two types of user objects:

- Local Users—Users with accounts that are managed by your security devices. You can create local user groups that include multiple users simplify user administration and make policies and VPNs easier to create.
- External Users and External User Groups—Users with accounts that are managed by external devices, such as RADIUS servers. You can use external users and groups to create group expressions (for details, see [“Configuring Group Expressions” on page 447](#)).

Configuring Local Users

Local user objects represent the user account on your security devices. To add a local user object:

1. In the navigation tree, double-click the **Object Manager**, select **User Objects**, then select **Local Users**. In the main display area, click the Add icon and select **New > User** to display the New Local User dialog box.
2. Enter a name, color, and comment for the local group.
3. Select **Enable** to enable authentication for this user, then configure the authentication methods for the user:

- XAuth. Enables XAuth authentication for this user. If you select this option, you must also enter an XAuth password for the user.



NOTE: All passwords handled by NSM are case-sensitive.

- IKE. Enables IKE authentication using one of the IKE proposals defined in the IKE proposal objects. If you select this option, you must also configure the IKE Share limit and authentication token.
- Auth. Enables local authentication against a username and password stored in a security device's local database. If you select this option, you must also enter an Auth password for the user.
- L2TP. Enables authentication in the L2TP tunnel that the user uses to connect to the device. If you select this option, you must also enter an L2TP password for the user.
- Click **OK** to save the user object.

Configuring Local User Groups

Organize local users in groups to add multiple users at one time to a security policy, and to manage the members without changing the policy. To add a local user group object:

1. In the navigation tree, double-click the **Object Manager**, select **User Objects**, then select **Local Users**. In the main display area, click the Add icon and select **New > Group** to display the New Local User Group dialog box.
2. Enter a name, color, and comment for the local user group.
3. Configure the members of the group:
 - To add members, select users from the Non-members list and click **Add**. Use Ctrl-click to select multiple users, or click **Add All** to add all users in Non-members list to the group.
 - To remove members, select users in the Members list and click **Remove**. Use Ctrl-click to select multiple users, or click **Remove All** to remove all users in Members list from the group.
4. Click **OK** to save the local user group.

Configuring External Users

External user objects represent users whose accounts are maintained and authenticated on devices that are not managed by NSM, such as an external RADIUS or SecureID server. When an external user is included in a security policy (under Authentication rule options), the security device uses the external server to authenticate that user.

To configure an external user:

1. In the navigation tree, double-click the **Object Manager**, select **User Objects**, then select **External Users**. In the main display area, click the Add icon and select **New** to display the New External User dialog box.
2. Enter a name, color, and comment for the external user.
3. Click **OK** to save the external user object.

Configuring External User Groups

External User Group objects represent user groups that are managed on non-security devices, such as an external RADIUS or SecureID server. When an external user group is included in a security policy (under Authentication rule options), the security device uses the external server to authenticate those users.

To use an external user group in a VPN, however, you must also create local user objects with IKE authentication for each external user. In phase 1 of IKE negotiations, the security device authenticates the external user group using the RADIUS server. In phase 2 of IKE negotiations, the device uses the local user object or local user group for authentication. Typically, you configure the local user object with IKE authentication and a U-FQDN (e-mail address); during phase 2, the device prompts the user for their U-FQDN for authentication.

To add an external user group object:

1. In the navigation tree, select **Object Manager > User Objects > External User Groups**. In the main display area, click the Add icon and select **New** to display the New External Group dialog box.
2. Enter a name for the external user group. The name must match the name of the user group as configured on the external server.
3. Enter a color and comment for the external user group.
4. Configure the authentication methods for the user group:
 - XAuth. Enables XAuth authentication for the user group.
 - Auth. Enables local authentication against a username and password stored in a security device's local database.



NOTE: All passwords handled by NSM are case-sensitive.

- L2TP. Enables authentication in the L2TP tunnel that users in the group use to connect to the device.
5. Click **OK** to save the new group.

Using Radius with User Groups

In this example, you configure an external RADIUS auth server named radius1 and define an external auth user group named auth_grp2. You define the external auth user group auth_grp2 in two places: External RADIUS auth server “radius1,” and in NSM. For the

RADIUS server, you enter the IP address 10.20.1.100 and change its port number from the default port number (1645) to 4500.

Next, you populate the auth user group “auth_grp2” with auth users on the RADIUS server only, leaving the group unpopulated in NSM. The members in this group are accountants who require exclusive access to a server at IP address 10.1.1.80. You create an address book entry for the server and name the address “midas.” Finally, you configure a security policy that permits only authenticated traffic from auth_grp2 to midas, both of which are in the Trust zone.

1. On the RADIUS server, load the Juniper Networks dictionary file and define auth user accounts. Use the Juniper Networks user group VSA to create the user group auth_grp2 and apply it to the auth user accounts that you want to add to that group.



NOTE: For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation. If you are using a Microsoft IAS RADIUS server, there is no dictionary file to load; you must manually define the correct vendor-specific attributes (VSAs) on the server.

2. In NSM, in the main navigation tree, select **Object Manager > Authentication Servers** and click the Add icon. Configure the server:
 - a. For name, enter **radius1**. Select a color and add a comment, if desired.
 - b. For Main Server, enter the IP **10.20.1.100**; for Primary Backup Server, enter IP **10.20.1.110**; for Secondary Backup Server, enter IP **10.20.1.120**.
 - c. For timeout, enter **30**.
 - d. Select **For Firewall Auth Users**.
 - e. For Server Type, select **RADIUS**, then configure the RADIUS server:
 - For server port, enter **4500** (default is 1645)
 - For secret, enter **A56hYY97kl**
 - For retry timeout, select **4**.
 - f. Click **OK** to save the RADIUS authentication server object.
3. Configure the External User Group in NSM:
 - a. In the Object Manager, select **User Objects > External User Groups**.
 - b. Click the Add icon to display the New External User Group dialog box. Configure the following, then click **OK**:

- For Name, enter **auth_grp2**.
 - For Color, select **red**.
 - For Comment, enter **Accountant Access**.
 - Enable **Auth**.
4. Add the address object that represents the Accounting Server:
 - a. In the Object Manager, select **Address Objects**. Click the Add icon and select **Host**. The New Host dialog box appears.
 - b. Configure the following, then click **OK**:
 - For Name, enter **Midas**.
 - For Color, select **orange**.
 - For Comment, enter **Accounting Server**.
 - Select **IP**, then enter the IP Address **10.1.1.80**.
 5. Configure a firewall rule to use the RADIUS authentication server object to authenticate traffic between the external user group and the Midas server.

Configuring VLAN Objects

Use VLAN objects to limit rule matching to packets within a particular VLAN.

VLAN objects can either target a specific VLAN tag, or a range of VLAN tags. You can use more than one VLAN object in a rule.

VLAN objects have the following components:

- Name: What the object is called in the NSM UI.
- Comment and Color: Useful for organizing and explaining the object to other users. Have no effect on the object in the system.
- ID Type: Specify whether the object will cover a single VLAN tag or a range of VLAN tags.
- Specific and Low/High: For a single VLAN tag, specify the tag. For a range of VLAN tags, specify the lowest and highest values in the range.

Configuring IP Pools

An IP pool object contains IP ranges (a range of IP addresses within the same subnet). You use IP Pool objects to assign IP addresses to L2TP users in L2TP VPNs or local users on a specific device. The IP pool you select for the VPN or the local user determines the range of IP addresses the device can assign to the L2TP RAS user when the user connects to the L2TP VPN.



NOTE: For more information about configuring XAuth and L2TP local users on a device, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

An IP range includes the following:

- Start IP—The beginning of the range of IP addresses included in the pool, inclusive. The Start IP must always be lower than the End IP.
- End IP—The end of the range of IP addresses included in the pool, inclusive. The End IP must always be higher than the Start IP.

Using Multiple IP Ranges

An IP Pool object can contain multiple, non-sequential IP ranges. You might need to use multiple ranges to accommodate large numbers of RAS users in a VPN.

You can configure up to 256 IP ranges within a single IP Pool object. You can add any number of IP Pool objects.



NOTE: Devices running ScreenOS 5.1 or earlier versions do not support multiple IP pool ranges. When you include a multi-range IP pool object in a device configuration or VPN for a device running ScreenOS 5.1 or earlier, the device automatically uses the first IP range defined in the IP Pool object.

To modify or delete an IP range from an IP Pool object, you must first ensure that no IP within the range is currently in use by any managed device. If you change or delete an IP range that contains a used IP address, the device using the IP generates an error during device update (error message appears with the Job Manager dialog box for the update).

In this example, you configure an IP pool with the ranges 1.1.1.1-1.1.1.10 and 2.2.2.2-2.2.2.20.

1. In the navigation tree, select **Object Manager > IP Pools**.
2. In the main display area, click the Add icon. The New IP Pool dialog box appears. Configure as follows:
 - For Name, enter **L2TP User Group 1**.
 - For Color, select **orange**.
 - For Comment, enter IPs for usergrp 1.
3. In the IP Pool dialog box, click the Add icon to configure the first IP pool range. The New IP Pool Name dialog box appears. Configure the Start IP and End IP, then click **OK**:

- For Start IP, enter **1.1.1.1**.
 - For End IP, enter **1.1.1.10**.
4. In the IP Pool dialog box, click the Add icon to configure the second IP pool range. The New IP Pool Name dialog box appears. Configure the Start IP and End IP, then click **OK**:
- For Start IP, enter **2.2.2.2**.
 - For End IP, enter **2.2.2.20**.
5. Click **OK** again to save the IP Pool object and return to Object Manager.

Configuring Group Expressions

Group expressions are statements that set conditions for authentication requirements, enabling you to combine multiple external user objects. You can create group expressions using the operator OR, AND, or NOT to combine user objects, user group objects, or other group expressions to define:

- Alternatives for authentication (“a” OR “b”)
- Requirements for authentication “a” AND “b”)
- Exclusions of a user group, or another group expression (NOT “c”).



NOTE: The user and user groups you reference in the group expressions must be external users that are stored on an external RADIUS server. (A RADIUS server enables a user to belong to more than one user group).

The operators have different meanings depending on the type of user object you are using in the security policy, as listed in [Table 42 on page 447](#).

Table 42: Group Expression Operators

User Objects	
OR	If the security policy defines authentication for “a” or “b” user objects, the security device authenticates the user if it is either “a” or “b” .
AND	Requires one of the two objects in the expression to be either a user group or a group expression (a single user cannot be both user “a” and user “b”). If the security policy defines authentication for “a” AND a member of group “b” , the security device authenticates the user only if those two conditions are met.
NOT	If the security policy defines authentication for any user object that is not the “c” user (NOT “c”), the security device authenticates all users except the “c” user.

Table 42: Group Expression Operators (continued)

User Objects	
User Groups	
OR	If the security policy defines authentication for user group "a" or user group "b", the security device authenticates the user if it belongs to either "a" or "b" user group.
AND	If the security policy defines authentication for user group "a" AND user group "b", the security device authenticates the user only if it belongs to both user groups.
NOT	If the security policy defines authentication for any user group that is not group "c" (NOT "c"), the security device authenticates all users except those that belong to the "c" user group.
Group Expressions	
OR	If the security policy defines authentication for user objects that match the description of group expression "a" OR group expression "b", the security device authenticates the user if either group expression references that user.
AND	If the security policy defines authentication for user objects that match the description of group expression "a" AND group expression "b", the security device authenticates the user only if both group expressions reference that user.
NOT	If the security policy defines authentication for user objects that do not match the description of group expression "c" (NOT "c"), the security device authenticates all users except those that match the group expression.

Because a group expression references external user objects and external user groups, you must first create those user object and groups before you can use them in a group expression. You cannot reference local user object or local user object groups in a group expression.

To add a group expression:

1. In the navigation tree, double-click **Object Manager** and select **Group Expressions**.
2. In the main display area, click the Add icon and select **New**. The New Group Expression dialog box appears.
3. Enter a name, color, and comment for the group expression.
4. Select the operator you want to use in the expression (OR, AND, NOT) and then configure the operands:

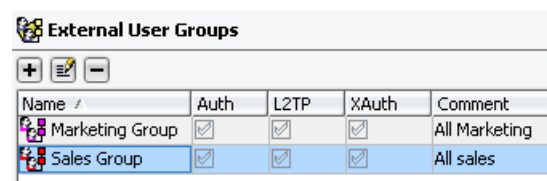
- For NOT expressions, use Operand 1 to select the user object, group, or expression that cannot be present for a successful match. Because the operation is exclusion, you do not need to configure Operand 2.
 - For AND expressions, use Operand 1 and Operand 2 to select the two user objects, groups, or expressions that must be present for a successful match.
 - For OR expressions, use Operand 1 and Operand 2 to select the two user objects, groups, or expressions, one of which must be present for a successful match.
5. Click **OK**. The group expression object appears in the Object Manager.

After you have created a group expression object, you can use that object in the Authentication rule options.

In this example, you configure a group expression to authenticate all users that belong to your Sales group and your Marketing group, then add the expression to a security policy that provides access to your protected networks.

1. First, create two external user group objects: one to represent the Sales users and the other to represent the Marketing users.

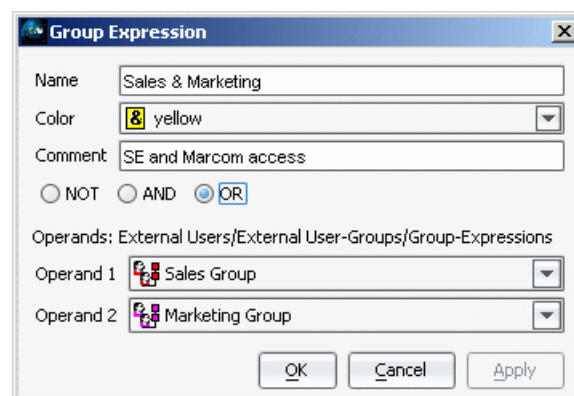
Figure 87: Configure External User Groups for Sales and Marketing



Name	Auth	L2TP	XAuth	Comment
Marketing Group	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	All Marketing
Sales Group	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	All sales

2. Next, create a group expression object that references both the Sales and Marketing groups.

Figure 88: Configure Group Expression for Sales and Marketing



Group Expression

Name: Sales & Marketing

Color: yellow

Comment: SE and Marcom access

☐ NOT ☐ AND ☒ OR

Operands: External Users/External User-Groups/Group-Expressions

Operand 1: Sales Group

Operand 2: Marketing Group

OK Cancel Apply

3. Finally, add the group expression object to your firewall rule in the Authentication rule option.

Configuring Remote Settings

A remote settings object defines the DNS and WINS servers that are assigned to L2TP RAS users after they have connected to the L2TP tunnel. You can use remote settings objects in an L2TP VPN, and when configuring a local user on a specific device.



NOTE: For information about configuring XAuth and L2TP local users on a device, see *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Security devices incorporate DNS (domain name server) and WINS support to permit the use of domain names as well as IP addresses for identifying locations. A DNS or WINS server keeps a table of the IP addresses associated with domain names. Using DNS or WINS, you can reference locations by their domain name (www.juniper.net) in addition to using a routeable IP address (such as 209.125.148.136).

Before you can use DNS or WINS for domain name/address resolution in a VPN, you must create remote settings for the DNS or WINS servers (primary and secondary).

To configure a remote setting, select Remote Settings and click the Add icon. Enter a name, color, and comment for the object, then configure the following parameters:

- DNS1—Enter the IP address of the primary DNS server.
- DNS2—Enter the IP address of the secondary DNS server.
- WINS1—Enter the IP address of the primary WINS server.
- WINS2—Enter the IP address of the secondary WINS server.

Configuring Routing Instance Objects

A routing instance is a collection of routing tables, interfaces contained in these routing tables, and routing option configurations. You can use Object Manager to configure a routing instance object. The routing instance objects configured in Object Manager can be included in the RADIUS server and LDAP server configurations within the access profile object. A routing instance object is a polymorphic object (similar to zone objects) that maintains the mapping between the actual routing instance and the device in which it is created. For details about polymorphic objects, see [“Polymorphic Objects” on page 576](#).

To view all routing instance objects, select **Routing Instance Objects** in the navigation tree. The Object Manager displays all the routing instance objects in the routing instance table. You can create, view, edit, or delete a routing instance object in the Object Manager. You can also perform a Find Usages operation, and view the version history of a routing instance object. For more information on configuring routing instances, see the *Junos Routing Protocols Configuration Guide*.



NOTE: The routing instance shared object is not supported on ScreenOS devices.

Viewing Routing Instance Objects

To view all routing instance objects, select **Routing Instance Objects** in the navigation tree. The Object Manager displays all the routing instance objects configured in NSM. Routing instance objects are listed in a table consisting of the following columns:

- Name—Name of the routing instance object.
- Domain—Domain where the routing instance object has been created.
- Service—The actual routing instance in the device.
- Comment—Description of the routing instance object.

Creating Routing Instance Objects

To create a routing instance object:

1. In the navigation tree, select **Object Manager > Routing Instance Objects**. The Object Manager displays all the routing instance objects configured in NSM.
2. In the main display area, click the Add icon. The New Routing Instance dialog box appears.
3. Enter a unique name for the routing instance object.
4. Select a color to represent the routing instance object.
5. Enter a comment or description about the routing instance object.
6. In the New Routing Instance dialog box, click the Add icon. The New Routing Instance Entry dialog box appears.
7. Enter the name of the domain where you want to create the routing instance object.
8. Enter the name of the device in which you want to create the routing instance.
9. Select a routing instance from the routing instance drop-down list box and click OK. If no routing instance is available, you need to create a routing instance using the Add icon in the New Routing Instance Entry dialog box. For details on adding routing instances, see the *Junos Routing Protocols Configuration Guide*.

Configuring Zone Group Objects

A number of zones grouped together forms a zone group object.

A use case scenario where this object proves useful is when you want to enforce the same firewall policies for the traffic among multiple zones. One zone group object can represent all zones from where the traffic originates and another zone group object can represent all zones to which traffic flows. When you add these zone group objects in the **From Zone** and **To Zone** columns of a firewall policy and run a summarize delta config,

NSM expands these objects and generates policies between each zone from the From Zone group to each zone in the To Zone group.

For example, consider you want to enforce the same firewall policies for the traffic flowing from $n1$ zones to $n2$ zones:

1. From **Object Manager**:

- a. Create a zone group object, **Zone_Group_Object_A**, which consists of a group of $n1$ zones.
- b. Create a zone group object, **Zone_Group_Object_B**, which consists of a group of $n2$ zones.

2. From **Policy Manager**:

- a. Select a policy.
- b. Under the **Zone based Firewall** tab, select the newly added zone group objects in the **From Zone** and **To Zone** columns.

In this example, select **Zone Group Object A** in the **From Zone** column and **Zone Group Object B** in the **To Zone** column.

3. From **Device Manager**:

- a. Select the device to which you want to apply this configuration.
- b. **Run Summarize Delta Config.**

NSM automatically generates $n1*n2$ firewall policies. The newly created policies are displayed in the **Job Information** dialog box.

The IDs of these policies are greater than 750000 and less than 998499.

Update the device to push these policies to the device.

Viewing Zone Group Objects

To view all zone group objects, select **Zone Group Objects** in the navigation tree. The **Object Manager** displays all the zone group objects configured in NSM.

Zone group objects are listed in a table consisting of the following columns:

- **Name**—Name of the zone group object.
- **Domain**—Domain where the zone group object has been created.
- **Device**—Device where the zone group object has been created.
- **Service**—The actual zone group in the device.

An alternate method is to view the zone group objects in the drop-down list in the **Shared Objects for Policy** section of the firewall policy window.

Adding a Zone Group Object

To create a zone group object:

1. In the navigation tree, select **Object Manager > Zone Group Objects**. The **Object Manager** displays all the zone group objects configured in NSM.
2. In the main display area, click the Add icon. The **New Zone** dialog box appears.
3. Enter a unique name for the zone group object.
4. Select a color to represent the zone group object.
5. Click the Add icon. The **New Zone Map Entry** dialog box appears.
6. Enter the name of the domain where you want to create the zone group object.
7. Enter the name of the device in which you want to create the zone group.
8. Select a zone from the **Zone** drop-down list box and click **OK**. If no zones are available, you need to create a zone using the Add icon in the **New Zone Map Entry** dialog box.

Configuring NAT Objects

The types of Network Address Translation (NAT) objects that are supported in NSM are legacy NAT objects for Screen OS devices and Junos NAT objects for Junos OS devices. For more information, see the following sections:

- [Configuring Legacy NAT Objects on page 453](#)
- [Configuring Junos OS NAT Objects on page 455](#)

Configuring Legacy NAT Objects

A global NAT object contains references to device-specific NAT configurations, enabling multiple devices to share a single object. Use the Device Manager to configure NAT for each device, then create a global NAT object that includes the device-specific NAT configuration. The single global NAT object represents multiple device-specific NAT objects; for example, a global dynamic IP (DIP) represents multiple device-specific DIPs. However, a global NAT object can contain only one device-specific NAT object from the same device.

Use global NAT objects in VPNs; when you install the VPN on a device, that device automatically replaces the global NAT object with its device-specific NAT configuration. Before you configure a shared NAT object, ensure that you have configured the mapped IP (MIP), virtual IP (VIP), or DIP on the device itself.

You cannot configure NAT objects for SRX Series Services Gateways and use them in security policies. For SRX Series gateways, NAT settings must be configured in the device.

For more information on DIP, MIP, and VIP objects, see the following sections:

- [Configuring DIP Objects on page 454](#)
- [Configuring MIP Objects on page 454](#)
- [Configuring VIP Objects on page 454](#)
- [Configuring Destination NAT Objects on page 455](#)

Configuring DIP Objects

To configure a DIP object:

1. In Object Manager, select **NAT Objects > DIP** and click the Add icon.
2. Enter a name, color, IP version (IPv4 or IPv6), and comment for the object, then click the Add icon to specify the device-specific DIP:
 - Device—Select the security device that includes the DIP.
 - Interface or DIP Group—Select the interface or DIP group for the device.
 - For interface, select the interface on the device and the dynamic IP address configuration for that interface.
 - For DIP group, select the dynamic IP group configuration for that device.

If no values appear in the pull-down menu for interface, DIP, or DIP group, make sure that you have configured DIP correctly in the Device Manager.

You can add multiple device DIPs to a single global DIP object (one DIP per device).

Configuring MIP Objects

To configure a MIP object:

1. In Object Manager, select **NAT Objects > MIP** and click the Add icon.
2. Enter a name, color, IP version (IPv4 or IPv6), and comment for the object, then click the Add icon to specify the device-specific MIP:
 - Device—Select the security device that includes the MIP.
 - Interface—Select the interface on the device that uses the mapped IP address.
 - MIP—Select the mapped IP address configuration for that interface.

If no values appear in the pull-down menu for interface or MIP, make sure that you have configured MIP correctly in the Device Manager. You can add multiple device MIPs to a single global MIP object.

For information about configuring a MIP object and an example, see the NSM Online Help description “Configuring Firewall/VPN Devices.”

Configuring VIP Objects

To configure a VIP object:

1. In Object Manager, select **NAT Objects > VIP** and click the Add icon.
2. Enter a name, color, and comment for the object, then click the Add icon to specify the device-specific VIP configuration:

- Device—Select the security device that includes the VIP.
- Interface—Select the interface on the device that uses the virtual IP address.
- VIP—Select the virtual IP address configuration for that interface.

If no values appear in the pull-down menu for interface or VIP, ensure that you have configured VIP correctly in the Device Manager. You can add multiple device VIPs to a single global VIP object.

For information about configuring a VIP object and an example, see the NSM Online Help description “Configuring Firewall/VPN Devices.”

Configuring Destination NAT Objects

To configure a destination NAT object:

1. In Object Manager, select **NAT Objects > Destination NAT** and click the Add icon.
2. Enter a name, color, and comment for the object, then click the Add icon to specify the device-specific destination NAT configuration:
 - Device—Select the security device.
 - Destination-nat—Select a value from the pull-down menu. If no values appear on the pull-down menu, click the Add icon to create a new value.

Configuring Junos OS NAT Objects

You can configure NAT objects (source NAT and destination NAT objects) for devices running on the Junos OS. These objects represent user-defined address pools and during network address translation, the original source or destination IP address within a packet is translated to an IP address within this pool. For more information on configuring these objects, see the following sections:

- [Configuring Source NAT Objects on page 455](#)
- [Configuring Destination NAT Objects on page 459](#)

Configuring Source NAT Objects

A source NAT object consists of a user-defined address pool and is used during source address translation. You can use this object while configuring a rule so that when the rule is matched, the source IP address of the packet is translated to an IP address from this pool.

You can create, edit, delete, and search for a source NAT object from **Object Manager**. For more information, see the following sections:

- [Adding a Source NAT Object on page 456](#)
- [Editing a Source NAT Object on page 457](#)
- [Deleting a Source NAT Object on page 458](#)

Adding a Source NAT Object

To add a source NAT object:

1. Select **Object Manager > Junos NAT Objects > Source NAT**.

The Source NAT page appears on the right pane. You can add, edit, delete, or search for a source NAT object using the icons at the top of this dialog box.

2. Select **(+)** to add a new source NAT object. The **New JunosSource NAT** dialog box appears.

- Enter a name, color, and comment.
- Select **(+)** to configure the parameters for the new source NAT object. A **New – Junos Source NAT** dialog box appears. Here, you must select the device that performs the translation and specify the address pool.

3. Select a device from the **Device** drop-down list.

- The devices are listed only if you have previously added these devices to NSM. To add a device, use the **Device Manager**.
- If a device is not selected, you cannot configure the **Proxy-ARP** and **Junossource-nat** fields.

4. If the proxy ARP functionality is required, select the interface which accepts the ARP requests, from the **Proxy ARP** drop-down list. If no values are listed, select **(+)** to configure a new value. The **New Interface** dialog box appears. Specify the logical interface on which to configure proxy ARP and its details. For more information, see [“Configuring Proxy ARP” on page 457](#).

5. Select **(+)** next to **Junossource-nat** to configure the address pool. The **New Pool** dialog box appears.

You will see a list of values to select from the drop-down list if you have previously configured address pools.

[Table 43 on page 456](#) lists the fields that are available and the action you need to perform on each fields.

Table 43: Source NAT Configuration Options

Tab	Field	Function	Action
General	Name	Descriptive name for the pool.	Type a name for the pool.
Routing Instance	Ri Name	Specify the routing instance to which the pool is bound.	Select the routing instance name. The values are listed only if you have added them previously. To add a new routing instance to a device, select Object Manager > Routing Instance Objects .
Address	IP Address/Ipaddr	Specify address prefixes, addresses, or a range of addresses or address prefixes.	Click the Add or Delete icons to configure or remove address entries.

Table 43: Source NAT Configuration Options (continued)

Tab	Field	Function	Action
Port Translation	Low/High	Specify whether port translation must be performed or not.	By default, port translation is enabled. Enter a port range. Select No Translation to disable port translation.
Host Address Base	Ipaddr	Specify the base address of the original source IP address range. This is used for IP shifting.	Enter the IP address.
Overflow Pool		Specify a source pool to use when the current address pool is exhausted. The pool can be a user-defined pool or the IP address of an interface.	<ul style="list-style-type: none"> • None—No overflow pool. • Pool-name—Select a user-defined pool. • Interface—Enter the IP address of the interface.

Configuring Proxy ARP

To configure proxy ARP:

1. In the **New Interface** dialog box, enter the name of the interface. To navigate to this dialog box, see steps 1 to 4 of [“Adding a Source NAT Object” on page 456](#).
2. Specify the hosts (range of IP addresses) whose ARP requests this device must accept, as follows:
 - Click **Address** and select (+) to configure the start of the address range in the **New** dialog box.
 - Click **To** and configure the end of the address range.
 - Click **OK**.
3. Click **OK**. Proxy ARP is now configured.

Editing a Source NAT Object

To edit a source NAT object:

1. Select the SRX Series device from the Device Tree.
2. Right-click and then select **Edit**.
3. Select **Security > NAT > Source**.
The Source NAT options appear.
4. Click **Pool** and select the existing source NAT pool.

5. Right-click and then select **Edit** to edit the source NAT pool.
6. Click **OK**.



NOTE: To edit or delete a source NAT object, you must select the object from the device through which the source NAT object was configured. The source NAT object that is available in **Object Manager**, allows only a new configuration.

Deleting a Source NAT Object

A source NAT object should be deleted from the device. To ensure the complete removal of the source NAT object, also delete the source NAT object from the corresponding NAT object under **Object Manager**.

To delete a source NAT object from the device:

1. Select the SRX Series device from the Device Tree.
2. Right-click and then select **Edit**.
3. Select **Security > NAT > Source**.
The Source NAT options appear.
4. Click **Pool** and then select the existing source NAT pool.
5. Right-click and then click **Delete**.
The source NAT object is deleted.

To delete a source NAT object from Object Manager:

1. Select **Object Manager > Junos NAT Objects > Source NAT**.
The Source NAT page appears on the right pane.
2. Right-click on the source NAT pool and click **Edit**.
3. Right-click on the device in which the NAT was removed, and click **Delete**.



NOTE: Alternatively, you can select (-) at the top of the page or dialog box to delete the entries.

Configuring Destination NAT Objects

A destination NAT object consists of a user-defined address pool and is used during destination address translation. You can use this object while configuring a rule so that when the rule is matched, the destination IP address of the packet is translated to an IP address from this pool.

You can add, edit, delete, and search for a destination NAT object from **Object Manager**. For more information, see the following sections:

- [Adding a Destination NAT Object on page 459](#)
- [Editing a Destination NAT Object on page 460](#)
- [Deleting a Destination NAT Object on page 461](#)

Adding a Destination NAT Object

To add a destination NAT object:

1. Select **Object Manager > Junos NAT Objects > Destination NAT**.

The Destination NAT page appears on the right pane. You can add, edit, delete, or search for a destination NAT object using the icons at the top of this dialog box.

2. Select (+) to create a new destination NAT object. The **New JunosDestination NAT** dialog box appears:
 - Enter a name, color, and comment for the new destination NAT object.
 - Select (+) to configure the parameters for the new destination NAT object. The **New – Junos Destination NAT** dialog box appears. Here, you must select the device that performs the translation and define the address pool.
3. Select a device from the **Device** drop-down list.
 - The devices are listed only if you have previously added these devices to NSM. To add a device, use the **Device Manager**.
 - If a device is not selected, you cannot configure the **Proxy-ARP** and **Junosdestination-nat** fields.
4. If the proxy ARP functionality is required, select the interface which accepts the ARP requests, from the **Proxy ARP** drop-down list. If there are no values listed, select (+) to configure a new value. The **New Interface** dialog box appears. Specify the logical interface on which to configure proxy ARP and its details. For more information, see [“Configuring Proxy ARP” on page 460](#).
5. Select (+) next to **Junosdestination-nat** to configure the address pool. The **New Pool** dialog box appears.

You will see a list of values to select from the drop-down list if you have previously configured address pools.

Table 44 on page 460 lists the fields that are available and the action you need to perform on each fields.

Table 44: Destination NAT Configuration Options

Tab	Field	Function	Action
General	Name	Descriptive name for the pool.	Type a name for the pool.
Routing Instance	Ri Name	Specify the routing instance to which the pool is bound.	Select the routing instance name. The values are listed only if you have added them previously. To add a new routing instance to a device, select Object Manager > Routing Instance Objects .
Address	IP Address tab: IP Address	Specify a single IP address.	Enter an IP address.
	To Range/Port tab:	<ul style="list-style-type: none"> To Address—Specify an address range. Port—Specify the optional port number. 	<ul style="list-style-type: none"> Enter an IP address in the To Address field. Enter a port number in the Port field.

Configuring Proxy ARP

To configure proxy ARP:

1. In the **New Interface** dialog box, enter the name of the interface. To navigate to this dialog box, see steps 1 to 4 of [“Adding a Destination NAT Object” on page 459](#).
2. Specify the hosts (range of IP addresses) whose ARP requests this device must accept, as follows:
 - Click **Address** and select (+) to configure the start of the address range in the **New** dialog box.
 - Click **To** and configure the end of the address range.
 - Click **OK**.
3. Click **OK**. Proxy ARP is now configured.

Editing a Destination NAT Object

To edit a destination NAT object:

1. Select the SRX Series device from the Device Tree.
2. Right-click and then select **Edit**.
3. Select **Security > NAT > Destination**.

The Destination NAT options appear.

4. Click **Pool** and then select the existing destination NAT pool.
5. Right-click and then select **Edit** to edit the destination NAT pool.
6. Click **OK**.

Deleting a Destination NAT Object

To delete a destination NAT object:

1. Select the SRX Series device from the Device Tree.
2. Right-click and then select **Edit**.
3. Select **Security > NAT > Destination**.

The Destination NAT options appear.

4. Click **Pool** and then select the existing destination NAT pool.
5. Right-click and then click **Delete**.

The destination NAT object is deleted.



NOTE: To edit or delete a destination NAT object, you must select the object from the device through which the destination NAT object was configured. The destination NAT object that is available in Object Manager, allows only a new configuration.

Configuring Certificate Authorities

A digital certificate is an electronic means for verifying your identity through the word of a trusted third party, known as a Certificate Authority (CA). NSM simplifies creating and managing certificates:

- Use the same CA server for multiple devices. Create a single CA object for each CA server you use, then use that object for those devices.
- Generate a local and CA certificate in one click using SCEP.
- Use OCSP to automatically check for revoked certificates (ScreenOS 5.0 or later devices only)
- Use a certificate chain that includes a root CA and subordinate CA (CA group)

A CA object represents the CA server you want to use to authenticate the identity of your VPN member. You can use an independent or internal CA server:

- Independent CA server—Owned and operated by an independent CA. The independent CA provides the IP addresses of their CA and CRL servers. You submit a local certificate request to the independent CA and provide your local certificate information.
- Internal CA server—Owned and operated by your company. You provide the IP addresses of the CA and CRL servers and local certificate information.

You can obtain a CA certificate file (.cer) from the CA that issued the local certification, then use this file to create a Certificate Authority object. Then, install this CA certificate on the managed device using NSM. Because the CA certificate is an object, however, you can use the same CA for multiple devices, as long as those devices use local certificates that were issued by that CA.

Alternatively, you can use SCEP to configure the device to automatically obtain a CA certificate at the same time it receives the local certificate. For details, see the NSM Online Help description of “Configuring Firewall/VPN Devices.”

Using Certificate Authorities

You must use obtain and install a CA certificate on each VPN member to authenticate the local device certificates on your managed devices.

Configuring Certificate Authorities

After you have obtained a CA Certificate file (.cer) from your CA, use this file to create a Certificate Authority object. In Object Manager, select Certificate Authorities, then click the Add icon to display the New CA Certificate dialog box. Enter a name for the CA Certificate, then click Load CA certificate and load the appropriate .cer file. NSM uses the information in the .cer file to automatically complete the Subject Name, Issued By, and Expired On fields.

Complete the remaining settings:

- X.509 Certificate Path Validation Level—X509 contains a specification for a certificate which binds an entity's distinguished name to its public key through the use of a digital signature.
 - Full. Use full validation to validate the certificate path back to the root.
 - Partial. Use partial validation to validate the certificate path only part of the way to the root.
- Revocation Check
 - Check for revocation. Select this option to enable revocation checking.
 - Do not check for revocation. Select this option to disable revocation checking.
- Revocation Checking Method—If you enabled revocation checking, you can select the checking method to use. If you did not enable revocation checking, these fields are unavailable.
 - CRL. Use a Certificate Revocation List when you want to keep a local copy of the revoked certificates on the managed device. This method enables the device to

check for revoked certificates quickly; to accept the certificate if no revocation information is found, also enable Best Effort.

- OCSP. Use the Online Certificate Status Protocol when you want the managed device to access a remote OCSP server to check for revoked certificates. Because the OCSP server dynamically updates its list of revoked certificates, this method provides the most up-to-date information; to accept the certificate if no revocation information is found, also enable Best Effort.
- Best Effort. Enable this option to check for revocation accept the certificate if no revocation information is found.
- CRL Settings—Configure the default setting for the Certificate Revocation List.
 - *Refresh Frequency*. Select the frequency that the device contacts the CA to obtain a new CRL list: Daily, Weekly, or Monthly.
 - *LDAP server*. Provide the IP address of the external LDAP server that manages the CRL.
 - *URL address*. Provide the URL address of your internal LDAP server that provides the CRL.
- OCSP—Configure the Online Certificate Status Protocol to dynamically check for revoked certificates.
 - Certificate Verification.
 - No revoke status check for CA delegated signing cert.
 - URL of OCSP Responder. Provide the URL address of the OCSP server.
- SCEP—Configure Simple Certificate Enrollment Protocol to get a local certificate automatically.
 - CA CGI. Enter the URL address of the Certificate Authority Certificate Generation Information.
 - RA CGI. Enter the URL address of the Registration Authority Certificate Generation Information that the security device contacts to request a CA certificate.
 - CA IDENT. Enter the name of the certificate authority to confirm certificate ownership.
 - Challenge. Enter the challenge words sent to you by the CA that confirm the security device identity to the CA.
 - CA Certificate Authentication. (Auto or Manual)
 - Polling Interval. (Poll or Do not poll).
 - Certificate Renewal. Define the number of times a certificate can be renewed.

Click **OK** to complete the CA object.

Configuring CRL Objects

A Certificate Revocation List (CRL) identifies invalid certificates. You can obtain a CRL file (.crl) from the CA that issued the local certification and CA certificate for the device, then use this file to create a Certificate Revocation List object.

You must install the CRL on the managed device using NSM. Because the CRL is an object, however, you can use the same CRL for multiple devices, as long as those devices use local and CA certificates that were issued by that CA.

Using CRLs

You can use a CRL object in a VPN to check for VPN members using revoked certificates.

Configuring CRLs

After you have obtained a CRL file (.crl) from your CA, use this file to create a Certificate Revocation object.

In Object Manager, select **CRLs**, then click the icon to display the New CRL dialog box. Enter a name for the CRL, then click **Load CRL** and load the appropriate .crl file. NSM uses the information in the .crl file to automatically complete the Issued By and Expire On fields. Click **OK** to complete the CRL object.

Configuring Extranet Policies

Extranet policies enable you to configure and manage extranet devices (that is, third-party router).

In this example, you want to update an existing policy on a third-party router to deny certain ftp traffic from a specific IP address. You can do this by creating a script that performs the required actions when you update the extranet device. You also need to create your rule in an Extranet Policy object.

To create an Extranet Policy object:

1. In the Object Manager, select Extranet Policies. The New ExtranetPolicyObject window appears.
2. Enter the name of the Extranet Policy, for example, Extranet Policy1. Add a comment in the Comments field.
3. Configure the Extranet Policy object:

- Click **New**. The New - Rule window appears.
- Use the up/down arrow to specify an ID for the rule.
- Add a comment for the rule.
- Click Deny in the Action field.
- Select a source address in the Source tab.
- Select a destination address in the Destination tab.
- Select **FTP** in the Service tab.
- Select the integer IDs that you created in the Custom Policy Field object in the Options tab.

4. Click **OK**.

When you create the extranet device in NSM, bind the policy to the appropriate interface and specify the script you want to perform the required update actions. When you update the device, NSM invokes the script. Any XML output appears in the Job Information window.

Configuring Binary Data Objects

Binary data objects provide an efficient way of handling large binary files of configuration data that are typical in Secure Access and Infranet Controller device configurations. These files are the only configuration objects that are not imported when you import a device configuration. Without this special handling, their size could overwhelm server resources.

These files must be copied separately onto the client UI device, and configured as shared objects. As a result, only the large binary files that you need to manage in NSM are imported, and those files can be shared across multiple devices.

Binary data files handled in this way include:

- Custom sign-in access pages
- Custom sign-in meeting pages
- Antivirus live update files
- Antivirus live update patch files
- Endpoint Security Assessment Plug-in (ESAP) packages
- Third-party host checker policies
- Secure virtual workspace wallpaper images
- Hosted Java applets
- Custom Citrix client CAB files

See ["Managing Large Binary Data Files \(Secure Access and Infranet Controller Devices Only\)" on page 289](#) for information about where to find sources for these files, how to

upload them to NSM, and how to link to them from to include them in a device object configuration.

Adding Binary Data Objects

Before creating the object in the Object Manager, you must copy the file from its source to the file system of your client UI device:

1. In the Configure panel of the NSM main navigation tree, select **Object Manager > Binary Data**.
2. Click the Add icon. In the New Binary Data dialog box, give the object a name, select a color for the icon, add an optional comment, and select the binary file that the object will reference by navigating to it in the client UI file system.
3. Click **OK** to add the object to the Binary Data list in the Object Manager.

Viewing, Editing, and Deleting Binary Data Objects

To view binary data objects, from the Configure panel of the main navigation tree, select **Object Manager > Binary Data**. The object manager lists each configured binary data object, and provides the following information about each object:

- The name given to the object for use in NSM.
- The pathname to the file on the client UI device.
- A comment provided by the administrator.

To edit a binary data object:

1. Double-click on the object in the Binary Data list. In the Binary Data dialog box, you can change the object name, the color of its icon, the comment, or the file on the client UI device that the object references.
2. Click **OK** when you are done.

To delete a binary data object, select the object in the Binary Data list and click the delete icon.

Configuring Protected Resources

A protected resource combines network components, network services, a traffic direction, and the security devices that protect those components and services. Protected resources are the source and destination addresses of a policy-based VPN.

Protected resources consist of the following elements:

- **IP Address**—The address represents the computer, network, or range of addresses to be considered part of this protected resource. The address can be an individual host, a network, or an address group.
- **Network Service**—Services are the protocols (HTTP, FTP) that communicate over a network. The service can be an individual service or a service group.
- **Traffic Direction**—Traffic direction is determined by the IP address that initiates the connection:
 - Client connections are outgoing (outbound) from the protected network.
 - Server connections are incoming (inbound) to the protected network.
 - To protect incoming and outgoing traffic, select **Both**.
- **Security Device**—The device that protects the network component and server. If the resource can be reached through more than one device, add multiple devices to the resource. When you add a protected resource to a VPN, the devices in the protected resource are included in the VPN.

Each protected resource represents an address or a range of addresses on your network. Each resource also can specify a service (such as FTP or NSF). Therefore, the protected resource is the destination for all traffic using the selected service to the selected address.

You can have more than one protected resource for a single address or range of addresses. That way you can individually manage different services traffic to the same destination separately.

Creating Protected Resources

To add a protected resource object:

1. In the navigation tree, select **VPN Manager > Protected Resources**. In the main display area, click the Add icon to display the Protected Resource dialog box.
2. Enter a name for the protected resource.
3. Select the services you want to permit to this resource, such as FTP, HTTP, NFS, and so on. Select **Any** to permit all services.
4. Select the initiator of the permitted service: Server, a Client, or Both.
5. Select the address object or address group for the resource.
6. Add the security device through which traffic can reach the protected resource:
 - a. In the Security Gateway area, click the icon to display the Security Gateway dialog box.
 - b. Select security device or device group
 - c. Select the security zone on the security device that contains the address objects.
 - d. Click **OK** to add the security gateway to the protected resource.

You can add multiple security gateways to provide redundant access for the protected resource.

Editing Protected Resources

You can edit protected resources to accommodate changes in your network:

- If you make changes to a protected resource object that is used in a VPN, NSM automatically generates new configuration and propagates your changes to all affected security devices.
- If you change the security device that protects a resource, NSM removes the previous security device from all affected VPNs and adds the new security device. However, NSM does not configure the VPN topology for the new security device—you must reconfigure the topology to include the new device manually.

Configuring IKE Proposals

In an AutoKey IKE VPN, you can use the Internet Key Exchange (IKE) protocol to generate and distribute encryption keys and authentication algorithms to all VPN nodes. IKE automatically generates new encryption keys for the traffic on the network, and automatically replaces those keys when they expire. Because IKE generates keys automatically, you can give each key a short life span, making it expire before it can be broken. By also exchanging authentication algorithms, IKE can confirm that the communication in the VPN tunnel is secure.

Because all security parameters are dynamically assigned, VPN nodes must negotiate the exact set of security parameters that will be used to send and receive data to other VPN nodes. To enable negotiations, each VPN node contains a list of proposals; each proposal is a set of encryption keys and authentication algorithms. When a VPN node attempts to send data through the VPN tunnel, IKE compares the proposals from each VPN node and selects a proposal that is common to both nodes. If IKE cannot find a proposal that exists on both nodes, the connection is not established.

IKE negotiations include two phases:

- In Phase 1, two members establish a secure and authenticated communication channel.
- In Phase 2, two members negotiate Security Associations for services (such as IPSec) that require key material and parameters.

By default, NSM includes several common IKE phase1 and phase2 proposals. To view these proposals, from VPN Manager, select **IKE Phase1 Proposals** or **IKE Phase2 Proposals**.

Creating Custom IKE Phase1 Proposals

Create a custom proposals for a specific combination of authentication and encryption that is not available in the predefined proposals, or to match the name of proposals on a non-security device.

To create a custom IKE Phase1 proposal, select **Custom IKE Phase** and click the icon. Enter a name and choose a color for the object, then configure the following settings:

- Authentication Method—Select the authentication method.

- Preshared Key. Use this option to generate an ephemeral secret and authenticate data using MD5 or SHA hash algorithms against the secret.
- RSA Certificate.
- DSA Certificate.
- Diffie-Hellman Group—The Diffie-Hellman group provides asymmetric encryption to encrypt the keys needed to decrypt the data. The larger the modulus of the group, the more secure the generated key is—and the more time it takes to generate the key. Select the group that meets your security requirements and user needs:
 - Group 1. Uses a 768-bit modulus.
 - Group 2. Uses a 1024-bit modulus
 - Group 5. Uses a 1536-bit modulus.
 - Group 14. Uses a 2048-bit modulus.
 - Group 19. Uses a 256-bit modulus.
 - Group 20. Uses a 384-bit modulus.
- Encryption Algorithm—Select the algorithm that meets your security requirements:
 - DES-CBC
 - 3DES-CBC
 - AES-CBC (128 Bits)
 - AES-CBC (192 Bits)
 - AES-CBC (256 Bits)



NOTE: Security devices use hardware encryption for DES and 3DES and use software encryption for AES.

- Hash Algorithm—Select the algorithm that meets your security requirements.
 - MD5. Authenticate data using Message Digest version 5.
 - SHA-1. Authenticate data with Secure Hash Algorithm-1.
 - SHA-2. Authenticate data with Secure Hash Algorithm-2 (minimum 256 bit).
- Lifetime—Enter the number of seconds before the key is regenerated. The default value is 28800 seconds (8 hours).

Click **OK** to add the custom IKE object to the management system.

Creating Custom IKE Phase 2 Proposals

Create a custom proposals for a specific combination of authentication and encryption that is not available in the predefined proposals, or to match the name of proposals on a non-security device.

- Perfect Forward Secrecy—PFS ensures that a single key permits access to data protected by that single key. The key used to protect transmission of data and the material used to create that key are used only once and are not used to derive additional keys. Select the DH group to encrypt the key:
 - No Perfect Forward Secrecy.
 - Diffie-Hellman Group 1.
 - Diffie-Hellman Group 2.
 - Diffie-Hellman Group 3.
 - Diffie-Hellman Group 14.
 - Diffie-Hellman Group 19.
 - Diffie-Hellman Group 20.



NOTE: You can create Phase 1 and Phase 2 proposals with Diffie-Hellman Group14 from VPN Manager>AutoKey IKE Parameters>Security>Phase2 Proposal>Security-Level (User-Defined).

You can only create Custom IKE Phase 1 and 2 proposals with Diffie-Hellman Group 14 on devices running ScreenOS 6.2 or later. On other devices, an error message is generated.

- Lifetime (Seconds)—Enter the number of seconds before the key is regenerated. The default value is 3600 seconds (8 hours).
- Lifesize (KB)—Enter the number of bytes permitted through the connection before the key is regenerated. A value of 0 (the default) means no limit.
- Encryption (ESP) or Authentication (AH) Algorithm.
 - Select ESP to configure encryption and authentication, then select the desired algorithms.
 - Select AH to configure authentication only, then select the desired algorithm.



NOTE: We strongly recommend that you do not use null AH with ESP.

Click **OK** to add the custom IKE object to the management system.

Configuring Dial-in Objects

Netscreen devices allow users to dialin and manage the box as a console. By switching the modem interface, you can both dial in, and dial out. You can use NSM to configure Dial-in details. You can create and edit white lists of allowed numbers, and black lists of

blocked numbers. You can set a policy for unknown CNIDs if you wish to change the default "denied" setting.

- [Creating a Dial-In Profile on page 471](#)
- [Linking the Dial-In Profile with the Device on page 471](#)
- [Setting the Time-out Period for the Modem Dial-In Authentication on page 471](#)

Creating a Dial-In Profile

1. Select **Object Manager > Dial-In**.
2. Select **Add Dial In Object**. The New Dial in window opens.
3. Click **+** in the Phone Settings table for either the White List or Black List. The New List Entry box opens.
4. Enter the new phone number.
5. Click **OK**.

Linking the Dial-In Profile with the Device

From the Device configuration editor:

1. Select **Network > Interface > Edit Serial Interface > Modem**.
2. Enter the name of the Dial-in profile in the **Dial-in Name** field.
3. Click **OK**.

Setting the Time-out Period for the Modem Dial-In Authentication

From the Device configuration editor:

1. Select **Device Admin > CLI Management**.
2. Set the time period in the **Modem Dial-In Authentication Time-out** field.
3. Click **OK**.

Configuring Border Signaling Gateway Objects

The Border Signaling Gateway (BSG) handles VoIP signaling in Junos OS based on policies you set. NSM allows you to define two shared BSG objects in the Object Manager that can be referenced in the BSG transaction rulebase in the Policy Manager. You can create, view, edit, and delete the following BSG objects from the Object Manager.

- **BSG Service Point object:** This is a polymorphic object for specifying the egress service point in a transaction term's route action. With the service point object, you can specify the device, a gateway in the device, and a service point in the gateway.
- **BSG Admission Controllers:** BSG Admission Controllers control Session Initiation Protocol (SIP) dialogs and transactions. You can define the following settings for both dialogs and transactions:
 - **Maximum Concurrent Number** This is the maximum number of concurrent dialogs that are allowed. You can set it from 0 to 100,000. Zero causes all calls to be rejected. The default setting is 100000
 - **Committed Attempts Rate** The maximum number of attempts per second to initiate an out-of-dialog transaction. You can set it from 0 to 100. The default setting is 100
 - **Committed Burst Size** The maximum number of dialogs allowed to burst above the committed-rate and still be accepted. You can set it from 0 to 200.

You can define one admission controller per gateway, and reference all of the admission controllers from the BSG transaction rulebase in the Policy Manager. Admission controller objects are listed on the transaction policy's shared-object menu, where you can drag and drop them into the transaction terms. When you import a device, the admission controller objects are also imported.

BSG objects are supported in Junos OS Release 9.5 and later. When updating devices running under earlier versions of Junos OS, the admission controller setting is dropped.

CHAPTER 9

Configuring Security Policies

Firewall rules define access to your network, including permitted services, users, and time periods. You can also use firewall rules to control the shape of your network traffic as it passes through the firewall or to log specific network events. Multicast rules permit multicast control traffic, such as IGMP or PIM-SM messages, to cross Juniper Networks security devices. Multicast rules permit multicast *control* traffic only; to permit *data* traffic (both unicast and multicast) to pass between zones, you must configure firewall rules.

Because all incoming and outgoing network traffic passes through your firewall, it is the ideal location to control the traffic flowing on your network. Creating security policies enables you to define what type of traffic should be permitted on your network, as well as how that traffic is treated while inside. A security policy can contain firewall rules (in the Zone and Global rulebases), multicast rules (in the Multicast rulebase), and IDP rules (in the Application Policy Enforcement (APE), IDP, Exempt, Backdoor Detection, SYN Protector, Traffic Anomalies, and Network Honeypot rulebases).

This chapter contains the following sections:

- [About Security Policies on page 474](#)
- [Creating a Security Policy on page 487](#)
- [Configuring Firewall Rules on page 492](#)
- [Configuring Multicast Rules on page 510](#)
- [Configuring Antivirus Rules on page 511](#)
- [Configuring Antispam Rules on page 512](#)
- [Configuring IDP Rules on page 512](#)
- [Configuring Application Policy Enforcement \(APE\) Rules on page 527](#)
- [Configuring Exempt Rules on page 535](#)
- [Configuring Backdoor Rules on page 538](#)
- [Configuring SYN Protector Rules on page 542](#)
- [Configuring Traffic Anomalies Rules on page 546](#)
- [Configuring Network Honeypot Rules on page 550](#)
- [Installing Security Policies on page 553](#)
- [Managing Rules and Policies on page 558](#)

- [Pre and Post Rules on page 572](#)
- [Polymorphic Objects on page 576](#)

About Security Policies

A security policy determines how your managed devices handle your network traffic. To display previously configured security policies, select **Configure > Policy Manager** and double-click **Policies**. When you edit a security policy, the name of that security policy appears in bold in the main navigation tree.

Using the Network and Security Manager (NSM) UI, you can configure rules in up to ten rulebases (Zone, Global, Multicast, IDP, Exempt, APE, Backdoor Detection, SYN Protector, Traffic Anomalies, and Network Honeypot) for each security policy.



NOTE: In the ScreenOS WebUI and CLI, a security policy is a single statement that defines a source, destination, zone, direction, and service. In NSM, those same statements are known as rules, and a security policy is a collection of rules.

After you create a security policy by building rules in one or more rulebases, you can assign that policy to specific devices. For information about assigning a policy to a device, see [“Assigning a Security Policy to a Device” on page 553](#).

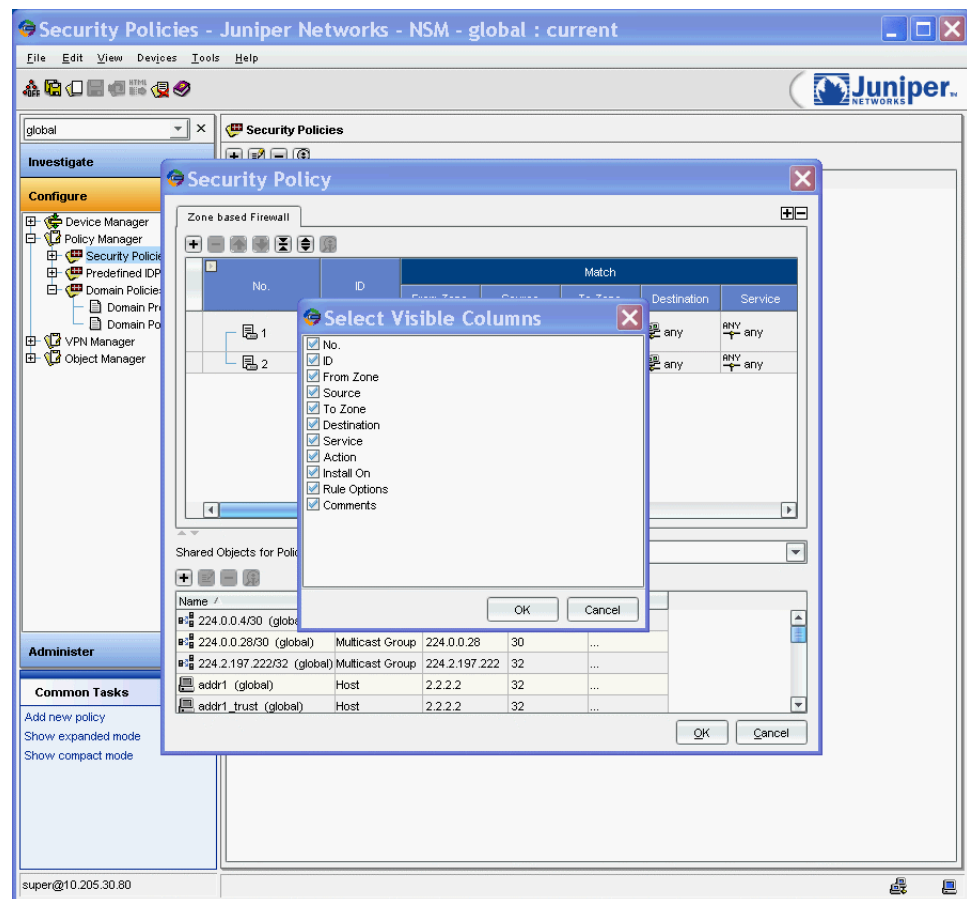
Viewing Rulebase Columns for a Security Policy

By default, each rulebase displays a subset of available columns for each rule. This mode, known as Compact Mode, contains columns in which you can configure typical rule parameters. To see all columns for the rulebase, change the mode of the security policy to Expanded: From the menu bar, select **View > Show Expanded Mode**. You can set a different mode for each security policy. You can also create Custom Mode views based on the columns shown in the Compact Mode or Expanded Mode view.

To create a Custom Mode view:

1. Select the mode from which you want to create a Custom Mode:
 - To create a custom mode based on the Compact Mode column options, select **View > Show Compact Mode**.
 - To create a custom mode based on the Expanded Mode column options, select **View > Show Expanded Mode**.
2. Move the cursor over a column header of the security policy. A small icon appears to the left above the No. column. Click on the icon to display the Select Visible Columns dialog box, as shown in [Figure 89 on page 475](#).

Figure 89: Displaying the Select Visible Columns Dialog Box



3. Select the visible columns (Destination, Service, Action, NAT, and so on) that you want to include in the Custom Mode view.
4. Click **OK**. The Custom Mode view shows the columns you selected from the Select Visible Columns dialog box.

Viewing and Editing Custom Policy Fields

NSM allows you to create multiple fields under **Rule Options**. You can customize these fields to save metadata, and you can edit and filter the values in each of these custom fields. You must create and save these custom policy fields as objects under the Object Manager before you can use them in policy. See [“Configuring Custom Policy Fields” on page 416](#) for details.

In extended mode, every custom field is displayed as a separate column nested under a header named **Custom Field**. In compact mode, the custom field values are listed in a single **Optional Field** column. Filters set in the **Optional Fields** column do not impact the custom fields.

Right-click on an individual custom field to edit or filter that particular value. A dialog box appears, displaying the values of the custom field in a tree structure. You can then search, add, delete or filter any value in that field.

About Rulebases

A rulebase is a set of rules that define how the managed device handles traffic. NSM supports three firewall rulebases and six IDP rulebases, as detailed in the following sections. A security policy can contain only one instance of any rulebase type.

By default, the predefined roles System Administrator, Domain Administrator, and IDP Administrator can view and edit all rulebases. The Read-Only System Administrator and Read-Only Domain Administrator can only view rulebases. When creating a custom role, you can include permissions to view or edit individual rulebases.

NSM supports the following firewall rulebases:

- **Zone**—Contains rules that apply to traffic from one specific zone to another. Create a firewall rule in the zone-specific rulebase when you need to control traffic between specific zones. The zone-specific rulebase can contain firewall rules and VPN rules and links.
- **Global**—Contains rules that are valid across all zones. Create a firewall rule in the global rulebase when you need to control specific traffic across the entire firewall. The global rulebase can contain only firewall rules.
- **Multicast**—Contains rules that enable IGMP proxy or PIM-SM multicast control traffic between zones.

NSM supports different kinds of IDP-capable devices that can provide firewall and IDP functionality: standalone IDP appliances, ISG gateways, J Series routers, SRX Series gateways, and MX Series routers.

NSM supports the following IDP rulebases:

- **IDP**—This rulebase protects your network from attacks by using attack objects to detect known and unknown attacks. Juniper Networks provides predefined attack objects that you can use in IDP rules. You can also configure your own custom attack objects.



NOTE: Juniper Networks updates predefined attack objects on a regular basis to keep current with newly-discovered attacks.

- **APE**—This rulebase is used by IDP devices to detect network traffic based on application signatures and to take specified action.
- **Exempt**—This rulebase works in conjunction with the IDP rulebase to prevent unnecessary alarms from being generated. You configure rules in this rulebase to exclude known false positives or to exclude a specific source, destination, or source/destination pair from matching an IDP rule. If traffic matches a rule in the IDP rulebase, IDP attempts to match the traffic against the Exempt rulebase before performing the action specified.

- **Backdoor Detection**—This rulebase protects your network from mechanisms installed on a host computer that facilitates unauthorized access to the system. Attackers who have already compromised a system typically install backdoors (such as Trojans) to make future attacks easier. When attackers send and retrieve information to and from the backdoor program (as when typing commands), they generate interactive traffic that IDP can detect.



NOTE: If you import an ISG2000 or ISG1000 gateway into NSM, the imported device configuration does not include the IDP, Exempt, or Backdoor rulebases.

- **SYN Protector**—This rulebase protects your network from SYN-floods by ensuring that the three-way handshake is performed successfully for specified TCP traffic. If you know that your network is vulnerable to a SYN-flood, use the SYN-Protector rulebase to prevent it.
- **Traffic Anomalies**—This rulebase protects your network from attacks by using traffic flow analysis to identify attacks that occur over multiple connections and sessions (such as scans).
- **Network Honeypot**—This rulebase protects your network by impersonating open ports on existing servers on your network, alerting you to attackers performing port scans and other information-gathering activities.

Rule Execution Sequence

The rules in all rulebases combine to create a security policy. Security devices process and execute firewall and VPN rules in the following order:

1. Zone rulebase
2. Global rulebase
3. Multicast rulebase

Managed devices process and execute IDP rules in the following order:

1. Exempt rulebase
2. IDP rulebase
3. APE rulebase
4. Backdoor rulebase
5. SYN Protector rulebase
6. Traffic Anomalies rulebase
7. Network Honeypot rulebase

About Rules

A rule is a statement that defines a specific type of network traffic. Traffic must meet the rule requirements before it is permitted to pass through the security device. By default, all security devices deny all traffic.

When traffic passes through the security device, the device attempts to match that traffic against its list of rules. Network traffic that matches this list of requirements is considered to “match” the rule, and the device performs the action specified in the rule. If any requirement is not met, the network traffic does not match, and is denied.

Using the NSM UI, you can create intrazone firewall rules, global firewall rules, multicast rules, VPN rules, and VPN links for all security devices. For ISG gateways, you can create IDP rules, APE rules, exempt rules, and backdoor detection rules. For standalone IDP appliances, you can create IDP rules, APE rules, exempt rules, backdoor detection rules, SYN protector rules, traffic anomalies rules, and network honeypot rules. NSM also supports J Series, SRX Series, and MX Series devices running Junos OS that support both firewalls and IDP policies. Each security policy (all rulebases combined) can contain a maximum of 40,000 rules.

About Firewall Rulebases

You create rules in the firewall rulebases to enable access across your networks by permitting or denying specific network traffic flowing from one zone to another zone. After you have added a device in NSM, you can create rules in the firewall rulebases of your security policy.

You can build multiple firewall rules in both firewall rulebases for a single device; these rules combine to create a security policy that determines how that device handles traffic. To simplify your security policy, use device groups to build access rules that apply to all your perimeter security devices, then apply the entire policy to the perimeter device group.



NOTE: When a firewall policy with network address objects is applied to Junos devices, the device update operation in NSM fails, because DMI devices do not support network address objects.

Firewall Rules (Zone and Global)

Within a firewall rule, you specify where the traffic is coming from, where it is going, and what service it is using. You can also use firewall rules to authenticate users, monitor network traffic flowing between zones, or set a schedule on a firewall rule that controls the time period that the rule is applied to network traffic.



NOTE: On Juniper Networks vsys devices, rules defined in the root system do not affect rules defined in virtual systems.

When creating firewall rules, consider the type, location, and functionality of each device in your network. Typically, a single security policy for multiple devices works well for devices that perform similar functions, such as perimeter firewalls. However, you might want to create a separate security policy per device when the management system contains separate administrators with regional responsibilities, or when you need to troubleshoot a device issue (use one security policy per device to enable an administrator to troubleshoot on one device without making policy changes on other devices).

A firewall rule must contain the following elements:

- **Direction**—The direction that the traffic flows between two zones; all traffic flows from a source zone to a destination zone. You can select any zone/zone group for source or destination; however, the zones must be valid for the security devices you select in the Install On column of the rule. You can also use zone exceptions to specify unique to and from zones for each device.
- **Source address**—The address that initiates the traffic.
- **Destination address**—The address that receives the traffic.
- **Service**—The application-level protocol that the traffic uses to transmit data.
- **Action**—The action the device performs when it receives traffic that matches the direction, source, destination, and service specified in the rule.
- **Install On**—The device on which the firewall rule is installed. You can install the same rule on multiple devices.

To begin configuring firewall rules for your managed devices, see [“Configuring Firewall Rules” on page 492](#).

Validating Firewall Policies

NSM firewall policy rules are loaded on NSM after they are validated. Validation displays warnings in the rules. Validation of all the rules during rules loading takes a very long time and adversely affects NSM performance. To improve the performance, disable validation of firewall policy during loading by using the option **Disable GUI validation > firewall policy** under **Tools > Preferences**. For more information about disabling GUI validation, see [“Disabling GUI Validation” on page 24](#). This option disables validation of warnings and validates errors only.

To validate warnings on firewall rules after all rules are loaded, click **Validate** at **Policy Manager > Policy > Zone based Firewall**. [Figure 90 on page 480](#) shows the zone-based firewall validation option on the toolbar. When the loaded firewall policies are validated, NSM scans all the firewall rules under the policy and validates the rules.

Figure 90: Firewall Validation Option

No.	ID	From Zone	Source
1	1	trust	A

Figure 91 on page 480 shows the validation of firewall rules for errors when you disable GUI validation in the preferences. If you do not disable this option, validation is triggered and warning messages are displayed.

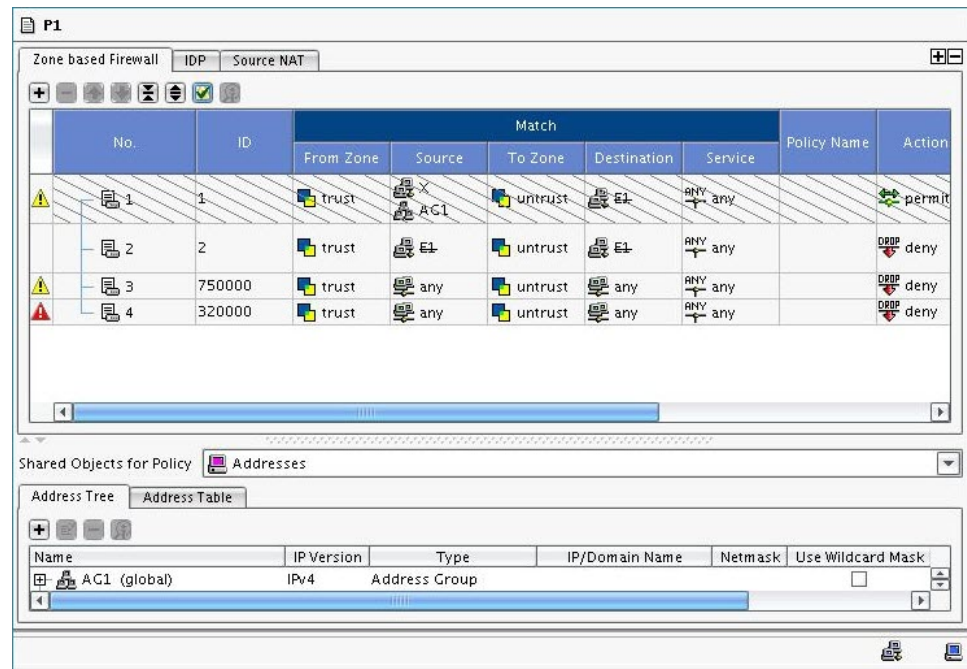
Figure 91: Firewall Error Validation Option

No.	ID	From Zone	Source	To Zone	Destination	Service	Policy Name	Action
1	1	trust	AC1	untrust	E1	ANY any		permit
2	2	trust	E1	untrust	E1	ANY any		deny
3	750000	trust	any	untrust	any	ANY any		deny
4	320000	trust	any	untrust	any	ANY any		deny

Name	IP Version	Type	IP/Domain Name	Netmask	Use Wildcard Mask
AC1 (global)	IPv4	Address Group			

Figure 92 on page 481 shows the validation of firewall rules for warnings after firewall policy is loaded and validated.

Figure 92: Firewall Warnings Validation Option



NOTE: This feature is also supported on zone based firewall, rulebase, IDP firewall, and exempt rulebase.

Global Firewall Policies (Central Policy Mode)

SRX Series devices running Junos OS Release 11.2 and later releases support global firewall policies. You can create or delete a new global rulebase for security policies on these devices. If a global policy is already present on a device, the policy will be imported under the global rulebase. All addresses used in the global policy and in the zone-based firewall will be updated to the global address book. If a global firewall policy is not present on the device and the **Use Global address book** option is selected, addresses used in the zone-based firewall will be sent to the global address book at the next update.



NOTE:

- You can create a global firewall policy without specifying From and To zones.
- Policy name is mandatory for Global Firewall.
- When you enter duplicate policy name in global rulebase for Junos OS an error is displayed.

Type of Global Firewall Policy

Global firewall policies support IPv4, IPv6, and mixed rules. The options **any-ipv4**, **any-ipv6**, and **Mixed Rule** are available when you right-click on the source and destination columns to add **any-ipv4** and **any-ipv6** addresses. For more information about adding mixed rules for SRX Series devices see, “[Mixed Rule for Zone-based Firewall Policy and Global Firewall Policies](#)” on page 482.

A global firewall rule must contain the following elements:

- Source address— Address that initiates the traffic.
- Destination address—Address that receives the traffic.
- Service—Application-level protocol that the traffic uses to transmit data.
- Action—Action the device performs when it receives traffic that matches the source, destination, and service specified in the rule.
- Install On—Device on which the firewall rule is installed. You can install the same rule on multiple devices.

Mixed Rule for Zone-based Firewall Policy and Global Firewall Policies

The mixed-rule option allows you to add IPv4 and IPv6 addresses to zone-based firewall rules and the global firewall policy as a single rule. You can either add a mixed rulebase using the policy manager or promote the existing zone-based firewall rule as a mixed rule.

The default address option in a mixed rulebase is **any**. However, **any-ipv4** and **any-ipv6** addresses can be configured for a rule.



NOTE:

- The existence of two or more identical mixed rules under a policy displays a validation error.
 - The mixed-rule option is supported only on SRX Series devices.
-

Global Address Book Overview

In Junos OS Release 11.2 and later releases, the address book is moved from the zone level to the device global level. This permits objects to be used across many zones and avoids inefficient use of resources. The Network and Security Manager (NSM) manages its address book at the global level, assigning objects to devices that are required to create policies. You can enable the **Global address Book** option for the compatible devices, for NSM to push address objects used in policies to the global address book of the device.



NOTE: NSM imports addresses present in the global address book to NSM.

Understanding Global Address Books

The global address book is supported in Junos OS Release 11.2 and later releases.

- An address book is not configured within a specific zone; therefore, one address book can be associated with multiple zones.
- If a global address book is defined, you cannot create zone-based address books.
- By default, there is an address book, called global, associated with all zones.
- A zone can be attached to only one address book in addition to the global address book, which contains all zones by default.
- Address name overlaps are possible between the global address book and custom address book.
- NAT rules can use address objects only from the global address book. They cannot use addresses from custom address books.



NOTE: Custom Address book is not supported in NSM Central Policy Mode.

Nested Address Group Support

In Junos OS versions earlier than Release 11.2, nested address group were not supported in NSM. Because of this, service groups in NSM were flattened to an application set when they were pushed to a device. This caused inefficient of object resource usage. Beginning in Junos OS Release 11.2, nested address groups (address set) are supported. In NSM 2012.2 Release NSM supports, child address group members within a address group for J Series and SRX Series devices. Nested address group support is present for the Zone based address Books and Global Address book.



NOTE: NSM resolves the child address group members to the parent group when you perform a summarized delta configuration or device update for J Series and SRX Series device for release lower than 11.2.

Nested Service Group Support for DMI Devices

In Junos OS versions earlier than Release 11.2, nested service group were not supported in NSM. Because of this, service groups in NSM were flattened to an application set when they were pushed to a device. This caused inefficient object resource usage. Beginning in Junos OS Release 11.2, nested service groups (application set) are supported in NSM 2012.2 Release. NSM now supports child service group members within a service group for an SRX Series device. With this enhancement, you can group multiple service group and also use them in zone based and global rule bases.



NOTE: NSM resolves the child service group members to the parent group when you perform summarize configuration delta or perform a device update for J Series and SRX Series device for release lower than 11.2.

Services Offload Option on High-End SRX Series Devices in NSM

The services-offload feature is available on SRX3400, SRX3600, SRX5600, and SRX5800 Services Gateways in Junos OS Release 11.4R1 and later releases. NSM deletes the configuration if services offload is enabled on SRX Series for the branch devices. Services offload is available under Rule Options or when you right-click or double-click on a rule. Available actions for a rule are **permit**, **deny**, **reject**.



NOTE:

- The services-offload option is available only if the rule has the action as permit. If the action field is configured as deny or reject, then NSM displays the following warning message: **Service Offload: Only Permit is allowed when the Services-offload is Enabled.**

VPN Links and Rules

The rules for your rule-based VPNs appear in the Zone rulebase.

- Use VPN Links for VPNs created in VPN Manager—By default, VPN Manager autogenerated rules are implicitly executed as the first rule in the Zone rulebase, even though they do not appear. Because VPN Manager autogenerates the access rules for the VPN, you do not need to manually create them in the rulebase itself. However, to specify the exact location of the autogenerated rules in your rulebase, you can add a VPN link anywhere in the Zone rulebase.
- Use VPN Rules for VPNs created manually—If you did not use VPN Manager to create a rule-based VPN, you must manually add the VPN rules to create the VPN tunnel. You can place VPN rules anywhere in the Zone rulebase.

Because route-based VPNs are on always-on connection between two or more termination points, you do not need VPN rules to create the routing-based VPN tunnel. However, you might want to create access rules to control the flow of traffic in a routing-based VPN tunnel.



NOTE: VPN rules are not validated by rule validation. Only firewall rules are validated by rule validation.

About Rule Groups

A rule group is a user-defined grouping of rules within the Zone rulebase. Combining rules into a rule group can help you better manage rules. For example, you might want to

combine your VPN rules in a single rule group, or combine all rules that manage traffic from a specific interface on the device.

You can add, edit, and delete rule groups; however, deleting a rule group also deletes all rules within that group. You can create multiple rule groups (40,000 rules max in a security policy). NSM supports one level of rule groups; you cannot create a rule group within a rule group.



NOTE: Rule groups can be created for all Policy Manager rulebases except global and APE rulebases.

For information about rule groups, see [“Using Rule Groups” on page 562](#).

About the Multicast Rulebase

By default, security devices do not permit multicast control traffic such as IGMP or PIM-SM messages. If you run IGMP proxy or PIM-SM on your network, you must configure rules in the Multicast rulebase to explicitly permit multicast control traffic between zones.

You can also configure multicast rules to translate multicast addresses. For example, to translate a multicast group address in an internal zone to a different address on the outgoing interface, specify both the original multicast address and the translated multicast group address in a multicast rule.

When you create a multicast rule, you must specify the following:

- Source zone—The zone from which traffic initiates.
- Destination zone—The zone to which traffic is sent.
- Multicast group—The multicast group or access list that specifies the multicast groups for which you want the security device to permit multicast traffic.

Multicast rules control the flow of multicast control traffic only. To permit data traffic (both unicast and multicast) to pass between zones, you must configure rules in a firewall rulebase.

To begin configuring multicast rules for your managed devices, see [“Configuring Multicast Rules” on page 510](#).

About IDP Rulebases on ISG Family Devices

For IDP-capable security devices, such as the ISG Series gateways running ScreenOS 5.0-IDP and later, you can enable IDP in a zone or global firewall rule to direct permitted traffic to the IDP rulebases. If you do not enable IDP in a firewall rule for a target device, you can still configure rules in IDP rulebases, but you cannot apply the IDP rules when you update the security policy on the target security devices.



NOTE: If you configure a J Series router to be managed in central manager mode and you select an IDP rulebase rule and specify an IP address for the source and destination instead of “any,” the rule policy is not be pushed to the router.

When configuring IDP in a firewall rule, consider the following:

- The firewall action must be permit. You cannot enable IDP for traffic that the security device denies or rejects.
- Only traffic that is permitted by the firewall rule is passed to the IDP rulebases. The security device does not forward denied traffic to IDP rulebases.
- You cannot configure deep inspection (DI) for the rule; when you install the IDP license on an ISG2000 or ISG1000 device running ScreenOS 5.0–IDP and later, DI is automatically disabled on the device.



NOTE: The Attack Profile Settings only apply to the DI feature on security devices.

To enable IDP in a firewall rule, right-click in the Rule Options column for the zone or global firewall rule and select **DI Profile/Enable IDP**. The DI Profile/Enable IDP dialog box appears (by default, IDP is disabled). Select **Enabled** to enable IDP for traffic that matches the firewall rule, then select the mode in which you want IDP to operate:

- In inline mode, which is the default, IDP is directly in the path of traffic on your network and can detect and block attacks. For example, you can deploy the device with integrated Firewall/VPN/IDP capabilities between the Internet and an enterprise LAN, WAN, or special zones such as DMZ. This is the default mode.
- In inline tap mode, IDP receives a copy of a packet while the original packet is forwarded on the network. IDP examines the copy of the packet and flags any potential problems. IDP's inspection of packets does not affect the forwarding of the packet on the network.

About IDP Rulebases on Standalone IDP Sensors

Standalone IDP Sensors only support IDP-specific rulebases—not firewall rulebases. You do not need to direct traffic to the IDP rulebases; all traffic passing through a standalone IDP Sensor is automatically examined for IDP-related issues.

You must configure the Sensor directly to operate in inline or sensor mode. Refer to the IDP Installer's Guide for configuration procedures.

- In inline mode, a Sensor is directly in the path of traffic on your network and can detect and block attacks. For example, you can deploy the Sensor between the Internet and an enterprise LAN, WAN, or special zones such as DMZ.
- In sensor mode, a Sensor receives a copy of a packet while the original packet is forwarded on the network. The Sensor examines the copy of the packet and flags any

potential problems. The Sensor's inspection of packets does not affect the forwarding of the packet on the network.

Enabling IPSec Null Encryption for IDP Inspection

In NSM release 2007.3 and later, you can enable or disable de-capsulation of IPSec packets with null encryption. This feature applies to ISG devices with IDP functionality running ScreenOS 6.1 and later..

To enable IPSec null encryption for IDP inspection:

1. In the NSM main navigation tree, click **Device Manager > Devices**.
2. Select an ISG device with IDP functionality running ScreenOS 6.1 or later, and click the Edit icon. The Devices dialog box appears.
3. Click **Security > IDP SM Settings**.
4. Click the **Run-time Parameters** tab.
5. Enable IPSec null encryption by selecting the **Enable IPSec ESP-NULl decapsulation support** check box, and then click **OK**.

Managing Security Policies

After you have created a security policy, you can:

- Modify individual rules in each rulebase, such as changing rule order (determine the order that rules are applied to network traffic by placing the rules in the desired sequential order), disabling a rule, negating source or destination addresses (ScreenOS 5.x devices only), and so on.
- Validate a security policy before installing it on your managed devices.
- Merge multiple security policies into a single policy for easier management. For example, after importing (or re-importing) devices into the management system, you might want to merge their imported policies into a single policy for all devices.
- Export the policy to an HTML file.

For information about managing your security policies, see [“Managing Rules and Policies” on page 558](#).

Creating a Security Policy

When creating a security policy, consider the following:

- **Objects**—Before creating a security policy, you should first use Object Manager to create objects representing your network components, custom services, custom attack objects, and so on. You use these objects when configuring rules within the policy.

If you are running an IDP-enabled device, you can use the profiler to monitor the traffic of interest on your network.

- **Pre-Existing Policies**—When creating a new policy, you can use an existing policy as a template. NSM comes with a collection of predefined IDP policies that you can use, or you can use a policy that was created earlier by your organization.
- **Rulebases**—When you initially create a security policy, only the Firewall rulebase and possibly the IDP rulebase appear by default. To create a rule in other rulebases, you must manually add those rulebases to the policy.

The following sections detail these options.

Configuring Objects for Rules

Objects are reusable logical entities that represent specific settings, configurations, or external pieces of hardware. You can reuse objects in multiple areas in the NSM GUI. Within rules, you use objects to define the source, destination, and service, as well as to specify settings for rule options, such as Web Filtering or attack protection.

For some object types, such as services and IDP attack objects, NSM contains a database of predefined objects. If the predefined objects do not meet your networking requirements, you can create custom objects and add them to the object database. For other object types, such as address objects, DI profiles, and Global MIPS, no predefined objects exist; before you can use one of these objects in a rule, you must create the object in Object Manager.

Applying the Same Object to Multiple Rules

You can apply the same object (column value) to a selection of policy rules. Rule groups must be in an expanded state to apply the same object to the rules of a rule group. Columns that disallow duplicate values, such as the rule ID and No. columns cannot be used to apply the same object to a selection of rules.



NOTE: You cannot apply the same object to a selection of rules for Predefined policies, VPN rules, or Central rules on a Regional Server regular policy.

To apply the same object to a selection of policy rules:

1. Select all the rules to which the column value will be applied:
 - To select all the rules in a rulebase, click on any rule in the rulebase and press **Ctrl + a**.
 - To select a contiguous range of rules in a rulebase, press **Shift + Ctrl** and select the rules.
 - To select a noncontiguous range of rules in a rulebase, press **Ctrl** and select the rules.
2. Right-click on the column value of the rule that you want to apply to the selected rules and select **Apply value to selected rules** from the menu.

The selected column value is applied to all selected rules.



NOTE: After you select the rules, a right-click on any column value displays the menu “Apply Value to selected rules,” and no other menu options are available for the selected column value.

Naming of Address Objects in a Security Policy That References Devices Running ScreenOS or Junos OS

Device updates might fail when a policy that references address objects for ScreenOS devices is assigned to a J Series device or an SRX Series device because the address object naming conventions in Junos OS are more restrictive than the naming conventions in ScreenOS. For devices running Junos OS, the address object name must be a string that begins with a letter and consists of letters, numbers, dashes, and underscores. For devices running ScreenOS, the address object name can include numbers, characters, and symbols. To ensure that a device running Junos OS can use the address objects referenced by the security policy that is assigned to the device, all address objects in that policy must follow the address object naming conventions for Junos OS. If the policy that is assigned to a device running Junos OS contains preexisting address objects for ScreenOS devices, these address objects must be renamed to follow the address object naming conventions for Junos OS.

Using the Policy Filter Tool

NSM provides a Policy Filter tool to filter policy rules-based on one or more filter conditions specified for rule attributes. One filter can contain several filter conditions for different attributes. The filter only applies to the current selected rulebase. The filter results are displayed in the same rulebase. Rules that do not match filter conditions are hidden. In the firewall rulebase, only open rule groups are filtered. When a filter is set and a closed rule group is expanded, only rules that match the filter will be displayed in the group. For information about using the Policy Filter tool, refer to the NSM Online Help.

Filtering the Comment Field

You can use filters for the comments field of your policy. By default, search finds an exact match unless used with a regular expression.

For example, you have two rules with the following two comments: test1 and juniper,\ntest1. If you want to find all the rules that have test1 in the comments field, you must use a regular expression. If you do not use the regular expression checkbox, the search returns rules with comment test1 only.

If you want to find all rules that end with the string test1, you can use one of the following regular expressions:

- `.*test1|.*\ntest1`
- `(.*|.*\n)test1`

Using a Predefined IDP Policy

When you create a new IDP security policy, you can select from the following predefined policies or use the Policy Creation Wizard, as described in the next section.



NOTE: IDP predefined policies are empty after an attack update. Relaunch the GUI to reinstate the policies.

For the standalone IDP Sensor and ISG with IDP devices, these policies are a good starting point for many common usage scenarios.

NSM includes the following security policy templates:

- `all_with_logging`—Includes all attack objects and enables packet logging for all rules.
- `all_without_logging`—Includes all attack objects but does not enable packet logging.
- `dmz_services`—Protects a typical DMZ environment.
- `dns_server`—Protects DNS services.
- `file_server`—Protects file sharing services, such as SMB, NFS, FTP, and others.
- `getting_started`—Contains very open rules. Useful in controlled lab environments, but should not be deployed on heavy traffic live networks.
- `idp_default`—Contains a good blend of security and performance.
- **Recommended**—Contains only the attack objects tagged as “recommended” by Juniper Networks security team. All rules have their Actions column set to take the recommended action for each attack object. By default, this policy is loaded onto all new IDP Sensors when they are added to NSM with the Add Device Wizard.
- `web_server`—Protects HTTP servers from remote attacks.

Each security policy template contains rules that use the default actions associated with the attack object severity and protocol groups. You should customize these templates to work on your network by selecting your own address objects as the Destination IP and choosing IDP actions that reflect your security needs.

Using the Policy Creation Wizard

This wizard guides you through the policy creation process. Use the wizard to specify the type of device the policy is for and the level of security you want. You can create a policy containing a zone-based firewall rulebase with one any-any-deny rule and/or an IDP rulebase. All other rulebases are optional and can be added to the policy based on need and access control permissions.

If you are logged in as an IDP Administrator, firewall-only rulebases are not available.

The Policy Creation wizard lets you select policies for the following devices:

- **Firewall/VPN**—Select this option to create a new policy containing a zone-based firewall rulebase with one any-any-deny rule. This option has only one set.

- **Stand Alone IDP**—Select this option to create a new policy containing the IDP rulebase.
- **Integrated Security Gateways/Security Routers**—Select this option to create a new policy containing a zone-based firewall rulebase with one any-any-permit IDP enabled rule as well as the IDP rulebase.



NOTE: If you do not have appropriate access-control permission and you attempt to create a policy, the wizard returns an error message stating that you do not have access to create rulebases.

In this example, you create a standalone IDP security policy that logs all levels of attack (Critical, Major, Minor, Warning, and Info) but drops connections only for critical and major attacks.

1. Click **Policies**, then go to the **File** menu and select **New Policy**.
2. Give the policy a name and add comments (optional), then click **Next**.
3. Select **Create New Policy for** (the default selection). Uncheck **Firewall/VPN Devices** and check **Stand Alone IDP Devices**, then click **Next**.
4. Select **Configure IDP Policy**, then click **Next**.
5. Check the boxes and select **Enable Logging** for all attack levels. Select **Drop Connection** for critical and major attacks. Click **Next** twice to continue.
6. Select the device to which you want to assign this policy, then click **Next**.
7. Click **Finish**.

Adding Rulebases

Security policies start with a minimum of rules and rulebases. You can add additional rules to the rulebases as needed.

To add a rulebase:

1. In the main navigation tree, select **Policies**, then double-click the policy name in the Security Policies window.
2. Click the Add icon in the upper right corner of the Policy window and select **Add <name> Rulebase**. The rulebase tab appears.
3. Configure a rule in the rulebase by clicking the Add icon on the left side of the Security Policy window. A default rule appears.
4. Add a new rulebase by clicking the Add icon in the upper right corner of the Security Policy window, then select the rulebase you want to add from the menu. You cannot add a rulebase more than once, so only rulebases that are not already in the policy are displayed.

The following sections explain how to configure rules in each rulebase.

Configuring Firewall Rules

The firewall rulebases enable you to create zone and global firewall rules that control the flow of traffic on your network. You can configure the following settings for a firewall rule:

- [Defining Match for Firewall Rules on page 492](#)
- [Defining Actions for Firewall Rules on page 496](#)
- [Selecting Devices for Firewall Rules on page 497](#)
- [Configuring Firewall Rule Options on page 498](#)
- [Comments for Firewall Rules on page 510](#)

For each rule, you must configure the rule parameters for the Match columns. The remaining columns are optional; however, the more specific you can be in defining rule parameters in each column, the more efficient your security policy can be when protecting your network.

Defining Match for Firewall Rules

A firewall rulebase controls traffic flow on your network, from one network component to another network component. To do this, the firewall must know the path that the traffic takes to reach its destination and the service the traffic uses to get there.

When creating your firewall rules, you must specify the areas in your network that the traffic passes through. These areas include the network components that originate and receive the traffic, and the firewall zones the traffic passes through. For firewall rules:

- The Destination Address, Source Address, Service, and Action are required for all rules in the Zone and Global rulebases.
- The To Zone, From Zone, and service are required for rules in the Zone rulebase.

You can create IPv6 rules with specific IPv6 source and destination addresses using the **Select Address Dialog** box. In this dialog box, you can populate hosts, networks, group addresses and polymorphic objects based on the context of the IP version selected. The policy filter is also enabled to support IPv6 addresses.

The following sections detail the Match columns of a firewall rule.

Configuring Source and Destination Zones for Firewall Rules

In the Zone rulebase, you create firewall rules to enable traffic to flow between zones (interzone) or between two interfaces bound to the same zone (intrazone). You must create zones on your device before you can create a rule for that device. In a single rule:

- You must select a single zone for the source zone and a single zone for the destination zone. These zones must be available on the devices covered by your policies
- You can also select multiple zone exceptions for both source and destination zones. A zone exception includes a specific zone and the device that contains that zone.

- You cannot create a rule that controls traffic between zones shared by vsys devices or by devices in an NSRP configuration.

In addition to the security zone, you can now also configure the self zone as the source zone in the security policy for all non-vsys devices running ScreenOS 6.2 and later. If you choose "self" as the source zone, then you must also configure the source address as "any". The system validates devices on which security policies with source zone "self" are configured. A validation error is generated for devices running versions below ScreenOS 6.2.

The Global rulebase does not contain source and destination zone columns. Because global rules permit or deny traffic flow between all zones on a device, both the source and destination zones are global and are not displayed.



NOTE: You can also configure "shared zones." The NSM Policy Manager uses shared objects, also known as "polymorphic objects," including zones to define various components of a policy rule. For more information, see ["Central Manager" on page 675](#).

Adding a Zone Group Object to the Firewall Policy

To add a zone group object to a firewall policy:

1. In the navigation tree, select **Policy Manager > Policies**.
2. Select the policy. The policy details appear in the main display area.
3. Right-click **From Zone** and click **Select Zone**. The **Select Zone or Zone Group** dialog box appears.
4. Do one of the following:
 - If the **Select Zone Object** drop-down list is displayed, select the zone group object from this drop-down list.
 - If the **Select Zone Object** drop-down list is not displayed, clear the **Use in Device Zone** check box.

The **Use Zone Group** check box appears. Select this check box and select a zone group object from the **Select Zone Object** drop-down list. If no zone group is available, you need to create a zone group object using the Add icon in the **Select Zone or Zone Group** dialog box.

5. Click **OK**.
6. Right-click **To Zone** and click **Select Zone**. The **Select Zone or Zone Group** dialog box appears.
7. Repeat Step 4 and Step 5.



TIP: When a policy includes a zone group object, make sure:

- The policy ID value is greater than **750000** and less than **998499**. Otherwise, NSM displays a validation error.

When you add a zone group object to a policy, you have to manually change the policy ID to be within this range.

- The zone group object contains at least one zone that is within the device to which the policy will be applied. Otherwise, the device update will fail. Only policies with zones that are relevant to a device are pushed during a device update.

Configuring Source and Destination Addresses for Firewall Rules

You create firewall rules to enable traffic to flow between two network components. In the NSM system, address objects are used to represent the components on your network: hosts, networks, and servers. When you add the address object to the rule, you are assigning it to a security zone on your security device.

You can add predefined address objects for the network components that originate and receive the traffic, or configure them as you create a firewall rule to control traffic between those components:

- To configure an address object as you are configuring the Source and Destination components of a rule, right-click in the Source or Destination column of a rule and select **Add Address**. Next, click the Add icon at the top of the New Source Addresses or New Destination Addresses dialog box and configure the desired address object.
- You can add an entire address group or select an individual address object from within the group.



TIP: When a Policy Manager tree table view includes an address group or service group, you can view the object (leaf member) count for the group by hovering over the group with the mouse. This feature is also supported for polymorphic objects in the address or service object category.

You can also negate all address objects in the source or destination columns of a rule. When the source or destination is negated, NSM considers all address objects defined for the current domain except the negated objects as part of the source or destination for that rule. To negate the source or destination, you must have previously added one or more address objects to the source or destination column of a rule.

You can add global MIP and VIP objects as the source or destination address in a rule; however:

- When installing the rule on devices running ScreenOS 5.0 and later, you can add multiple MIPs.
- When installing the rule on devices running ScreenOS 5.3 and later, you can add multiple MIPs and VIPs.

- When installing the rule on devices running ScreenOS 5.0 and later, you can add a single MIP object per rule. To use multiple MIP objects for these devices, you must use a separate rule for each global MIP object.

If you select multiple MIP or VIP objects in the source or destination column of a rule that includes devices running non-ScreenOS 5.3 and later in the Install On column, a validation message appears, indicating that those devices do not support multiple MIPs or VIPs within a single rule.

To control incoming Internet traffic to your trusted network, set the From Zone to Untrust and the To Zone to **Trust**. Set the source address as **any** and the destination to the address object that represents your trusted network.

To create a broader rule that controls traffic between multiple network components, create address object groups and use them in your firewall rules as you would other address objects. However, because security devices running ScreenOS 5.0 and later apply firewall rules to each address object separately, using address object groups can quickly decrease the number of available internal logical rules. If you must use address groups for both the source and destination, ensure that these groups are as small and as specific as possible.

To control traffic from your Marketing servers to your Engineering Servers, set the To Zone to Engineering and the From Zone to Marketing. Set the source address as the address group object that represents your Marketing servers, and the destination address to the address group object that represents your Engineering servers.

The more specific you are in defining the source and destination address in a firewall rule, the better your firewall performance will be.

To permit incoming traffic to your Engineering department network from any network except the Sales network, set the From Zone to Untrust and the From Zone to Trust. Set the source address group as the address group that represents Outside Sales network, and the destination address to the address group the represents your Engineering server network. Finally, right-click inside the source address column for the rule and select **Negate**.

Support for Any-IPv6 as a Source Address

With NSM support for any IPv6, you can now configure ISG devices running ScreenOS 6.2-IDP and later, and devices running Junos 10.2 and later to inspect data containing IPv6 addresses. The keyword "Any-IPv6" has been added to the IDP and firewall policies. In the context of source and destination addresses, the previous keyword "Any" will be treated as "Any-IPv4" on the device. You can continue to configure policies for IPv4 addresses.

To enable IPv6 functionality, you should set the environment variable IPv6 on the device to "yes" and then reboot the device. Since NSM does not manage environment variables, you cannot set this in NSM.

The Any-IPv6 functionality is supported on ISG family devices running ScreenOS 6.2-IDP and later versions, and devices running Junos 10.2 and later versions.

Configuring Services for Firewall Rules

Services are application layer protocols that define how data is structured as it travels across the network. In NSM, service objects represent the services running on your network. In a firewall rule, you specify which services are supported by the destination address object.



NOTE: All services rely on a transport layer protocol to transmit data. NSM includes services that use TCP, UDP, RPC, and ICMP transport layer protocols.

NSM comes with several service objects based on industry-standard services already created for you. You use these predefined service objects in firewall rules to specify the services that traffic can use to traverse your network.



TIP: When a Policy Manager tree table view includes an address group or service group, you can view the object (leaf member) count for the group by hovering over the group with the mouse. This feature is also supported for polymorphic objects in the address or service object category.

To control FTP traffic from the Engineering Server in the trust zone to the corporate Web Server in the DMZ zone, select the FTP, HTTP, IMCP ANY, and TELNET service objects.

You can create your own service objects to use in rules using the Object Editor, such as service objects for protocols that use nonstandard ports.

If you use a nonstandard port (8080) for your HTTP services, create an HTTP service object on port 8080. Add this service object to your firewall rule. NSM uses the specified service object, HTTP on port 8080, and considers all connections to TCP/8080 to be HTTP connections.

If the service of the network traffic matches a service selected in the rule, the firewall performs the action.



NOTE: For firewall rules installed on a ScreenOS 5.x device, if you use a custom service to relocate an application to a nonstandard port, you must also enable the Application option in the Rule Options > Miscellaneous > ScreenOS 5.x devices. For details, see [“ScreenOS 5.x and Later Options” on page 503](#).

Defining Actions for Firewall Rules

You can specify the action that your security device performs against traffic that matches the zones, address objects, and services specified in the firewall rule. You can set different actions for each rule:

- **Permit**—The managed device permits the traffic to pass through the firewall to its destination address.
- **Deny**—The managed device does not permit the traffic to pass through the firewall and drops all associated packets. No notification is returned to the sender.
- **Reject**—The managed device does not permit the traffic to pass through the firewall and drops all associated packets. For TCP and UDP packets, the device returns a notification message to the packet sender:
 - When the device drops a TCP packet, it returns a TCP RST packet to the sender.
 - When the device drops a UDP packet, it returns an ICMP port unreachable error to the sender.

For non-TCP and non-UDP packets, no notification is returned to the sender.

When you permit traffic, you can also:

- Use logging to monitor suspicious or abnormal uses of permitted traffic (such as excessive Web surfing).
- Use Antivirus to detect viruses in permitted traffic.
- Use Web Profiles to detect and prevent access to malicious or undesirable URLs.
- Use DI Profiles to detect and prevent attacks in permitted traffic.

For J Series and SRX Series devices, you can also use the NSM GUI to enable or disable DI IDP and Application Services. To use this feature:

1. Select a zone based firewall policy and right-click on the Rule Options column.
2. When the DI/Enable IDP/Appl Srvcs dialog box appears, select the applicable options.
 - Attack Profile Settings — Select an option.
 - IDP Option — Keep the Enabled setting or select Disabled.
 - Enabled — Keep Inline or select Inline Tap.
 - Application Services — Keep None, select Redirect WX, or select Reverse Redirect WX.
3. Click **OK** to save your settings.

The new settings (for example, “RWX”) appear next to the Rule Option entry.

Selecting Devices for Firewall Rules

In the install on column, select the devices that receive and use this rule. You can select multiple security devices on which to install the firewall rule. After you have created the security policy and assigned it to a device, NSM installs the rule only on the devices specified in the Install Column of the rule, enabling you to use a single security policy for multiple security devices.

To see the exact rules that are applied to a specific device, in Device Manager, right-click a device and select **Policy > View Pending Device Policy**.



NOTE: If a device specified in the Install column does not support a specific rule option configured for the rule, you can still install the security policy on the device, but the rule option is not enabled for that device. Additionally, during policy validation, a warning appears for each unsupported rule option. For details, see [“Validating Security Policies” on page 554](#).

Configuring Firewall Rule Options

Rule options enable you to configure additional protection mechanisms and other miscellaneous features. You can configure the following rule options:

- [Enabling NAT on page 498](#)
- [Enabling GTP for Firewall Rules on page 499](#)
- [Configuring Traffic Shaping in a Security Policy on page 499](#)
- [Enabling Logging and Counting for Firewall Rules on page 501](#)
- [Miscellaneous on page 502](#)
- [ID on page 503](#)
- [Configuring Web Filtering for Firewall Rules on page 504](#)
- [Configuring Authentication for Firewall Rules on page 505](#)
- [Configuring Antivirus for Firewall Rules on page 506](#)
- [Configuring a DI Profile/Enable IDP for Firewall Rules on page 507](#)
- [Configuring the Session Close Notification Rule on page 509](#)

To quickly configure all rule options, right-click the **Rule Options** column and select **Configure All Options**. The Configure Options dialog box appears; select the option tab you want to configure for the rule.

Enabling NAT

You can configure a policy-based network address translation (NAT) for a firewall rule. NAT enables the security device to translate the IP address of incoming or outgoing traffic so that the packets are routeable on the network.

Edit Source NAT

You can configure the security device to translate the source IP address:

- To translate the source IP address using a predefined range of IP addresses, select **NAT** and choose a Dynamic IP pool (DIP) object. For each matching packet, the device translates the original source address into a IP address selected from the DIP pool.
- To translate the source IP address using the IP address of the outgoing interface on the security device, select **Use Interface**.

Edit Destination NAT

You can configure security devices running ScreenOS 5.x and later, to translate the destination IP address. Enable Destination NAT and enter the destination IP address you want to translate to.

Other destination NAT options include:

- **Destination Port**—Your security devices can perform one-to-one destination NAT without changing the destination port numbers. However, you can configure the device to map the original destination port number in the segment header to another port number.
 - To enable destination port translation, select **Destination Port** and enter the port number you want to translate to.
 - To use the original destination port number, leave the default of None.
- **Upper IP Address**—Your device can also translate the destination IP address to a range of IP addresses. Select the **Upper IP Address** and enter the upper IP address. The device uses an address shifting mechanism to maintain the relationships among the original range of destination addresses after translating them to the new range of addresses.

Using the Device Manager, you can also implement NAT on any device interface in any zone except Untrust. For details, see NSM Online Help “Configuring Firewall/VPN Devices”.

For J Series devices, you can configure a NAT for a policy rule as one of the following:

- An interface
- A pool of a specific device interface
- A PoolSet defined under the “source NAT” setting for a device (collection of IP ranges)

You cannot configure NAT settings for SRX Series gateways using Policy Manager. NAT settings must be configured in the device for SRX Series gateways. However, if the device is managed in central management mode, you can right-click the device and select **Policy > View Pending Device Policy** to view all security policies that include NAT settings.

Enabling GTP for Firewall Rules

You can use a GTP object in a firewall rule to control how your security devices handle GPRS traffic. To add a GTP object, you must have already configured the object in Object Manager.

Configuring Traffic Shaping in a Security Policy

Traffic shaping enables you to control the amount of bandwidth that is available to the matching network traffic in a rule. You can also define a priority that defines how the security device handles the matching network traffic that exceeds the defined maximum bandwidth. For security devices running ScreenOS 5.3 and later, you can also manage the flow of traffic through the security device by limiting bandwidth at the point of ingress.

You can configure the following traffic shaping parameters:

- **Traffic Shaping Mode**—The traffic shaping mode is automatically determined by the security device, but you can set it to on or off.
- **Bandwidth**—You can control the amount of bandwidth that is available to the matching network traffic. When traffic shaping is enabled, you can configure the minimum, or guaranteed bandwidth allowed, by setting the number of kilobits per second (Kbps) using the Guaranteed Bandwidth field. This setting guarantees that this minimum amount of throughput is allowed to pass through the security device. In a similar manner, you can set the maximum bandwidth allowed using the Maximum Bandwidth field. For matching traffic that falls between the guaranteed and maximum settings, the security device passes traffic based on the priority setting.



NOTE: We recommend that you set the maximum bandwidth to greater than 10 Kbps. When the bandwidth is set to less than 10 Kbps, the security device might drop packets or the source address might attempt to resend the traffic repeatedly.

For security devices running ScreenOS 5.3 and later, you can also manage the flow of traffic through the security device by limiting bandwidth at the point of ingress. To configure the maximum amount of traffic allowed at the interface ingress, you need to first enable Use Policing Bandwidth, and then set the number of Kbps using the Policing Bandwidth field. This setting allows you to manage the maximum amount of traffic allowed to pass through the ingress interface.

- **Priority**—You can set a priority for each firewall rule in your security policy. Your security device passes permitted traffic according to the priority level specified in the matching rule. The higher the priority level of the rule, the faster the matching traffic for that rule passes. You can configure the mappings of eight priority levels to the first three bits in the DiffServ field or to the IP precedence field in the ToS byte in the IP packet header. By default, the highest priority (priority 0) on the security device maps to 111 in the IP precedence field. The lowest priority (priority 7) maps to 000 in the IP precedence field.
- **DSCP Class Selector**—NSM uses the Differentiated Services Code Point (DSCP) mechanism to set priority levels. Using DSCP, you can mark traffic at a position within a hierarchy of priority. You can map eight priority levels to the DiffServ system: Priority 0 is the highest priority, and priority 7 is the lowest priority. Each priority level maps to a specific set of bits in the DiffServ field or the IP precedence field in the ToS byte of the IP packet header. The class selector controls the number of bits affected in the DiffServ field. By default, the priority levels affect only the first three bits in the eight bit DiffServ field. The remaining bits are untouched, but can be altered by an upstream router, which might change the IP priority preference.

When the DSCP class selector is enabled, the class selector zeroes the remaining five bits in the DiffServ field, which prevents upstream routers from altering priority levels.

- DiffServ code point Group Marking—Enable this option by selecting the DiffServ code point Group Marking check box.
- DSCP Group—Click Add, Edit, or Delete to create or modify a DSCP group.

For information about changing the default mappings between priority levels and the DiffServ system, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.



NOTE: You can only apply traffic shaping to rules whose destination zone has a single interface bound to it. Security zones that contain sub interfaces or that contain more than one physical interface do not support traffic shaping.

For more information about Traffic Shaping, refer to the Concepts & Examples ScreenOS Reference Guide.

Enabling Logging and Counting for Firewall Rules

A good security policy generates enough log entries to fully document only the important security events on your network. However, if you need to keep a record of all log entries for archiving and accountability, you can design your rule to log every security event. For critical events, you might even want to be notified immediately by e-mail or set an alert to appear in the log entry.

Log entries appear in real-time in the Log Viewer and are also used in the Log Investigator for cross-tabulation of security events. Your goal is to fine-tune the notifications in your security policy to your individual security needs.

Configuring Logging and Alerts

To log an event for a rule, enable logging. Each time your security device matches network traffic to the rule, the device creates a traffic log entry that describes that event and NSM displays the traffic log entry in the Log Viewer. You can enable logging when a session is initialized, closed or both on a security device.

Depending on your security needs, you might want NSM to provide additional notification when a rule is matched, such as an alert in the log entry. An alert is a notification icon that appears in a log entry in the Log Viewer. When you enable alerts in your firewall rule and traffic matches that rule, the device generates a traffic log entry that includes an alert. Alerts can help you quickly identify specific network traffic, such as critical severity attacks.

You must enable logging before you can enable alerts.

Configuring Counting and Alarms

Counting and alarms work together to help you track the amount of traffic that is matching your firewall rule. Counting enables the device to count the number of bytes in network traffic that matches the firewall rule. Using this data, the device can then generate alarms that notify you when the matching network traffic falls outside your predefined byte range.

To set an alarm, enable counting and specify the minimum and maximum byte thresholds for matching network traffic. You can specify a predefined number of bytes per second, number of Kilobytes per minute, or both. Each time your security device detects network traffic that exceeds the alarm threshold in the rule, the device generates an alarm log entry for that describes that event and displays it in the Log Viewer.

You must enable counting before you can enable alarms. Although you can enable counting without also enabling alarms, NSM does not use the counting data except to trigger alarms. If you do not intend to use alarms, you should leave counting disabled. Additionally, because counting can impact performance during heavy traffic periods, you should enable counting and alarms only for firewall rules that detect important activity.

Configuring Log Actions

Use the Log Actions tab that appears when you select Log/Count in the Rule Options column to configure the following actions to occur when a log is generated from a specific rule:

- **Sending SNMP Trap**—Selecting this option directs the system to output logs to an SNMP server in SNMP format.
- **Sending Syslog Messages**—Selecting this option directs the system to output logs to a syslog server in syslog format.
- **Writing CSV files**—Selecting this option and specifying a filename directs the system to output logs using in CSV format.
- **Writing XML Files**—Selecting this option and specifying a filename directs the system to output logs using XML.
- **Sending Email**—Selecting this option directs the system to output logs to an e-mail address in SMTP format. You must specify the recipient e-mail address(es) that receives the exported log records.
- **Running Scripts**—Selecting this option directs the system to execute a script and report output status. You must specify the script that receives the exported log records (script must be located in the `/usr/netscreen/DevSvr/var/scripts/global` directory). In the event that the script fails, you can also configure the system to retry or skip running the script again.

You can configure log actions to occur for all rulebases, such as the IDP or Backdoor rulebases, that include logging options.

You can configure parameters for forwarding logs to SNMP, Syslog, Email, CSV and XML in the Action Parameters node of the Action Manager.

Miscellaneous

The following sections detail the Miscellaneous rule options.

Schedule

To control the time period that your security device applies the rule to your network traffic, you can define a schedule for the rule. If you define a schedule, the security device applies

the rule to your network traffic only during the time period specified in the schedule; if you do not specify a schedule, the rule is always applied to your network traffic.

In NSM, schedules are represented by schedule objects. Before you can define a schedule for a rule, you must create a schedule object that describes a time period. The schedule object defines the start time and date, end time and date, and frequency (recurring or one-time) of the time period.

You can use schedules to control the flow of network traffic at a time-sensitive level, and also enhance your network security.

To prevent employees from downloading large files during business hours, set the service object to FTP, the Action to deny, and configure traffic shaping to limit bandwidth. Using the Object Manager, create a schedule object called Business Day that describes the time period of 9:00 AM to 7:00 PM, M-F, recurring weekly. Right-click the schedule column in the rule and select the Business Day schedule object.

HA Session Backup

NetScreen-5XT and NetScreen-5GT security devices can disable active firewall rules that permit traffic if the session switches over to the modem link. This feature is ON by default.

ScreenOS 5.x and Later Options

For security devices running ScreenOS 5.x and later, you can configure additional rule options.

- **Application**—You can configure the security device to handle the service for the firewall rule as a known Layer 4 protocol service. If you are using application relocation (using a nonstandard port to handle an application service), enable this option to ensure that the security device correctly checks traffic.

ID

The rule ID is a number that uniquely identifies a rule within the rulebase and security policy. After you install a rule as part of a security policy on a security device, you can view that rule by logging in locally to the device with the WebUI or CLI where the rule appears as an individual policy. The individual policy on the device has the same ID as the rule in the management system, which helps you keep track of which rules are on which devices.

You can configure a rule ID for any zone-based firewall rule or VPN rule:

- For new rules, NSM automatically assigns a unique ID to that rule. You can change this ID, if desired, or leave the ID number.
- For rules that are already installed on a device, NSM has already created a unique ID for the rule. You can change this predefined ID if desired, to an ID number, or leave the ID set to “none”, which preserves the autogenerated ID number.



NOTE: When the ID is set to “none”, NSM uses a hashing algorithm on the source zone, destination zone, source address, destination address, and service fields for the rule to generate a unique ID.

- For VPN rules that are automatically created by VPN Manager, NSM creates a unique ID for each VPN rule. You can change this predefined ID, if desired, to a ID number, or leave the predefined ID set to “ none”, which preserves the autogenerated ID number.

When you copy and paste a rule within a rulebase, NSM automatically creates a new unique ID for the pasted rule.

You are not required to set a ID for a rule.



NOTE: The NSM GUI ID column now accepts alphabetic as well as numeric IDs.

Configuring Web Filtering for Firewall Rules

After you create a Web Filtering profile and you have enabled Web Filtering on your device, you need to bind it to your firewall rule. You need to select one of the following options:

- Web Filtering Through SurfControl SCFP/WebSense (Redirect)—With this option, the security devices sends the first HTTP request in a TCP connection to either a Websense server or a SurfControl server, enabling you to block or permit access to different web sites based on their URLs, domain names, and IP addresses.
- Web Filtering Through SurfControl CPA (Integrated)—With this option you permit or block access to a requested website by binding the default ns-profile or custom profile you created to a firewall rule.

When a profile is bound to the firewall rule, the security device matches the URL in the incoming HTTP request to the categories in the profile in the following sequence:

- Black List
- White List
- Custom URL Lists
- Predefined Web categories

If no custom profile is bound to the firewall rule, the security device uses the default profile ns-profile. If the security device does not find the category of the requested URL, then it performs the default action, to permit access to the URL (unless otherwise configured).

In this example, you will bind the predefined Web Filtering profile to a firewall rule.

1. Click **Policies** in the navigation tree. Select the device you want to bind to the Web filter profile.
2. In the Zone based Firewall Rules main display area, right-click under Rule Options. A pull-down menu appears.
3. Select **Web Filtering**.
4. In the Edit Web filter dialog box, click **Enable**.
5. Select **Web Filtering Through SurfControl CPA (Integrated)**. The Select SC-CPA Profile box appears.
6. Select the profile ns_profile to bind to the firewall rule.



NOTE: You can only bind one Web Filtering profile to a firewall rule.

7. Click **OK**.

Configuring Authentication for Firewall Rules

You can authenticate the identity of the user who is generating the network traffic. When you enable authentication in the rule, the user must authenticate future network traffic by supplying a user name and password in an initial, separate HTTP, FTP, or Telnet connection. If the user fails to authenticate using one of these services or provides incorrect credentials, the authentication requirement for the rule is not met and the network traffic is denied. (Typically, when you enable authentication, you also use the permit action.)



NOTE: You cannot enable authentication for a rule that includes the DNS/53 service object.

Configuring Authentication

Authentication enables you to control which RAS users can connect to the protected network and how they can connect. When you select an authentication server, you must also configure the users that authentication server authenticates.

Select the authentication mechanism:

- **No Authentication**—Use this option to enable the specified RAS users to connect without authentication.
- **Authentication**—Use for RAS users that use HTTP, FTP, or Telnet services to connect to the protected network. You can select an access profile as an authentication option from the Access Profile drop-down list box .
- **Web Authentication**—Use for RAS users using HTTP to connect to the protected network.
- **Infranet Authentication**—Use this option to enable specified RAS users to connect using a Juniper Networks Infranet Controller.

An unauthenticated user trying to access a UAC protected resource via HTTP, is usually redirected to a URL of an authenticating IC. The redirect URL is a global parameter specified per controller. On devices running ScreenOS 6.2 or later, you can additionally configure a redirect URL per policy, ensuring that traffic is efficiently handled.

- If you define a policy-based redirect URL, and enable redirect in the policy, unauthenticated HTTP traffic matching the policy is redirected to the policy-based redirect URL even if a global redirect URL is configured.
- If you do not define a policy-based redirect URL, and redirect is enabled in the policy, unauthenticated HTTP traffic matching the policy is redirected to the global redirect URL.

Configuring Users

RAS users are represented by user objects. Before you can authenticate a user in a firewall rule, you must create a user object that defines the user name, user password, and the authentication location (local or external). For Authentication and Web Authentication, configure the users:

- User—Select the User object that represents the user you want to authenticate.
- User Group—Select the User Group object that represents the users you want to authenticate.
- Group Expression—Select the Group Expression object.
- Allow Any—Use this option to authenticate any user or user group.

To authenticate RAS users with Authentication, you must include HTTP, FTP, or Telnet service objects in the Service column of the rule. You can include other services as well, or select any to specify all services. To make a connection to the destination IP address in the rule, the RAS user first initiates an HTTP, FTP, or Telnet connection to the destination address; the security device intercepts the request packet and responds with a login prompt for user credentials.

- If the destination address is a subnet, the remote user must authenticate for each IP address in that subnet.
- If the source address supports multiple remote user accounts (such as a Unix host running Telnet) OR is located behind a NAT device that uses a single IP address for all NAT assignments, only the first remote user from that source address must initiate and authenticate an HTTP, FTP, or Telnet connection. All subsequent remote users from that source address do not need to authenticate, and can pass matching network traffic to the destination address.

To authentication RAS users with Web Authentication, you must include HTTP service object in the Service column of the rule. To make a connection to the destination address in the rule, the RAS user first initiates an HTTP connection to the Web Authentication server. The security device responds with a login prompt for user credentials.

Configuring Antivirus for Firewall Rules

To configure Antivirus protection for a firewall rule:

- **None**—No Antivirus protection enabled.
- **Use External AV Server**—Uses an external antivirus scanner. Select an external policy object that defines an external scanner.
- **Use Scan Manager**—Scan Manager is an embedded scanning engine. To use Scan Manager, the security device you install the policy on must be a NetScreen-5GT or NetScreen-Hardware Security Client device running ScreenOS 5.0 - 5.2. If you install a policy that uses Scan Manager on a different device, the device executes and processes traffic according to the rule, but does not detect viruses using the embedded scanning engine.
- **Use Scan Manager with Profile**—Scan Manager is an embedded scanning engine. This setting tells the device to use the global profile specified. This setting only works for devices running ScreenOS 5.3.
- **Use ICAP Profile**—ICAP is a method of redirecting traffic to an ICAP-capable server running AV software. This feature is available on devices running ScreenOS 5.4 and higher.

Configuring a DI Profile/Enable IDP for Firewall Rules

Use the DI Profile/Enable IDP rule options to configure Deep Inspection (DI) or Intrusion Detection and Prevention (IDP) functionality within the rule.

This function applies to firewall and ISG devices only. Standalone IDP devices do not use firewall rules. DI and IDP are mutually exclusive. When you install the IDP license key on a security device, DI is automatically disabled.

Configuring DI Profile for a Rule

Security devices running ScreenOS 5.x and later, include Deep Inspection attack protection that can detect malicious network traffic at the application level. To configure attack protection, select a DI Profile object in your firewall rule to detect intrusion attempts within permitted traffic.

Attacks are specific patterns of malicious activity within a network connection, and an attack object uses selected sections of the attack pattern to detect the attack itself. NSM contains a database of predefined attack objects that detect known and unknown attacks against your network. You can use these predefined attack objects (and your own custom attack objects) to create a DI Profile object, which you then use in a firewall rule. When configuring a DI Profile, you can also defined the action that the device performs against those attacks when detected in permitted traffic.

You can configure one DI Profile for each rule. When the device detects a match between the permitted network traffic and an attack object within the selected DI Profile, the device generates an attack log entry.

To use a DI Profile:

- The firewall action must be permit. You cannot detect attacks in traffic that the device denies or rejects.

- The security device you install the policy on must be running ScreenOS 5.0 and later. If you install a policy that contains a DI Profile on an earlier ScreenOS device, the device executes and processes traffic according to the rule, but does not detect application-level attacks.

Configuring IDP for a Firewall Rule

When configuring a rule for an IDP-capable device, such as the ISG2000 or ISG1000 security gateway running ScreenOS 5.0 IDP1, you can enable IDP and select an IDP mode in the DI Profile/Enable IDP rule options. Enabling IDP directs the security device to pass all traffic permitted by the firewall rule to the IDP rulebase.

DI and IDP are mutually exclusive. When you install the IDP license key on a security device, DI is automatically disabled.

When configuring the firewall rule, consider the following:

- Traffic that is denied by a firewall rule cannot be passed to IDP rules. To enable IDP in a firewall rule, the action must be permit.
- For firewall rules that pass traffic to the IDP rulebases, the Install On column must include IDP-capable devices only.

To forward traffic to the IDP rulebases, enable IDP and select one of the following modes:

- In inline mode, IDP is directly in the path of traffic on your network and can detect and block attacks. For example, you can deploy the ISG2000 or ISG1000 with integrated Firewall/VPN/IDP capabilities between the Internet and an enterprise LAN, WAN, or special zones such as DMZ.
- In inline tap mode, IDP can detect attacks and provide notification. IDP receives a copy of a packet while the original packet is forwarded on the network. IDP examines the copy of the packet and flags any potential problems. IDP's inspection of packets does not affect the forwarding of the packet on the network.

You must deploy the ISG2000 or ISG1000 device inline. You cannot connect a device that is in inline tap mode to an external TAP or SPAN port on a switch.

Selecting either mode enables IDP for the firewall rule, and configures the security device to forward all permitted traffic to the IDP rulebases for further processing.

Limiting Sessions per Policy from Source IPs

With the session-limit option, you can restrict sessions from a particular Source IP address to all your devices running ScreenOS 6.1 and later. In NSM, you can set the following options from the Session Limit tab in the Configure Options window of the device.

- Session limit per src-ip on policy
- Session count
- Alarm without drop packet

When the sessions reach the threshold limit, the system drops all subsequent sessions. If you enable the "alarm without drop packet" option, the packet is not dropped, but an

alarm message is raised. If you do not set a source IP, the device lists the session counts of all the source IP addresses in the policy.

In cross-vsystraffic, since there is one policy per vsys to permit traffic, each cross-vsyst session is permitted by two policies. However, the session limit policy is only for the ingress vsys. You must configure the session limit in the ingress vsys policy to limit the session count.

In a synchronized NSRP setup, the session limit policy also counts sessions in the slave device, which does not impose any limit. When the slave becomes the master, a new session is created only if the existing session count does not exceed the threshold. If the threshold is exceeded, the packet is dropped. A new session can be created only when the session counts drop below the threshold when existing sessions are aged out.

Configuring the Session Close Notification Rule

An idle TCP connection remains established until terminated by either the client or the server. If, for any reason, the client or an intermediate device shuts down, the server continues to wait on the connection. As an intermediate security device, a device running ScreenOS maintains a session for each TCP connection until it times out. Traffic can resume if a client sends an RST (reset) packet, but the client needs to be informed of the situation in order to do so. If the TCP keep-alive option is activated on the server, it can be used to query the status of the connection.

NSM offers the option of configuring the SSG Series Secure Services Gateways, ISG Series Integrated Security Gateways, and the NetScreen Series Security Systems running ScreenOS 6.3 and later to send a notification to both the client and the server when a TCP session is closed. By default, this option is disabled. Before you can enable the Session Close Notification feature on NSM for a device, you must first set the following options:

- a. From **Device > Advanced > Packet flow >**:
 - Disable **Skip TCP sequence number check**.
 - Enable one or both of these options:
 - **Check TCP SYN bit before create/refresh session after TCP handshake**
 - **Check TCP SYN bit before Create session**
 - Set the number of seconds in the option **Notify threshold**.
- b. From **Device > Network > Edit the From / To Zone**, enable TCP/RST.

Configuring the Session Close Notification option:

1. Select **Policy Manager > Security Policy > Policy on device > Rule Options > Session Close Notification**. A **Session Close Notification** window opens.
2. Check the option – **Notify both ends if TCP session isn't normally terminated**.
3. Click **OK**.

configure the Session Close Notification option by selecting **Policy Manager > Security Policy > Policy on device > Rule Options > Configure All Options Session Close Notification**.

Comments for Firewall Rules

The Comments column of a rule contains the rule title, which is also the ScreenOS policy name (the name of the policy when viewing the device configuration using the WebUI).

You can also enter comments in the Comment Field, if desired.

Configuring Multicast Rules

A multicast rule is a statement that defines a specific type of multicast control traffic. When multicast control traffic passes through a security device, the device attempts to match that traffic against its list of rules. If a rule is matched, the device permits the traffic to pass through.

On security devices, you secure multicast control traffic using access lists. First, you create an access list, which defines one of the following:

- The multicast groups a host can join.
- The sources from which traffic can be received.

After creating an access list, you reference these access lists in a multicast rule in the security policy for the device.

Configuring Source and Destination Zones

In the Multicast rulebase, you create rules to enable multicast control traffic to flow between zones. You must create zones on your security device before you can create a rule for that device. In a single rule:

- You must select a single zone for the source zone and a single zone for the destination zone. These zones must be available on the security devices on which you install the policy.
- You can also select multiple zone exceptions for both source and destination zones. A zone exception includes a specific zone and the device that contains that zone.

Configuring Source and Destination Groups

When you create a multicast rule, specify the multicast groups for which you want to permit multicast traffic using one of the following methods:

- Specify a multicast group IP address, and optionally, the multicast group address on the outgoing interface
- Specify the access list that identifies the permitted multicast groups
- Select “any” to accept traffic for all multicast groups

Configuring Rule Options

Rule options enable you to specify the type of multicast control traffic to which this rule applies and whether the rule is bidirectional.

A rule can apply to either IGMP messages or PIM-SM messages:

- When running IGMP proxy on the security device, configure a rule that permits IGMP messages to flow between zones.
- When running PIM-SM on the security device, configure a rule that permits PIM-SM messages.

In this example, you define a multicast rule that permits IGMP messages from the Trust zone to the Untrust zone. You specify the original multicast group address object and a different destination multicast group object.

1. In the main navigation tree, select **Object Manager > Address Objects**.
2. In the main display area, click the Add icon and select **Multicast Group**. In the New Multicast Group dialog box, configure the following then click **OK**:
 - For Name, enter **mcast1**.
 - For Color, select **green**.
 - For IP Address, enter **232.1.1.1**.
 - For Netmask, enter **16**.
3. In the main display area, click the Add icon and select **Multicast Group**. In the New Multicast Group dialog box, configure the following then click **OK**:
 - For Name, enter **mcast2**.
 - For Color, select **red**.
 - For IP Address, enter **232.1.1.2**.
 - For Netmask, enter **16**.



NOTE: NSM validation prevents you from setting a 32 bit netmask in multicast.

- In the main navigation tree, select **Policies**, then create a new multicast rule in the Multicast rulebase of a new or existing security policy.
- Right-click in the Source Group column and select **Configure Source/Destination**.

Configuring Antivirus Rules

Antivirus settings are stored in a profile.

To assign an antivirus profile to a policy, do the following:

1. Double-click the Rule Options cell in a rule.
2. In the Configure Options dialog, click the Antivirus tab.
3. Select an Antivirus option:
 - **None**—Turns off antivirus scanning for that rule.
 - **Use External AV Server**—Indicates that you want to use an External AV Server. You must select the external AV server you wish to use.
 - **Use Scan Manager**—Tells the device to use the settings on the device, instead of a profile. Necessary for ScreenOS 5.0–5.2.
 - **Use Scan Manager with Profile**—Tells the device to use the indicated antivirus profile. Necessary for ScreenOS 5.3 and later.
 - **Use ICAP Profile**—Tells the device to use the indicted ICAP AV profile. Available with ScreenOS 5.4 and later.

Configuring Antispam Rules

Antispam settings are stored in profiles. Initially, NSM will have only one antispam profile available: ns-profile.

To assign an antispam profile to a policy, do the following:

1. Double-click the Rule Options cell in a rule.
2. In the Configure Options dialog, click the Antispam tab.
3. Check the **Enable Antispam profile** check box.
4. Select **ns-profile** in the Profile Name pull-down menu.
5. Click **OK**.

Configuring IDP Rules

The IDP rulebase protects your network from attacks by using attack objects to identify malicious activity and take action. Creating an IDP rule involves the following steps:

- [“Defining Match for Firewall Rules” on page 492](#) (does not apply to rulebases for standalone IDP Sensors) —The type of network traffic you want IDP to monitor for attacks, such as source/destination zones, source/destination address objects, and the application layer protocols (services) supported by the destination address object. You can also negate zones, address objects, or services.

Standalone IDP Sensors do not use firewall rules.

- [“Configuring Terminal IDP Rules” on page 516](#)—By default, rules in the IDP rulebase are non-terminal, meaning that IDP examines all rules in the rulebase and all matches are executed. You can specify that a rule is terminal; if IDP encounters a match for the source, destination, and service specified in a terminal rule, it does not examine any

subsequent rules for that connection. Note that the traffic does not need to match the attacks specified in the terminal rule. Terminal rules should appear near the top of the rulebase, before other rules that would match the same traffic. Use caution when specifying terminal rules.

- [“Configuring Attack Objects in IDP Rules” on page 519](#)—The attacks you want IDP to match in the monitored network traffic. Each attack is defined as an attack object, which represents a known pattern of attack. Whenever this known pattern of attack is encountered in the monitored network traffic, the attack object is matched. You can add attack objects by category, operating system, severity, or individually.
- [Configuring Actions](#)—The action you want IDP to take when the monitored traffic matches the rule’s attack objects. You can specify the action you want the device to perform against the current connection (see [“Defining Actions For IDP Rules” on page 517](#)) and future connections from the same source IP address (see [“Choosing an IP Action” on page 522](#)).
- [“Configuring Notification in IDP Rules” on page 523](#)—Disable or enable logging for the IDP rule.

The following sections detail each step.

Defining Match For IDP Rules

When creating your IDP rules, you must specify the type of network traffic that you want IDP to monitor for attacks. These characteristics include the network components that originate and receive the traffic, and the firewall zones the traffic passes through.

You must specify the From Zone, Source, User Role, To Zone, Destination, and Service in their respective Match columns for all rules in the IDP rulebase. The Terminate Match selection allows you to designate a rule as terminal; if IDP encounters a match for the other Match columns in a terminal rule, no other rules in the rulebase are examined. The matching traffic does not need to match the attacks specified in a terminal rule. (For more information on terminal rules, see [“Configuring Terminal IDP Rules” on page 516](#).)

The following sections detail the Match columns of an IDP rule.

Configuring Source and Destination Zones for IDP Rules (Does not apply to Standalone IDP Sensor rulebases)

You can select multiple zones for the source and destination, however these zones must be available on the security devices on which you will install the policy. You can specify “any” for the source or destination zones to monitor network traffic originating or destined for any zone.

For standalone IDP rulebases, the zones are always set to “any.”



NOTE: You can create custom zones for some security devices. The list of zones from which you can select source and destination zones includes the predefined and custom zones that have been configured for all devices managed by NSM. Therefore, you should only select zones that are applicable for the device on which you will install the security policy.

Configuring Source and Destination Address Objects for IDP Rules

In the NSM system, address objects are used to represent components on your network: hosts, networks, servers, etc. Typically, a server or other device on your network is the destination IP for incoming attacks, and can sometimes be the source IP for interactive attacks (see [“Configuring Backdoor Rules” on page 538](#) for more information on interactive attacks). You can specify “any” to monitor network traffic originating from any IPv4 address and “AnyIPv6” to monitor network traffic originating from any IPv6 address. You can also “negate” the address objects listed in the Source or Destination column to specify all sources or destinations except the excluded objects.

You can create address objects either before you create an IDP rule or while creating or editing an IDP rule. To select or configure an address object, right-click either the Source or Destination column of a rule and select **Select Address**. In the **Select Source Addresses** dialog box, you can either select an already-created address object or click the **Add** icon to create a new host, network, or group object.

To detect incoming attacks that target your internal network, set the **From Zone** to **Untrust**, and the **Source IP** to any IP. Then, set the **To Zone** to **dmz** and **trust**. Next, select the address object that represents the host or server you want to protect from attacks as the **Destination IP**.

To detect attacks between two network, select multiple address objects for the **Source** and **Destination**.

The more specific you are in defining the source and destination of an attack, the more you reduce false positives.

Configuring User Roles for IDP Rules

You can use role-based IDP policy to define roles and related access privileges, and apply an application policy to them that is effective regardless of where the user logs in. Role-based access control facilitates a dynamic network and access to partners. This feature is supported on the ISG1000 and ISG2000 gateways with SM devices running ScreenOS 6.3 and later.

To support role-based IDP policy, you must select both **Infranet Auth** and **IDP Enabled** in the **Firewall Rule Options**. When it receives a packet, the firewall verifies the role name of the user against the list of user roles and user role groups provided before forwarding the packet. You can configure either IP-based rules or role-based rules in an IDP policy but not both. Role-based rules have higher precedence than IP-based rules. Therefore, if roles have been specified for a session, the firewall first tries to match role-based rules and then tries to match IP-based rules. If roles are not configured for a session, the firewall searches for IP-based rules.

You can configure this feature by selecting **Policy Manager > Policies**. Select a device policy and add an IDP rulebase. Right-click on the **User Role** column. You can then **Select**, **Filter** or **Edit** user roles. If you select user roles, the **Select User Roles** dialog box opens. Select the device from the drop-down list in the **Device** field. Click the add icon (+) in the **Selected User Roles** to add either **New User Roles** or **New User Role Groups**. You can enter a user role in the **New User Define** box and click **OK** to create a new user role. The **Select**

User Roles dialog box allows you to view all the created user roles and add or remove them from the IDP policy. Similarly, you can create user role groups in the **New User Defined User Role Group** dialog box, view them, and add or remove them from the policy.

When you right-click on the **User Roles** column, you can also use the **Filter** and **Edit** options provided. With the **Filter** option, you can choose to apply a filter (true or false, negate, or ignore objects in group) to the user role values. The **Edit** option allows you to cut, copy, or paste the user role name in the column.

After making your changes, save the policy, and then update the device. Ensure that the device reflects the correct user role information.

The role-based access control feature has the following limitations:

- The role names in IDP policy must match those of the Infranet Controller (IC).
- Username-based IDP policy is not supported. The firewall must map either a source IP or the username to a user role before it can forward a packet.
- While the firewall supports 200 roles for one user, the IDP policy supports only 100 roles for each user.
- JUMBO FRAME or IPv6 mode is not supported.
- SYN Proxy or First UDP packet with fragment is not supported.
- Vsys is not supported.

Configuring Services for IDP Rules

Services are application layer protocols that define how data is structured as it travels across the network. Because the services you support on your network are the same services that attackers must use to attack your network, you can specify which services are supported by the destination IP to make your rule more efficient.



NOTE: All services rely on a transport layer protocol to transmit data. IDP includes services that use TCP, UDP, RPC, and ICMP transport layer protocols.

Service objects represent the services running on your network. NSM includes predefined service objects that are based on industry-standard services. You use these service objects in rules to specify the service an attack uses to access your network. You can also create custom service objects to represent protocols that are not included in the predefined services.

In the Service column you select the service of the traffic you want IDP to match:

- Select **Default** to accept the service specified by the attack object you select in the Attacks column. When you select an attack object in the Attack column, the service associated with that attack object becomes the default service for the rule. To see the exact service, view the attack object details.
- Select **Any** to set any service.

- Select **Service** to choose specific services from the list of defined service objects.

You want to protect your FTP server from FTP attacks. Set the service to Default, and add an attack object that detects FTP buffer overflow attempts. The Service column in the rule still displays “Default”, but the rule actually uses the default service of TCP-FTP, which is specified in the attack object.

Your mail server supports POP3 and SMTP connections but not IMAP. Set POP3 and SMTP service objects as services that can be used to attack that server. Because IMAP is not supported, you do not need to add the IMAP service object.

If you are supporting services on nonstandard ports, you should choose a service other than default.

You use a nonstandard port (8080) for your HTTP services. Use the Object Manager to create a custom service object on port 8080.

Add this service object to your rule, then add several HTTP attack objects, which have a default service of TCP/80. IDP uses the specified service, HTTP-8080, instead of the default, and looks for matches to the HTTP attacks in TCP traffic on port 8080.

You can create your own service objects to use in rules, such as service objects for protocols that use nonstandard ports. However, you cannot match attack objects to protocols they do not use.

Configuring Terminal IDP Rules

The normal IDP rule-matching algorithm starts from the top of the rulebase and checks traffic against *all* rules in the rulebase that match the source, destination, and service. A terminal rule is an exception to this normal rule-matching algorithm. When a match is discovered in a terminal rule for the source, destination, and service, IDP does not continue to check subsequent rules for the same source, destination, and service. It does not matter whether or not the traffic matches the attack objects in the matching rule.

You can use a terminal rule for the following purposes:

- To set different actions for different attacks for the same Source and Destination. This is illustrated by rules 3 and 6 in the “Setting Terminal Rules” example below.
- To disregard traffic that originates from a known trusted Source. Typically the action is None for this type of terminal rule. This is illustrated by rule 1 in the “Setting Terminal Rules” example below.
- To disregard traffic that is sent to a server that is only vulnerable to a specific set of attacks. Typically, the action is Drop Connection for this type of terminal rule.

Use caution when defining terminal rules. You can inadvertently leave your network open to attacks by creating an inappropriate terminal rule. Remember that traffic matching the source, destination, and service of a terminal rule is not compared to subsequent rules, even if the traffic does not match an attack object in the terminal rule. Use a terminal rule only when you want to examine a certain type of traffic for one specific set of attack objects and no others. Be particularly careful about terminal rules using “any” for both the source and destination.

Terminal rules should appear near the top of the rulebase, before other rules that would match the same traffic. You set a rule as terminal by selecting the box in the Terminate Match column of the Security Policy window when the rule is created or modified.



NOTE: In many cases, you can use an exempt rule instead of a terminal rule. You might find it easier and more straightforward to configure an exempt rule than a terminal rule. See [“Configuring Exempt Rules” on page 535](#).

In the example IDP rulebase shown below, rules 1, 3 and 5 are configured as terminal rules:

- Rule 1 terminates the match algorithm if the source IP of the traffic originates from the Security Network, a known trusted network. If this rule is matched, IDP disregards traffic from the Security Network and does not continue monitoring the session for malicious data.
- Rules 3 and 6 set different actions for different attacks when the destination IP is the Corporate or Europe E-mail server. Rule 3 terminates the match algorithm when the attack is an e-mail that uses the SMTP context Confidential. Rule 6 closes the server when the attack is an SMTP attack.
- Rule 5 terminates the match algorithm when the source is the Internal Network and the attack is a Critical, High, or Medium Trojan Backdoor. The rule ensures that IDP closes both the client and server and does not continue to match the connection.

Defining Actions For IDP Rules

You can define actions for the security device to perform against attacks that match rules in your security policy. For each attack that matches a rule, you can choose to either take action on the packet containing the attack (permit or drop packet) or take action on the connection or session (permit, ignore, drop or close connection). Refer [Table 45 on page 518](#) for details.

Remember, that the device can drop the packet containing the attack only when IDP is enabled in the inline mode.

When IDP is enabled in the inline tap mode on ISG-IDP devices, and the action defined is *drop packet* or *drop connection*, IDP causes the firewall to drop the session upon detection of an attack. However, it cannot prevent the attack packet from reaching its destination because in the inline tap mode, the IDP only receives a copy of the packet while the original packet is sent to its destination.

When standalone IDP sensors are deployed in the inline tap or sniffer mode, IDP cannot perform a drop action and there is no disruption to the session carrying attack traffic.

[Table 45 on page 518](#) lists actions for IDP rules:

Table 45: IDP Rule Actions

Action	Description
None	IDP inspects for attacks but takes no action against the connection if an attack is found. If a rule that contains an action <i>None</i> is matched, the corresponding log record displays <i>accept</i> in the action column of the Log Viewer.
Ignore	IDP completely ignores the session if the rule does not specify an attack. If an attack is specified in the rule, IDP inspects the session and generates a log for the first attack detected. Subsequently, IDP ignores the rest of that session and neither inspects the session for attacks nor generates attack logs. Use with caution.
Drop Packet	<p>IDP drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source IP address.</p> <p>Depending on the protocol in use and its mode, IDP behaves differently when you define this rule.</p> <ul style="list-style-type: none"> • If using UDP in the inline mode, the IDP drops the packet whereas it dismisses the action if functioning in the inline tap mode. • If using TCP, in the inline mode, the IDP drops the connection. In the inline tap mode, though the connection is dropped, the attack packet might still have got through.
Drop Connection	<p>IDP drops the connection without sending a RST packet to the sender, preventing the traffic from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.</p> <p>Depending on the protocol in use and its mode, IDP behaves differently when you define this rule.</p> <ul style="list-style-type: none"> • If using UDP in the inline mode, the IDP drops the session. In the inline tap mode, the session is dropped but the attack packet would have been let through. • If using TCP in the inline mode, the IDP drops the connection. In the inline tap mode, the IDP drops the connection but the attack packet might have got through.
Close Client	IDP closes the connection to the client, but not to the server.
Close Server	IDP closes the connection to the server, but not to the client.
Close Client and Server	IDP closes the connection and sends a RST packet to both the client and the server. If IDP is operating in inline tap mode, IDP sends a RST packet to both the client and server but does NOT close the connection.
Diffserv Marking	IDP assigns the service differentiation value indicated to the packet, then passes it on normally. The value is set in the dialog that appears when you select this action in the rulebase.

Table 45: IDP Rule Actions (continued)

Action	Description
Recommended	<p>IDP takes the action recommended by Juniper Networks. With IDP 4.1 and later, attack objects have a recommended action associated with them. If a packet triggers more than one attack object, IDP applies the most secure of the recommended actions. Available with IDP 4.1 and later.</p> <p>This setting has no meaning for IDP 4.0 or earlier. Rules with this setting will not be loaded onto devices running earlier versions of IDP.</p>

Configuring Attack Objects in IDP Rules

Attack objects represent specific patterns of malicious activity within a connection, and are a method for detecting attacks. Each attack object detects a known or unknown attack that can be used to compromise your network. .

To add attack objects to a rule, right-click the Attacks column of the rule and select **Select Attacks**. In the Add Attacks dialog box, you can add attacks using one or both of the following options:

- **Attack List**—Select this option to add individual attack objects from an alphabetically list of all predefined and custom attack objects. Attack objects are listed alphabetically by name of attack.

Selecting individual attacks is a good option if you know the exact name of the attack you want to add to a rule. To locate a specific word or string in the attack object name, use the integrated search function in NSM.

- **Attack Groups**—Select this option to add attack object groups from three predefined dynamic attack groups (Category, OS, Severity); if you have created a custom dynamic group, that group is also listed.

Selecting attack groups is a good option when you are unsure of the exact attack you want to add to a rule, but you know the type of attack protection you want the security device to provide. Within the Attack Groups, you can:

- Add all attack objects (select All Attacks). Consider carefully before selecting this option; using all attack objects in a rule can severely impact performance on the security device.
- Add one or more attack groups (hold down CTL to select multiple groups). Predefined dynamic groups might contain subgroups as well.
- Add individual attack objects (hold down CTL to select multiple objects)

The following sections detail each predefined dynamic attack group.

Adding IDP Attack Object Groups by Category

The Category group includes attack objects organized by services. Services are application layer protocols that define how data is structured as it travels across the network. A

protocol is a specification that indicates how communication between two entities (applications, servers, Ethernet cards, etc.) occurs.

When attacking a system, attackers use the protocol of a supported service to communicate their malicious activity to the server. However, attackers can only use protocols that are supported by the system they are attacking. You can add a category group to the Attacks column in your rule; however, you need to select only the categories that are used by the address objects you are protecting with the rule.

For example, if you rely extensively on FTP and HTTP for file transfers to and from your Web servers, choose the FTP and HTTP category groups to carefully monitor all traffic that uses these services.



NOTE: As of Release 2007.3, a few of the entries in the IDP attack group table, starting with the Response category, are removed to enhance the performance of IDP devices. See the latest NSM Release Notes for information on the Response category removed from the IDP attack group table.

Adding IDP Attack Objects by Operating System

The Operating System group includes attack objects for several predefined operating systems to help you choose the attack objects that are the most dangerous to specific components on your network. You can choose BSD, Linux, Solaris, or Windows.

Adding IDP Attack Objects by Severity

The Severity group includes five attack object groups organized by severity level. You can select one or more groups to include in your rule. To protect critical address objects or “popular” attacker targets, such as your mail server, use multiple severity levels to ensure maximum protection.

We recommend using the following actions and notification settings listed in [Table 46 on page 520](#) when using severity-based dynamic attack groups in a rule:

Table 46: Severity Levels, Recommended Actions and Notifications

Severity	Cause	Recommended Action	Notification
Critical	Attacks attempt to evade an IDS, crash a machine, or gain system-level privileges.	Drop Packet	Logging Alert
Major	Attacks attempt to crash a service, perform a denial-of-service, install or use a trojan (1c), or gain user-level access to a host.	Drop Packet Drop Connection	Logging Alert
Minor	Attacks attempt to obtain critical information through directory traversal or information leaks.	(no recommended action)	Logging

Table 46: Severity Levels, Recommended Actions and Notifications (continued)

Severity	Cause	Recommended Action	Notification
Warning	Attacks attempt to obtain noncritical information or scan the network with a scanning tool. They can also be obsolete attacks or anomalous (but probably harmless) traffic.	(no recommended action)	Logging
Info	Attacks are normal, harmless traffic containing URLs, DNS lookup failures, and SNMP public community strings. You can use informational attack objects to obtain information about your network.	(no recommended action)	(no recommended notification)

You configure actions in the Action column of the rule; see [“Defining Actions For IDP Rules” on page 517](#). You configure notification settings in the Notification column of the rule; see [“Configuring Notification in IDP Rules” on page 523](#).

Adding Custom Dynamic Attack Groups

You can add previously created custom dynamic attack groups to a rule.

Additionally, after you have added the custom group to a rule, you can edit the settings for the dynamic group by double-clicking the group icon in the rule.

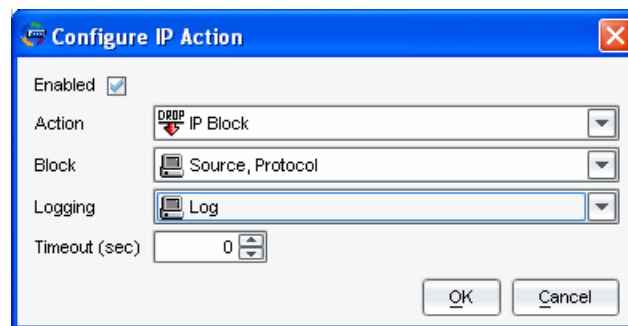
Configuring IP Actions in IDP Rules

This column only appears when you view the security policy in Expanded Mode. To change the security policy view from Compact Mode to Expanded Mode, from the menu bar, select **View > Expanded Mode**.

If the current network traffic matches a rule, the security device can perform an IP action against future network traffic that uses the same IP address. IP actions are similar to other actions; they direct the device to drop or close the connection. However, because you now also have the attacker’s IP address, you can choose to block the attacker for a specified amount of time. If attackers cannot immediately regain a connection to your network, they might try to attack easier targets.

Use IP actions in conjunction with actions and logging to secure your network. In a rule, first configure an action to detect and prevent current malicious connections from reaching your address objects. Then right-click in the IP Action column of the rule and select **Configure**. The Configure IP Action dialog box appears, as shown in [Figure 93 on page 522](#). Enable and configure an IP action to prevent future malicious connections from the attacker’s IP address.

Figure 93: Configure IP Action



Choosing an IP Action

For each IP action option, an IP action is generated by NSM. The IP action instructs the security device to perform the specified task. Select from the following options:

- IDP Notify—The security device does not take any action against future traffic, but logs the event. This is the default.
- IDP Drop—The security device drops the matching connection and blocks future connections that match the criteria set in the Block list.
- IDP Close—The security device closes future connections that match the criteria in the Block list.

Choosing a Block Option

Each block option follows the criteria you set in the Actions box. Block options can be based on the following matches of the attack traffic:

- Source, Destination, Destination Port and Protocol—The security device blocks future traffic based on the source, destination, destination port, and protocol of the attack traffic. This is the default.
- Source—The security device blocks future traffic based on the source of the attack traffic.
- Destination—The security device blocks future traffic based on the destination of the attack traffic.
- From Zone, Destination, Destination Port and Protocol—The security device blocks future traffic based on the source zone, destination, destination port, and protocol of the attack traffic.
- From Zone—The security device blocks future traffic based on the source zone of the attack traffic.

Setting Logging Options

When the security device detects attack traffic that matches a rule and an IP action is triggered, the device can log information about the IP action that was taken or create an alert in the Log Viewer. By default, there are no logging options set.

Setting Timeout Options

You can set the number of seconds that you want the IP action to remain in effect after a traffic match. For permanent IP actions, leave the timeout at 0 (this is the default).

Configuring Notification in IDP Rules

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, you can also set an alert flag to appear in the log record.

To log an attack for a rule, right-click the Notification column of the rule and select **Configure**. The Configure Notification dialog box appears.

The first time you design a security policy, you might be tempted to log all attacks and let the policy run indefinitely. Don't do this! Some attack objects are informational only, and others can generate false positives and redundant logs. If you become overloaded with data, you can miss something important. Remember that security policies that generate too many log records are hazardous to the security of your network, as you might discover an attack too late or miss a security breach entirely due to sifting through hundreds of log records. Excessive logging can also affect throughput, performance, and available disk space. A good security policy generates enough logs to fully document only the important security events on your network.



NOTE: J Series and SRX Series devices do not send packet data to NSM. If your policy rules attempt to do so, then NSM does not log the data.

- **Setting Logging**—In the Configure Notification dialog box, select **Logging** and then click **OK**. Each time the rule is matched, the NSM system creates a log record that appears in the Log Viewer.
- **Setting an Alert**—In the Configure Notification dialog box, select **Alert** and then click **OK**. If Alert is selected and the rule is matched, the security device places an alert flag in the Alert column of the Log Viewer for the matching log record.
- **Logging Packets**—You can record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack and its purpose, whether or not the attack was successful, and any possible damage to your network.



NOTE: To improve performance, log only the packets after the attack.

If multiple rules with packet capture enabled match the same attack, the security device captures the maximum specified number of packets. For example, you configure Rule 1 to capture 10 packets before and after the attack, and Rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.



NOTE: Packet captures are restricted to 256 packets before and after the attack.

Setting VLAN Tags for IDP Rules

You can choose to apply rules to traffic on certain VLANs only. Normally, for a rule to take effect, it must match the packet source, destination, service, and attack objects. If the VLAN cell is populated with a value other than any, then the rule will also consider the packet's VLAN tag when determining a match.

The IDP, Exempt, Backdoor, SYN Protector, Traffic Anomalies, and Network Honeypot rulebases support VLAN matching. VLAN matching is only supported in Transparent and Sniffer modes.



NOTE: VLAN matching is supported in IDP 4.1 and later. Rules with a VLAN Tag field set to anything other than any are removed from the rulebase before NSM sends the security policy to an IDP device that does not support VLAN tags.

VLAN tag matching can be set to any, none, a particular VLAN tag value, or a range of VLAN tag values. Use VLAN objects to create individual VLAN tags or ranges of VLAN tags. You can assign more than one VLAN object to a rule. To assign a VLAN object to a rule, or to set the VLAN Tag value to none, right-click in the VLAN Tag cell of the rule.

VLAN matching works as follows:

- Any: Matches traffic with any or no VLAN tag (default)
- Single tag: Matches traffic with that specific tag only
- Range of tags: Matches traffic with any tag in that range
- None: Matches only traffic that has no VLAN tag

Setting Severity for IDP Rules

(This column only appears when you view the security policy in Expanded Mode. To change the security policy view from Compact Mode to Expanded Mode, from the menu bar, select **View > Expanded Mode**.)

You can override the inherent attack severity on a per-rule basis within the IDP rulebase. You can set the severity to either Default, Info, Warning, Minor, Major, or Critical.

To change the severity for a rule, right-click the Severity column of the rule and select a severity.

Setting Target Devices for IDP Rules

For each rule in the IDP rulebase, you can select which devices the rule applies to. When you install the security policy that the rule belongs to, the rule becomes active only on the devices you selected in the Install On column of the rulebase.



NOTE: NSM supports IDP on ISG gateways running ScreenOS 5.0.0-IDP1 and later, standalone IDP appliances running IDP 5.0 and later, J Series routers, SRX Series gateways, and MX Series routers .

Entering Comments for IDP Rules

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select **Edit Comments**. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

You can deploy an ISG2000 or ISG1000 gateway as a standalone IDP security system protecting critical segments of your private network. For example, you might already have a security device actively screening traffic between the Internet and your private network (some device can optionally use Deep Inspection to inspect this traffic), but you still need to protect internal systems, such as mail servers, from attacks that might originate from user machines in an otherwise trusted network. In this case, you need a security system that provides IDP instead of firewall functions.

Standalone IDP Sensors function in this mode by default and do not have to be specifically configured for it.

In this example, you are deploying an ISG2000 device as a standalone IDP security system between the Trust zone and the custom "Data_Center" zone in your network. Your company's file, mail, and database servers reside in the Data_Center zone. While you want to allow users in the Trust zone to be able to access the servers in the Data_Center zone, you also need to protect the servers from attacks that inadvertently might have been introduced into a user machine in the Trust zone. You create a firewall rule from the Trust to the Data_Center zone that allows traffic from any source to any destination for any service, then enable IDP in the Rule Options column.

You would then add and configure IDP rulebases for the security policy to detect possible attacks against servers in the Data_Center zone.

Configuring multiple IDP policies for an MX Series Router

You can configure multiple IDP policies for an MX Series device by associating existing IDP rules in the security policy assigned to the device, to multiple IDP policies. IDP services on MX series routers allow administrators to provide security services to service provider subscribers. Multiple IDP policies allow administrators to reference a service set associated with a subscriber to a pre-configured IDP policy. This IDP policy is used to enforce security inspection for traffic per subscriber. Service set configuration is supported in-device in

MX series devices and IDP policies can be associated with service sets using the configuration node **Services > Service Interface Pool > Service Set**.

To create a new IDP policy:

- 1. In the main navigation tree, select Policies, then double-click the policy name in the Security Policies window. The Security Policy window appears.
- 2. Select the IDP tab in the Security Policy window.
- 3. Click **Add** in the Policies panel.
- 4. Enter a name for the policy and comments if desired, in the pop-up menu, and click **OK**. The new IDP policy is added to the Policies panel.

To add rules to the IDP policy:

- 1. In the main navigation tree, select **Security Policies**, then double-click the policy name in the Security Policies window. The Security Policy window appears.
- 2. Select the **IDP** tab in the Security Policy window.
- 4. Right-click on the policy name in the Policies panel and select **Add Rule**. The rule will be added to the IDP policy.



NOTE: If you select an IDP rule associated with multiple IDP policies from the IDP rule table in a Security Policy window, the Policies panel displays the multiple IDP policies to which the rule is associated.

To remove rules from the IDP policy:

- 1. In the main navigation tree, select **Security Policies**, then double-click the policy name in the Security Policies window. The Security Policy window appears.
- 2. Select the **IDP** tab in the Security Policy window.
- 3. Select a rule from the IDP rule table.
- 4. Right-click on the policy that includes the rule in the Policies panel, and select **Remove Rule**. The rule will be removed from the corresponding IDP policy in the Policies panel but will remain in the IDP rule table.



NOTE: For other devices which do not support multiple IDP policies, an IDP rule's association with multiple IDP policies on the Policies panel is ignored.



NOTE: From-Zone and To-Zones are not applicable to MX series devices and these values will be trimmed or ignored if configured.

Configuring Application Policy Enforcement (APE) Rules

You can configure APE rules to detect network traffic based on application signatures (rather than services, service contexts, and signatures) and to take a specified action. APE rules are supported on IDP standalone devices running IDP release 5.0.

You complete the steps in the following sections to create an APE rulebase:

- [“Adding the APE Rulebase Using the Policy Manager” on page 527](#) or [“Adding the APE Rulebase to a Policy Using the Application Profiler” on page 528](#)—Create, modify, or delete APE rules from the Policy Manager or you can select one or more traffic flows on the Application Profiler tab to create APE rules.
- [“Defining Matches For APE Rules” on page 529](#)— Define the type of network traffic you want IDP to monitor for applications, such as source/destination zones, source/destination address objects, and the application layer protocols (services) supported by the destination address object. You can also negate zones, address objects, or services.
- [“Configuring Actions For APE Rules” on page 531](#)— Specify the action you want IDP to take when the monitored traffic matches the rule's application objects. You can specify the action you want the security device to perform against the current connection and future connections from the same source IP address (see [Choosing an IP Action](#)).
- [“Configuring Notification in APE Rules” on page 534](#)— Disable or enable logging for the IDP rule.



NOTE: All APE rules are terminal. When a match is discovered in a terminal rule for the source, destination, service, and application, IDP does not continue to check subsequent rules for the same source, destination, service, and application.

Adding the APE Rulebase Using the Policy Manager

You can create APE rules based on Layer-7 applications and protocols. Before you can configure a rule in the APE rulebase, you need to add the APE rulebase to a security policy.

To configure an APE rulebase and APE rules:

1. In the main navigation tree, select **Policies**. Double-click the policy name in the security policies window or click the policy name and then select the Edit icon.
2. Click the Add icon in the upper right corner of the Security Policy window and select **Add APE Rulebase** to enable the APE rulebase tab.

3. To configure an APE rule, click the Add icon on the left side of the Security Policy window to open a default APE rule. You can modify the rule as necessary.
4. Click **OK**.



NOTE: Policy updates include custom applications on the IDP devices with application identification support, such as the ISG Series running ScreenOS 6.3 and IDP 5.0 or later.

Adding the APE Rulebase to a Policy Using the Application Profiler

From the Application Profiler view, you select from traffic flows to create corresponding APE rules in the APE rulebase. When you configure a new APE rule, the APE rulebase is automatically created.



NOTE: If you do not have appropriate access-control permission and you attempt to create APE rules, the wizard returns an error message stating that you do not have access to create rulebases.

To create APE rules for a policy from the Application Profiler:

1. From the Investigate panel, select **Security Monitor > Profiler**.
2. Select a traffic flow (row) from the Application Profiler view and right-click on a column row.
3. Right-click on the traffic flow row.
4. Select **Create Application Rules > For Policies**.

The New Application Rules dialog box is displayed.



NOTE: If an APE rulebase is not already configured, the rulebase is automatically configured when you add an APE rule to the security policy.

5. Select one or more policies to which you want to add application rules, and click **Next**.
6. From the New Application Rules dialog box, configure one or more application rules.
7. Click **Next**.
8. Verify that the new rules have been correctly configured in the policy, and click **Finish**.

Defining Matches For APE Rules

When creating your APE rules, you must specify the type of network traffic that you want IDP to monitor for applications.

The match columns From Zone, Source, To Zone, Destination, and Service are required for all rules in the APE rulebase. If IDP encounters a match for the other Match columns in an APE rule, no other rules in the rulebase are examined. .

The following sections describe the Match columns of an APE rule.

Configuring Applications for APE Rules

You can select one or more applications or an application group to monitor network traffic originating from or destined for any application.

1. To select or configure an application, right-click the Application column of a rule and select **Select Application**. The Select Application dialog box is displayed.
2. From the Select Application dialog box, check the check box next to each predefined or custom application, or an application group that you wish to add.

Configuring Source and Destination Zones for APE Rules (Does not Apply to Standalone IDP Sensor Rulebases)

You can select multiple zones for the source and destination. However, these zones must be available on the security devices on which you will install the policy. You can specify “any” for the source or destination zones to monitor network traffic originating or destined for any zone.



NOTE: You can create custom zones for some security devices. The list of zones from which you can select source and destination zones includes the predefined and custom zones that have been configured for all devices managed by NSM. Therefore, you should only select zones that are applicable for the device on which you will install the security policy.

Configuring Source and Destination Address Objects for APE Rules

In the NSM system, address objects are used to represent components on your network: hosts, networks, servers, and so on. Typically, a server or other device on your network is the destination IP for incoming attacks, and can sometimes be the source IP for interactive attacks (see “[Configuring Backdoor Rules](#)” on page 538 for more information on interactive attacks). You can specify “any” to monitor network traffic originating from any IP address. You can also “negate” the address objects listed in the Source or Destination column to specify all sources or destinations except the excluded objects.

You can create address objects either before you create an APE rule or while creating or editing an APE rule. To select or configure an address object, right-click either the Source or Destination column of a rule and select **Select Address**. In the Select Source Addresses

dialog box, you can either select an already created address object or click the Add icon to create a new host, network, or group object.



NOTE: You can select either a user role or a source IP address for the APE rule, but not both.

To detect incoming attacks that target your internal network, set the From Zone to Untrust, and the Source IP to any IP. Then set the To Zone to dmz and trust. Next, select the address object that represents the host or server you want to protect from attacks as the Destination IP.

To detect attacks between two network, select multiple address objects for the Source and Destination.

The more specific you are in defining the source and destination of an attack, the more you reduce false positives.

Configuring User Roles for APE Rules

User roles are configured in conjunction with source IP addresses. You select either a user role or a source IP address for the APE rule. If you configure a user role in a APE rule, you must also set the source address to “any”. NSM does not automatically set the source address to “any” when a user role is configured in the rule but displays a message that only a user role or a source address can be specified in a rule.

1. To select or configure a user role, right-click the User Role column of a rule and select **Select User Role**.
2. From the Select User Roles dialog box, select a device from the Device drop down menu.
3. Use the **Add** or **Remove** button to add or remove user roles.
4. Click **OK**.

Configuring Services for APE Rules

Services are application layer protocols that define how data is structured as it travels across the network. Because the services you support on your network are the same services that attackers must use to attack your network, you can specify which services are supported by the destination IP to make your rule more efficient.



NOTE: All services rely on a transport layer protocol to transmit data. IDP includes services that use TCP, UDP, RPC, and ICMP transport layer protocols.

Service objects represent the services running on your network. NSM includes predefined service objects that are based on industry-standard services. You use these service

objects in rules to specify the service an attack uses to access your network. You can also create custom service objects to represent protocols that are not included in the predefined services.

In the Service column you select the service of the traffic you want IDP to match:

- Select **Default** to accept the service specified by the attack object you select in the Attacks column. When you select an attack object in the Attack column, the service associated with that attack object becomes the default service for the rule. To see the exact service, view the attack object details.
- Select **Any** to set any service.
- Select **Service** to choose specific services from the list of defined service objects.

For example, to take some action on FTP traffic, set the service to Default and add the application object FTP. The Service column in the rule still displays “Default,” but the rule actually uses the default service of TCP-FTP, which is specified in the application object.

You can create your own service objects to use in rules, such as service objects for protocols that use nonstandard ports. However, you cannot match application objects to protocols that they do not use.

Configuring Actions For APE Rules

You can tell the security device which actions to perform against attacks that match rules in your security policy. For each attack that matches a rule, you can choose to ignore, drop, or close the current packets or connection. If the rule is triggered, the device can perform actions against the connection.

Remember that the device can drop traffic only when IDP is enabled in inline mode; when IDP is enabled in inline tap (sniffer) mode, it cannot perform drop or close actions.

Table 47 on page 531 lists actions for APE rules:

Table 47: APE Rule Actions

Action	Description
None	IDP takes no action against the connection. If a rule that contains an action of None is matched, the corresponding log record displays “accept” in the action column of the Log Viewer.
Drop Connection	IDP drops the connection without sending an RST packet to the sender, preventing the traffic from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
Close Client	IDP closes the connection to the client, but not to the server.
Close Server	IDP closes the connection to the server, but not to the client.

Table 47: APE Rule Actions (continued)

Action	Description
Close Client and Server	IDP closes the connection and sends a RST packet to both the client and the server. If IDP is operating in inline tap mode, IDP sends an RST packet to both the client and server but does not close the connection.
Diffserv Marking	IDP assigns the service differentiation value indicated to the packet, then passes it on normally. The value is set in the dialog that appears when you select this action in the rulebase.
Rate Limiting	<p>IDP enforces a rate limit for all current sessions that match the rule. If the limit has not been reached, the IDP appliance forwards the packets. If the limit has been reached, the IDP appliance behaves as if bandwidth is unavailable; it drops packets until the aggregate bandwidth falls below the limit. When the IDP appliance drops packets, the TCP or UDP endpoints identify the packet loss and slow down the transmission rate.</p> <p>Specify rate limits depending on the bandwidth for your links. If you have a 1 Gbps link, and want no more than 10% available to peer-to-peer traffic, the sum of the rate limits you specify for all peer-to-peer rules should be less than 102.4 Mbps (in each direction). You configure separate rate limits for client-to-server and server-to-client directions. For peer-to-peer traffic, we recommend you set the same rate for each direction.</p> <p>Standalone IDP sensors running 5.1 and later assign the rate-limit for client-to-server and server-to-client traffic in Kbps. Make sure that you enable the rate-limit feature for the device. For more information on enabling this feature, see <i>Intrusion Detection and Prevention Administration Guide</i>.</p> <p>NOTE: For TFTP traffic, all traffic is counted as client-to-server traffic. A TFTP server responds to <i>get</i> requests by establishing an ephemeral port from which to send the reply. In this case, both directions appear to the IDP appliance as client-to-server flows. We recommend you set the same rate for each direction.</p>
Diffserv Marking & Rate Limiting	<p>IDP enforces the user-prescribed rate limit (in Kbps) on the session that matches the rule and assigns the user-specified service differentiation value to the packets in that session.</p> <p>The service differentiation value and the rate-limit value are set in the dialog that appears when you select this action in the APE rulebase.</p> <p>Only standalone IDP sensors running IDP 5.1 and later support this feature. When you select this action in an APE rule installed on a device running IDP 5.0 or earlier, NSM displays a message warning the user that if unsupported, this APE action might cause a device update failure.</p>

Configuring IP Actions in APE Rules

This column only appears when you view the security policy in Expanded Mode. To change the security policy view from Compact Mode to Expanded Mode, from the menu bar, select **View > Expanded Mode**.

If the current network traffic matches a rule, the security device can perform an IP action against future network traffic that uses the same IP address. IP actions are similar to other actions; they direct the device to drop or close the connection. However, because you now also have the attacker's IP address, you can choose to block the attacker for a specified amount of time. If attackers cannot immediately regain a connection to your network, they might try to attack easier targets.

Use IP actions in conjunction with actions and logging to secure your network. In a rule, first configure an action to detect and prevent current malicious connections from reaching your address objects. Then right-click in the IP Action column of the rule and select **Configure**. Enable and configure an IP action to prevent future malicious connections from the attacker's IP address.

Choosing an IP Action

For each IP action option, an IP action is generated by NSM. The IP action instructs the security device to perform the specified task. Select from the following options:

- IDP Notify—The security device does not take any action against future traffic, but logs the event. This is the default.
- IDP Drop—The security device drops the matching connection and blocks future connections that match the criteria set in the Block list.
- IDP Close—The security device closes future connections that match the criteria in the Block list.

Choosing a Block Option

Each block option follows the criteria you set in the Actions box. Block options can be based on the following matches of the attack traffic:

- Source, Destination, Destination Port and Protocol—The security device blocks future traffic based on the source, destination, destination port, and protocol of the attack traffic. This is the default.
- Source—The security device blocks future traffic based on the source of the attack traffic.
- Destination—The security device blocks future traffic based on the destination of the attack traffic.
- From Zone, Destination, Destination Port and Protocol—The security device blocks future traffic based on the source zone, destination, destination port, and protocol of the attack traffic.
- From Zone—The security device blocks future traffic based on the source zone of the attack traffic.

Setting Logging Options

When the security device detects attack traffic that matches a rule and an IP action is triggered, the device can log information about the IP action that was taken or create an alert in the Log Viewer. By default, there are no logging options set.

Setting Timeout Options

You can set the number of seconds that you want the IP action to remain in effect after a traffic match. For permanent IP actions, leave the timeout at 0 (this is the default).

Configuring Notification in APE Rules

You can log an attack and create log records with attack information that you can view in realtime in the Log Viewer. For more critical attacks, you can also set an alert flag to appear in the log record.

To log an attack for a rule, right-click the Notification column of the rule and select **Configure**.

The first time you design a security policy, you might be tempted to log all attacks and let the policy run indefinitely. Do not do this! Some attack objects are informational only, and others can generate false positives and redundant logs. If you become overloaded with data, you can miss something important. Security policies that generate too many log records are hazardous to the security of your network, because you might discover an attack too late or miss a security breach entirely due to sifting through hundreds of log records. Excessive logging can also affect throughput, performance, and available disk space. A good security policy generates enough logs to fully document only the important security events on your network.



NOTE: J Series routers and SRX Series gateways do not send packet data to NSM. If your policy rules attempt to do so, NSM does not log the data.

- **Setting Logging**—In the Configure Notification dialog box, select **Logging** and then click **OK**. Each time the rule is matched, the NSM system creates a log record that appears in the Log Viewer.
- **Setting an Alert**—In the Configure Notification dialog box, select **Alert** and then click **OK**. If the rule is matched, the security device places an alert flag in the Alert column of the Log Viewer for the matching log record.
- **Logging Packets**—You can record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack and its purpose, whether or not the attack was successful, and any possible damage to your network.



NOTE: To improve performance, log only the packets after the attack.

If multiple rules with packet capture enabled match the same attack, the security device captures the maximum specified number of packets. For example, you configure Rule 1 to capture 10 packets before and after the attack, and Rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.



NOTE: Packet captures are restricted to 256 packets before and after the attack.

Setting VLAN Tags for APE Rules

You can specify that the rule be applied only to packets from particular VLANs. See “Setting VLAN Tags for IDP Rules” on page 524 more information.

Setting Severity for APE Rules

This column only appears when you view the security policy in Expanded Mode. To change the security policy view from Compact Mode to Expanded Mode, from the menu bar, select **View > Expanded Mode**.

You can override the inherent attack severity on a per-rule basis within the APE rulebase. You can set the severity to Default, Info, Warning, Minor, Major, or Critical.

To change the severity for a rule, right-click the Severity column of the rule and select a severity.

Setting Target Security Devices for APE Rules

For each rule in the APE rulebase, you can select the security device on which the rule is installed. When you install the security policy that the rule belongs to, the rule becomes active only on the devices you selected in the Install On column of the rulebase.



NOTE: NSM supports APE only on standalone IDP devices running IDP release 5.0.

Entering Comments for APE Rules

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select **Edit Comments**. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

Configuring Exempt Rules

The Exempt rulebase works in conjunction with the IDP rulebase. Before you can create exempt rules, you must first create rules in the IDP rulebase. If traffic matches a rule in the IDP rulebase, IDP attempts to match the traffic against the Exempt rulebase before performing the specified action or creating a log record for the event.



NOTE: If you delete the IDP rulebase, the Exempt rulebase is also deleted.

You might want to use an exempt rule when an IDP rule uses an attack object group that contains one or more attack objects that produce false positives or irrelevant log records. To prevent unnecessary alarms, you might want to use an exempt rule to exclude a specific source, destination, or source/destination pair from matching an IDP rule.

When you create an exempt rule, you must specify the following:

- Source and destination for traffic you want to exempt. You can set the source or destination to “any” to exempt network traffic originating from any source or sent to any destination. You can also specify “negate” to specify all sources or destinations except the specified addresses.
- The attacks you want IDP to exempt for the specified source/destination addresses. You must include at least one attack object in an exempt rule.



NOTE: The Exempt rulebase is a non-terminal rulebase. That is, IDP attempts to match traffic against all rules in the Exempt rulebase and all matches are executed.

Adding the Exempt Rulebase

Before you can configure a rule in the Exempt rulebase, you need to add the Exempt rulebase to a security policy.

1. In the main navigation tree, select **Policies**. Open a security policy by double-clicking the policy name in the security policies window or click the policy name and then select the Edit icon.
2. Click the Add icon in the upper right corner of the Security Policy window and select **Add Exempt Rulebase** to enable the Exempt rulebase tab.
3. To configure an exempt rule, click the Add icon on the left side of the Security Policy window to open a default exempt rule. You can modify this rule as necessary.

Defining a Match

You specify the traffic you want to exempt from attack detection. The Match columns From Zone, Source, To Zone, and Destination are required for all rules in the exempt rulebase.

The following sections detail the Match columns of an exempt rule.

Configuring Source and Destination Zones

You can select multiple zones for the source and destination, however these zones must be available on the security devices on which you will install the policy. You can specify “any” for the source or destination zones to monitor network traffic originating or destined for any zone.



NOTE: You can create custom zones for some security devices. The list of zones from which you can select source and destination zones includes the predefined and custom zones that have been configured for all devices managed by NSM. Therefore, you should only select zones that are applicable for the device on which you will install the security policy.

Configuring Source and Destination Address Objects

In the NSM system, address objects are used to represent components on your network: hosts, networks, servers, etc. You can specify "any" to monitor network traffic originating from any IPv4 address and "AnyIPv6" to monitor network traffic originating from any IPv6 address. You can also negate the address objects listed in the Source or Destination column to specify all sources or destinations except the excluded object.

You can create address objects either before you create an exempt rule or while creating or editing an exempt rule. To select or configure an address object, right-click either the Source or Destination column of a rule and select **Select Address**. In the Select Source Addresses dialog box, you can either select an already-created address object or click the Add icon to create a new host, network, or group object.

To improve performance and eliminate false positives between your Internal Lab devices and your Engineering desktops, you want to exempt attack detection.

Setting Attack Objects

You specify the attacks you want IDP to exempt for the specified source/destination addresses. You must include at least one attack object in an exempt rule.

You consistently find that your security policy generates false positives for the attack HTTP Buffer Overflow: Header on your internal network. You want to exempt attack detection for this attack when the source IP is from your internal network.

Specifying VLANs

You can specify that the rule be applied only to packets from particular VLANs. See [“Setting VLAN Tags for IDP Rules” on page 524](#) more information.

Setting Target Devices

For each rule in the rulebase, you can select the IDP-capable device that will use that rule to detect and prevent attacks. Alternatively, you can use Device Manager to assign policies to devices.

Entering Comments

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select Edit Comments. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

Creating an Exempt Rule from the Log Viewer

You can also create a rule in the Exempt rulebase directly from the NSM Log Viewer. You might want to use this method to quickly eliminate rules that generate false positive log records. .

To create an exempt rule from the Log Viewer:

1. View the IDP/DI logs in the Log Viewer.
2. Right-click a log record that contains an attack you want to exempt and select Exempt.

The Exempt rulebase for the security policy that generated the log record is displayed, with the exempt rule that is associated with the log entry. The source, destination, and attack settings for the rule are automatically filled in based on the information in the log record.



NOTE: If the Exempt rulebase does not already exist when you create an exempt rule from the Log Viewer, the rulebase is automatically created and the rule is added.

You can modify, reorder, or merge an exempt rule created from the Log Viewer in the same manner as any other exempt rule that you create directly in the Exempt rulebase.

Configuring Backdoor Rules

A backdoor is a mechanism installed on a host computer that facilitates unauthorized access to the system. Attackers who have already compromised a system can install a backdoor to make future attacks easier. When attackers type commands to control a backdoor, they generate interactive traffic.

Interactive traffic is traffic that indicates human involvement in a normally automated process, such as a user typing commands. Interactive traffic looks different than other traffic because humans are manually controlling one end of the connection. In a connection between two programs, the data transfer is automated; TCP packets can be batched and sent in bulk for efficiency. In a connection between a program and a user, packets are sent when they become available; characters display as they are typed (not after the word is complete). Interactive programs transmit several short IP packets containing individual keystrokes and their echoes, reflecting the real-time actions of a user (or attacker).

When attackers type commands to control a backdoor, they generate interactive traffic that IDP can detect. Unlike antivirus software, which scans for known backdoor files or executables on the host system, IDP detects the interactive traffic that is produced when backdoors are used. This method ensures that IDP can detect all backdoors, both known and unknown. If interactive traffic is detected, IDP can perform IDP actions against the connection to prevent the attacker from further compromising your network.

When you configure a backdoor rule, you must specify the following:

- Source and destination addresses for traffic you want to monitor. To detect incoming interactive traffic, set the Source to “any” and the Destination to the IP address of network device you want to protect. To detect outgoing interactive traffic, set the Source to the IP address of the network device you want to protect and the Destination to “any”.
- Services that are offered by the Source or Destination as well as interactive services that can be installed and used by attackers.



NOTE: Do not include TELNET, SSH, RSH, NETMEETING, or VNC as services, as these services are often used to legitimately control a remote system. Including these services can generate false positives.

- Action that the IDP is to perform if interactive traffic is detected. Set the Operation to “detect”. If you are protecting a large number of network devices from interactive traffic, you can create a rule that “ignores” accepted forms of interactive traffic from those devices, then create another rule that “detects” all interactive traffic from those devices.



NOTE: The Backdoor rulebase is a terminal rulebase. That is, when IDP finds a match on a rule in the Backdoor rulebase, it does not execute succeeding rules.

Adding the Backdoor Rulebase

Before you can configure a rule in the Backdoor rulebase, you need to add the Backdoor rulebase to a security policy.

1. In the main navigation tree, select **Policies**. Open a security policy by double-clicking the policy name in the security policies window or click the policy name and then select the Edit icon.
2. Click the Add icon in the upper right corner of the Security Policy window and select **Add Backdoor Rulebase**. The Backdoor rulebase tab appears.
3. Configure a backdoor rule by clicking the Add icon on the left side of the Security Policy window. A default backdoor rule appears. You can modify this rule as needed.

Defining a Match

You specify the traffic you want to IDP to monitor for indications of backdoors. The Match columns From Zone, Source, To Zone, Destination, and Service are required for all rules in the Backdoor rulebase.

The following sections detail the Match columns of a backdoor rule.

Configuring Source and Destination Zones

You can select multiple zones for the source and destination, however these zones must be available on the security devices on which you will install the policy. You can specify "any" for the source or destination zones to monitor network traffic originating or destined for any zone.



NOTE: You can create custom zones for some security devices. The list of zones from which you can select source and destination zones includes the predefined and custom zones that have been configured for all devices managed by NSM. Therefore, you should only select zones that are applicable for the device on which you will install the security policy.

Configuring Source and Destination Address Objects

In NSM, address objects are used to represent components on your network: hosts, networks, servers, and so on. Typically, a server or other device on your network is the destination IP for incoming attacks, and can sometimes be the source IP for interactive attacks. You can specify "any" to monitor network traffic originating from any IPv4 address and "AnyIPv6" to monitor network traffic originating from any IPv6 address. You can also negate the address objects listed in the Source or Destination column to specify all sources or destinations except the excluded address object.

You can create address objects either before you create a backdoor rule or while creating or editing an backdoor rule. To select or configure an address object, right-click either the Source or Destination column of a rule and select Select Address. In the Select Source Addresses dialog box, you can either select an already-created address object or click the Add icon to create a new host, network, or group object.

Configuring Services

Select interactive service objects. Be sure to include services that are offered by the source or destination IP as well as interactive services that are not; attackers can use a backdoor to install any interactive service. Do not include telnet, SSH, RSH, netmeeting, or VNC, as these services are often used to remotely control a system legitimately and their inclusion might generate false positives.

Setting Operation

Set the Operation to detect or ignore. If you select detect, choose an action to perform if backdoor traffic is detected. If you are protecting a large number of address objects from interactive traffic, you can create a rule that ignores accepted forms of interactive traffic from those objects, then create a succeeding rule that detects all interactive traffic from those objects.

Setting Actions

Choose an action to perform from [Table 48 on page 541](#) if IDP detects interactive traffic:

Table 48: Actions for Backdoor Rule:

Action	Description
Accept	IDP accepts the interactive traffic.
Drop Connection	IDP drops the interactive connection without sending a RST packet to the sender, preventing the traffic from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
Close Client and Server	IDP closes the interactive connection and sends a RST packet to both the client and the server. If the IDP is in sniffer mode, IDP sends a RST packet to both the client and server but does NOT close the connection.
Close Client	IDP closes the interactive connection to the client, but not to the server.
Close Server	IDP closes the interactive connection to the server, but not to the client.

Setting Notification

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, you can also set an alert flag to appear in the log record.

To log an attack for a rule, right-click the Notification column of the rule and select **Configure**. The Configure Notification dialog box appears.

The first time you design a security policy, you might be tempted to log all attacks and let the policy run indefinitely. Don't do this! Some attack objects are informational only, and others can generate false positives and redundant logs. If you become overloaded with data, you can miss something important. Remember that security policies that generate too many log records are hazardous to the security of your network, as you might discover an attack too late or miss a security breach entirely due to sifting through hundreds of log records. Excessive logging can also affect IDP throughput, performance, and available disk space. A good security policy generates enough logs to fully document only the important security events on your network.

Setting Logging

In the Configure Notification dialog box, select **Logging** and then click **OK**. Each time the rule is matched, the IDP system creates a log record that appears in the Log Viewer.

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, however, you might want to be notified immediately by e-mail, have IDP run a script in response to the attack, or set an alarm flag to appear in the log record. Your goal is to fine-tune the attack notifications in your security policy to your individual security needs.

Setting an Alert

In the Configure Notification dialog box, select **Alert** and then click **OK**. If Alert is selected and the rule is matched, IDP places an alert flag in the alert column of the Log Viewer for the matching log record.

Logging Packets

You can record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack and its purpose, whether or not the attack was successful, and any possible damage to your network.



NOTE: To improve IDP performance, log only the packets after the attack.

If multiple rules with packet capture enabled match the same attack, IDP captures the maximum specified number of packets. For example, you configure Rule 1 to capture 10 packets before and after the attack, and Rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.



NOTE: Packet captures are restricted to 256 packets before and after the attack.

Setting Severity

You can override the inherent attack severity on a per-rule basis within the Backdoor rulebase. You can set the severity to either Default, Info, Warning, Minor, Major, or Critical.

To change the severity for a rule, right-click the Severity column of the rule and select a severity.

Specifying VLANs

You can specify that the rule be applied only to packets from particular VLANs. See [“Setting VLAN Tags for IDP Rules” on page 524](#) for more information.

Setting Target Devices

For each rule in the rulebase, you can select the IDP-capable device that will use that rule to detect and prevent attacks. Alternatively, you can use Device Manager to assign policies to devices.

Entering Comments

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select **Edit Comments**. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

Configuring SYN Protector Rules

The SYN-Protector rulebase protects your network from SYN floods by ensuring that the three-way handshake is performed successfully for specified TCP traffic. If you know

that your network is vulnerable to a SYN flood, use the SYN Protector rulebase to prevent it.

The TCP Handshake

When a TCP connection is initiated, a three-way handshake takes place:

- A client host sends a SYN packet to a specific port on the server to request a connection.
- Next, the server sends the client host a SYN/ACK packet, which both acknowledges (ACK) the original SYN packet from the client host and forwards a new SYN packet. The potential connection is now in a SYN_RECV state.
- Finally, the client host sends an ACK packet to the server to acknowledge receipt of the SYN/ACK packet. The connection is now in an ESTABLISHED state.

This three-way handshake contains an inherent, exploitable vulnerability that attackers can use to disable the system: a SYN flood. Most systems allocate a large, but finite number of resources to a connection table that is used to manage potential connections. While the connection table can sustain hundreds of concurrent connections across multiple ports, attackers can generate enough connection requests to exhaust all allocated resources.

SYN-Floods

Attackers initiate a SYN flood by manipulating the basic three-way handshake:

- A client host sends a SYN packet to a specific port on the server. However, the attacker ensures that the client host's IP address is a spoofed IP address of an unreachable system.
- Next, the server sends the client host (spoofed address) a SYN/ACK packet. The potential connection is now in a SYN_RECV state.
- Since the system is unreachable, the server never receives an ACK or RST packet back from the client host. The potential connection is now in the SYN_RECV state, and is placed into a connection queue while it waits for an ACK or RST packet. This potential connection remains in the queue until the connection-establishment timer expires (when it will be deleted).
- The attacker sends another SYN packet to the server, requesting another connection. And then another. And another. The connection table fills to capacity and cannot accept new SYN requests. The server is overwhelmed, and quickly becomes disabled.

By default, the SYN Protector rulebase is only activated when the number of SYN packets per second is greater than 1020. This number is the sum of two parameters that you can set in the Sensor Settings Run-Time Parameters:

- Lower SYN's-per-second threshold below which SYN Protector will be deactivated (the default value is 1000)
- Upper SYN's-per-second threshold above which SYN Protector will be activated (the default value is 20)

Once the SYN Protector rulebase is activated, it remains active until the number of SYN packets per second is less than the Lower SYN's-per-second threshold (which is 1000 by default).

Adding the SYN Protector Rulebase

Before you can configure a rule in the SYN Protector rulebase, you need to add the SYN Protector rulebase to a security policy.

1. In the main navigation tree, select **Policies**. Open a security policy by double-clicking the policy name in the Security Policies window or click the policy name and then select the Edit icon.
2. Click the Add icon in the upper right corner of the Security Policy window and select **Add SYN Protector Rulebase** to open the SYN Protector rulebase tab.
3. Configure a SYN Protector rule by clicking the Add icon on the left side of the Security Policy window to open a default SYN Protector rule. You can modify this rule as needed.

Defining a Match

Specify the traffic you want IDP to monitor for SYN floods.

Configuring Source and Destination Address Objects

Set the Source Object to Any. Set the Destination Object to any address objects you want to protect.

Configuring Services

The default service, TCP-any, looks for SYN floods in all TCP-based traffic.

Always set the SYN Protector service value to TCP-any. Selecting individual services can cause unpredictable interactions with other rulebases.

Setting Mode

Select the mode that indicates how IDP handles TCP traffic:

- **None.** IDP takes no action, and does not participate in the three-way handshake.
 - **Relay.** IDP acts as the middleman, or relay, for the connection establishment, performing the three-way handshake with the client host on behalf of the server. Relay mode guarantees that the server allocates resources only to connections that are already in an ESTABLISHED state. The relay is transparent to both the client host and the server.

IDP receives the initial SYN packet sent by the client host and returns a SYN/ACK packet. If the client host sends an ACK packet, IDP completes the three-way handshake and allows the connection to move to an ESTABLISHED state. If IDP does not receive an ACK packet from the client host, as would be the case during a SYN flood attack, IDP does not complete the three-way handshake and the connection is not established.

- **Passive.** IDP handles the transfer of packets between the client host and the server, but does not actively prevent the connection from being established. Instead, IDP uses a timer to ensure that connections are established promptly, minimizing the use of server resources. The timer IDP uses for the connection establishment is shorter than the timer the server uses for the connection queue.

IDP transfers the SYN packet sent by the client host to the server, then transfers the SYN/ACK packet sent by the server to the client host. If the client host sends an ACK packet to the server before the IDP connection timer expires, the connection is established. If the client host does not send an ACK packet to the server, as would be the case during a SYN flood attack, the IDP connection timer expires. IDP resets the connection to free resources on the server.

Setting Notification

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, you can also set an alert flag to appear in the log record.

To log an attack for a rule, right-click the Notification column of the rule and select **Configure**. The Configure Notification dialog box appears.

The first time you design a security policy, you might be tempted to log all attacks and let the policy run indefinitely. Don't do this! Some attack objects are informational only, and others can generate false positives and redundant logs. If you become overloaded with data, you can miss something important. Remember that security policies that generate too many log records are hazardous to the security of your network, as you might discover an attack too late or miss a security breach entirely due to sifting through hundreds of log records. Excessive logging can also affect IDP throughput, performance, and available disk space. A good security policy generates enough logs to fully document only the important security events on your network.

Setting Logging

In the Configure Notification dialog box, select **Logging** and then click **OK**. Each time the rule is matched, the IDP system creates a log record that appears in the Log Viewer.

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, however, you might want to be notified immediately by e-mail, have IDP run a script in response to the attack, or set an alarm flag to appear in the log record. Your goal is to fine-tune the attack notifications in your security policy to your individual security needs.

Setting an Alert

In the Configure Notification dialog box, select **Alert** and then click **OK**. If Alert is selected and the rule is matched, IDP places an alert flag in the alert column of the Log Viewer for the matching log record.

Logging Packets

You can record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. Viewing the packets used in an attack on your

network can help you determine the extent of the attempted attack and its purpose, whether or not the attack was successful, and any possible damage to your network.



NOTE: To improve IDP performance, log only the packets after the attack.

If multiple rules with packet capture enabled match the same attack, IDP captures the maximum specified number of packets. For example, you configure Rule 1 to capture 10 packets before and after the attack, and Rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.



NOTE: Packet captures are restricted to 256 packets before and after the attack.

Setting Severity

You can override the inherent attack severity on a per-rule basis within the SYN Protector rulebase. You can set the severity to either Default, Info, Warning, Minor, Major, or Critical.

To change the severity for a rule, right-click the Severity column of the rule and select a severity.

Specifying VLANs

You can specify that the rule be applied only to packets from particular VLANs. See [“Setting VLAN Tags for IDP Rules” on page 524](#) for more information.

Setting Target Devices

For each rule in the rulebase, you can select the IDP-capable device that will use that rule to detect and prevent attacks. Alternatively, you can use Device Manager to assign policies to devices.

Entering Comments

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select **Edit Comments**. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

Configuring Traffic Anomalies Rules

Traffic anomaly rules protect your network from attacks by using traffic flow analysis to identify attacks that occur over multiple connections and sessions (such as scans).

Before attempting to enter an unknown network, attackers often gather information about the network and analyze any weaknesses to help them choose the best attack method. A port scan or network scan is often the first reconnaissance performed. Attackers typically use a scanning tool that attempts to connect to every port on a single machine

(port scanning) or connect to multiple IP addresses on a network (network scanning). By determining which services are allowed and responding on your network, attackers can gain valuable information about your network configuration.

To detect scans and other distributed network attacks, the Traffic Anomalies Rulebase looks for patterns that indicate abnormal network activity. Attackers often use scanning tools to automate their port scans, allowing them to scan multiple ports quickly and efficiently. IDP can detect these scans by counting the number of ports scanned in a specified time period. You can also set a session limit threshold, which defines the maximum number of sessions for a single host.

Detecting TCP and UDP Port Scans

To detect TCP and UDP port scans, set a port count (number of ports scanned) and the time threshold (the time period that ports are counted) in seconds.

Example: Traffic Anomalies Rule

You want to create a Traffic Anomalies rule that looks for port scans on your internal network. You set both the TCP and UDP Port Count to 20 and the Time threshold to 120 seconds. The rule is matched if the same Source IP scans 20 TCP ports on your internal network within 120 seconds, or if the same Source IP scans 20 UDP ports on your internal network within 120 seconds.

Detecting Other Scans

In addition to port scans, the attacks can occur over multiple connections and sessions:

- **Distributed Port Scans.** Use multiple Source IP addresses to scan ports.
- **ICMP Sweeps.** Use a single Source IP to ping multiple IP addresses.
- **Network Scans.** Use a single Source IP to scan multiple IP addresses.

To detect these attacks, set the IP Count (the number of times attempts to scan or ping ports on your network occur) and the Time (the time period that IP addresses are counted) in seconds.

Example: Traffic Anomalies Rule

To create a Traffic Anomalies rule that looks for distributed port scans on your internal network, set the IP Count to 50 and the Time to 120 seconds. If 50 IP addresses attempt to scan ports on your internal network within 120 seconds, the rule is matched.

Example: Traffic Anomalies Rule

You want to create a Traffic Anomalies rule that looks for network scans and ICMP sweeps on your internal network. You set the IP Count to 50 and the Time to 120 seconds for ICMP sweeps and network scans. The rule is matched if:

- The same Source IP attempts to scan 50 IP addresses on your internal network within 120 seconds
- The same Source IP attempts to ping 50 IP addresses on your internal network within 120 seconds

Session Limiting

You can set a session limit threshold that defines the maximum number of sessions allowed from a single host within a second. For each source IP specified in the rule, the Sensor tracks the sessions per second; if the session rate exceeds the user-defined maximum, the Sensor generates a `SCAN_SESSION_RATE_EXCEEDED` event log record, which appears in the Log Viewer. To take action when this event is triggered, configure an IP action in the rule.

Example: Session Limiting

Your internal network typically has a low volume traffic. To detect a sudden increase in traffic from a specific host (which might indicate a worm), set the source IP to your Internal Network and the configure the session count as 200 session/sec. To block traffic that exceeds the session limit, set an IP action of IDP Block and chose Source, Protocol from the Blocking Options menu.

Adding the Traffic Anomalies Rulebase

Before you can configure a rule in the Traffic Anomalies rulebase, you need to add the Traffic Anomalies rulebase to a security policy.

1. In the main navigation tree, select **Policies**. Open a security policy by double-clicking the policy name in the Security Policies window or by clicking the policy name and then selecting the Edit icon.
2. Click the Add icon in the upper right corner of the Security Policy window and select **Add Traffic Anomalies Rulebase** to open the Traffic Anomalies rulebase tab.
3. Configure a Traffic Anomalies rule by clicking the Add icon on the left side of the Security Policy window to open a default Traffic Anomalies rule. You can modify this rule as needed.

Defining a Match

You specify the traffic you want IDP to monitor for network anomalies.

Configuring Source and Destination Address Objects

Set the Source Object to Any. Set the Destination Object to any address objects you want to protect.

Configuring Services

Set the Service to Any, unless you want to tailor different rules to different services.

Setting Detect Options

Right-click the rulebase cell in the Traffic anomalies column and select **Detect**. In the View Detect Options dialog, set the Port Count and Time Threshold values for each value you want to monitor. The values are measure in number of hits (Port Count) in a particular number of seconds (Time Threshold).

Setting Response Options

The IP Action column governs what action the IDP Sensor takes when it finds a matching condition.

Right-click the rulebase cell in the IP Action column and select **Configure**. The Configure IP Action dialog displays.

Configure your IP Action settings as appropriate for your network.

Setting Notification

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, you can also set an alert flag to appear in the log record.

To log an attack for a rule, right-click the Notification column of the rule and select **Configure**. The Configure Notification dialog box appears.

The first time you design a security policy, you might be tempted to log all attacks and let the policy run indefinitely. Don't do this! Some attack objects are informational only, and others can generate false positives and redundant logs. If you become overloaded with data, you can miss something important. Remember that security policies that generate too many log records are hazardous to the security of your network, as you might discover an attack too late or miss a security breach entirely due to sifting through hundreds of log records. Excessive logging can also affect IDP throughput, performance, and available disk space. A good security policy generates enough logs to fully document only the important security events on your network.

Setting Logging

In the Configure Notification dialog box, select **Logging** and then click **OK**. Each time the rule is matched, the IDP system creates a log record that appears in the Log Viewer.

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, however, you might want to be notified immediately by e-mail, have IDP run a script in response to the attack, or set an alarm flag to appear in the log record. Your goal is to fine-tune the attack notifications in your security policy to your individual security needs.

Setting an Alert

In the Configure Notification dialog box, select **Alert** and then click **OK**. If Alert is selected and the rule is matched, IDP places an alert flag in the alert column of the Log Viewer for the matching log record.

Logging Packets

You can record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack and its purpose, whether or not the attack was successful, and any possible damage to your network.



NOTE: To improve IDP performance, log only the packets after the attack.

If multiple rules with packet capture enabled match the same attack, IDP captures the maximum specified number of packets. For example, you configure Rule 1 to capture 10 packets before and after the attack, and Rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.



NOTE: Packet captures are restricted to 256 packets before and after the attack.

Setting Severity

You can override the inherent attack severity on a per-rule basis within the SYN Protector rulebase. You can set the severity to either Default, Info, Warning, Minor, Major, or Critical.

To change the severity for a rule, right-click the Severity column of the rule and select a severity.

Specifying VLANs

You can specify that the rule be applied only to packets from particular VLANs. See [“Setting VLAN Tags for IDP Rules” on page 524](#) for more information.

Setting Target Devices

For each rule in the rulebase, you can select the IDP-capable device that will use that rule to detect and prevent attacks. Alternatively, you can use Device Manager to assign policies to devices.

Entering Comments

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select **Edit Comments**. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

Configuring Network Honeypot Rules

The Network Honeypot protects your network by impersonating open ports on existing servers on your network, alerting you to attackers performing port scans and other information-gathering activities.

Impersonating a Port

Attackers view ports as entry points into your network. You can create counterfeit ports on existing servers to trick attackers who are attempting to break into your network. A counterfeit port can appear to offer notoriously vulnerable services to make the port attractive to attackers.

- You create a counterfeit port in the Network Honeypot Rulebase by specifying an existing network object and choosing a port and service to impersonate. You can also set an IP Action to perform against the Source IP. If an attacker attempts to communicate with your counterfeit port, the rule matches and the IP action triggers.

Adding the Network Honeypot Rulebase

Before you can configure a rule in the Network Honeypot rulebase, you need to add the Network Honeypot rulebase to a security policy.

1. In the main navigation tree, select **Policies**. Open a security policy by double-clicking the policy name in the Security Policies window or click the policy name and then select the Edit icon.
2. Click the Add icon in the upper right corner of the Security Policy window and select **Add Network Honeypot Rulebase**. The Network Honeypot rulebase tab appears.
3. Configure a Network Honeypot rule by clicking the Add icon on the left side of the Security Policy window. A default Network Honeypot rule appears. You can modify this rule as needed.



NOTE: Honeypot rulebase is not supported in IDP 4.2. If the assigned policy has honeypot rule included, NSM will remove it before pushing the policy to the device and provide a warning in the Job manager.

Defining a Match

You specify the traffic you want IDP to monitor for network anomalies.

Configuring the Source

Set the Source object to Any.

Configuring Destination Address Objects and Services

Set the Destination Address and Service to the service that will appear to be available on the indicated address object.

Setting Operation

Right-click the cell in the Operation column and select Impersonate. This tells the IDP Sensor to impersonate the indicated services on the indicated device.

Setting Response Options

The IP Action column governs what action the IDP Sensor takes when it finds a matching condition.

Right-click the rulebase cell in the IP Action column and select Configure. The Configure IP Action dialog displays.

Configure your IP Action settings as appropriate for your network.

Setting Notification

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, you can also set an alert flag to appear in the log record.

To log an attack for a rule, right-click the Notification column of the rule and select **Configure**. The Configure Notification dialog box appears.

Setting Logging

In the Configure Notification dialog box, select **Logging** and then click **OK**. Each time the rule is matched, the IDP system creates a log record that appears in the Log Viewer.

You can choose to log an attack and create log records with attack information that you can view real-time in the Log Viewer. For more critical attacks, however, you might want to be notified immediately by e-mail, have IDP run a script in response to the attack, or set an alarm flag to appear in the log record. Your goal is to fine-tune the attack notifications in your security policy to your individual security needs.

Setting an Alert

In the Configure Notification dialog box, select **Alert** and then click **OK**. If Alert is selected and the rule is matched, IDP places an alert flag in the alert column of the Log Viewer for the matching log record.

Logging Packets

You can record the individual packets in the network traffic that matched a rule by capturing the packet data for the attack. Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack and its purpose, whether or not the attack was successful, and any possible damage to your network.



NOTE: To improve IDP performance, log only the packets after the attack.

If multiple rules with packet capture enabled match the same attack, IDP captures the maximum specified number of packets. For example, you configure Rule 1 to capture 10 packets before and after the attack, and Rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.



NOTE: Packet captures are restricted to 256 packets before and after the attack.

Setting Severity

You can override the inherent attack severity on a per-rule basis within the SYN Protector rulebase. You can set the severity to either Default, Info, Warning, Minor, Major, or Critical.

To change the severity for a rule, right-click the Severity column of the rule and select a severity.

Specifying VLANs

You can specify that the rule be applied only to packets from particular VLANs. See [“Setting VLAN Tags for IDP Rules” on page 524](#) for more information.

Setting Target Devices

For each rule in the rulebase, you can select the IDP-capable device that will use that rule to detect and prevent attacks. Alternatively, you can use Device Manager to assign policies to devices.

Entering Comments

You can enter notations about the rule in the Comments column. Anything you enter in the Comments column is not pushed to the target devices. To enter a comment, right-click the Comments column and select **Edit Comments**. The Edit Comments dialog box appears. You can enter up to 1024 characters in the Comments field.

Installing Security Policies ---

After you have successfully verified your security policy, you must:

1. Assign the policy to your managed devices.
2. Validate the policy.
3. Install the policy on your managed devices.

The following sections detail each step.

Assigning a Security Policy to a Device

New devices do not have an existing or default security policy. However, when you import a device configuration, NSM automatically imports all existing policies for the device. To simplify policy management, you can merge these multiple device policies into a single security policy that you install on several devices at one time. For details, see [“Merging Policies” on page 563](#).

After you have created a security policy, you must assign that policy to a device. Assigning a policy to a device links the device to that policy, enabling NSM to install the policy on that device. To assign an existing policy to a device, use one of the following methods:

- Right-click a device and select **Policy > Assign Policy**. Select the policy you want to assign to the device.
- Double-click a device to open the device configuration. In the Info tab, under Policy for device, select the policy you want to assign to the device.

You can use a single security policy to control multiple security devices. Each rule in a security policy contains an Install On column that specifies the devices the rule is applied

to. This means that you can assign a security policy to a device, but only some of the rules in that policy are actually installed on that device during a device update.

You can also create multiple policies for a single device, but only one security policy can be active on the device. When you update a device configuration, NSM installs the active policy on the security device. By default, NSM considers the active policy to be the policy that was most recently edited.



NOTE: If you delete and then re-import a device, you must reassign a policy to the device.

Validating Security Policies

You should validate a security policy to identify potential problems before you install it. NSM contains a Policy Validation tool to help you locate common problems, such as:

- Rule Duplication—Occurs when one or more rules in the security policy are identical. For more information, see [“Rule Duplication” on page 554](#).
- Zone Mismatch—Occurs when the source or destination zone you have chosen in a rule is not available on the device you selected in the Install column.
- Rule Shadowing—Occurs when a strict rule has no effect on traffic because it follows a broader ruler. For more information, see [“Rule Shadowing” on page 555](#).
- Unsupported Options—Occurs when a device in the Install column of a rule does not support a specific rule option configured for the rule. For details, see [“Unsupported Options” on page 555](#).

To use the Policy Validation tool to validate a security policy, you must first assign the security policy to a device. Then, to validate a policy, from the menu bar click **Devices > Policy > Validate Policy**. A Job Manager window displays job information and progress. Policy validation analyzes the source and destination addresses, the to and from zones, and the service when validating. If NSM identifies any problems in the policy during policy validation, it displays information about the problem at the bottom of the selected rulebase.



NOTE: We highly recommend that you validate a policy before installing it. A security policy that has internal problems can leave your network vulnerable.

Rule Duplication

Rule duplication occurs when an administrator configures the same rule in a rulebase more than once. Rule duplication can also occur during the rule validation process for devices running ScreenOS 5.0 and later. NSM treats each element of the rule as a separate rule. For example, when a rule with two service objects (AOL and DNS) is sent to the device, NSM sends it as two rules, one rule with AOL and another with DNS.



NOTE: For ScreenOS 5.0 and later, NSM sends rules with multiple objects or elements. For example, NSM can send a rule with two or more service objects as one rule.

You should delete all duplicate rules to maintain policy lookup efficiency.

A ScreenOS 5.0 and later device passes the policy validation process for HTTP; however, Rule 2 is not needed. To correct this problem, you should delete Rule 2.

Rule Shadowing

Rule shadowing occurs when an administrator selects or configures a policy in such a way that the next rules have no effect on traffic. Rule shadowing can introduce system vulnerabilities and packet dropping. Policy validation identifies rule shadowing. You should modify or delete all rules that overshadow others.

When a packet comes in, a security device compares it to the first rule in the policy. If a match occurs, the device executes the action associated with the rule. If no match occurs, the rule has no effect. Then, the device compares the packet to the next rule in the policy (unless the prior rule was a “terminal” rule.) So, each packet gets compared to every rule in the policy until a match occurs or a terminal rule ends the match process.

For example, if Rule 1 is a terminal rule, and a packet matches Rule 1, then the device will never compare the packet to the next rules. Or, if Rule 1 causes the packet to be dropped, and Rule 2 adds a diffserv marking, the diffserv marking will never be added.

In [Table 49 on page 555](#) Rule 1 shadows Rule 2. Rule 1 allows any service to a web server, but Rule 2 denies the service HTTP. When the security device receives a packet requesting HTTP service with the web server, Rule 1 allows the traffic. Rule 2 which denies HTTP is never checked.

Table 49: Rule Shadowing Example

Rule	From Zone	Source	To Zone	Destination	Service	Action
1	Untrust	Any	DMZ	Web server	Any	Allow
2	Untrust	Any	DMZ	Web server	HTTP	Deny

Unsupported Options

Policy Validation can also identify unsupported options in your security policy. Because different security devices and system support different features and options, policy validation checks the rules in the policy to ensure that the devices specified in the Install On column of the rule can support the Rule Options configured for the rule.

Some examples of unsupported option messages are included below:

- “Permit/Tunnel” Rules from home zone to work zone are not allowed on a Dial 2 device (except when NSRP Lite enabled).



NOTE: Because the “reject” firewall action is supported only by devices running ScreenOS 5.1 and higher, when NSM installs this rule on a device running an earlier OS, the action is automatically changed to “deny”.

- Schedule option is not supported on a vsys device.

For example, if you configure a firewall rule option (such as Antivirus protection or Deep Inspection) that is not supported by the security device in the Install column of the rule, policy validation displays an information message that describes the unsupported feature.

Installing New Security Policies

Before you install a new security policy, ensure that you have:

- Assigned the policy to your devices—After you have created a security policy, you must assign that policy to the devices you want to use that policy. Assigning a policy to a device links the device to that policy, enabling NSM to install the policy on that device.
- Selected the correct devices for the Install On column of each rule—A security device can only use one security policy at a time; when you install a new policy, it overwrites all existing policies on the security device.
- Configured each device in the Install On column of each rule correctly—When you push a policy to a device, you also push the device configuration to the device. Any changes made (by you or another administrator) to the device configuration are pushed to the device along with the policy.
- Configured rules in each rulebase correctly—The management system installs rules from all rulebases on the specified device. For information about rule installation and rule execution sequence, see [“Rule Execution Sequence” on page 477](#).
- Configured the VPN rules or VPN links in the policy correctly—The management system installs all VPN rules in the policy.

NSM does not validate VPN rules.

Additionally, to help you identify possible problems in your policy, you might want to run a Delta Config Summary before pushing the policy.

During policy installation, NSM installs the rules in the policy on the security devices you selected in the Install On column of each rule. The install process occurs between the management system and your managed devices. First, the GUI Server creates the ADM file that contains all policies for all devices selected for update (although the ADM file collects information from all policies, it does not merge the policies) The GUI Server sends the ADM to the Device Server. Next, the NSM Device Server receives the ADM and uses it to create a separate, individual DM for each device that you selected for update:

- For 5.0 and later devices, the Device Server sends the DM to the managed device, which translates the information in the DM into commands and runs those commands on the devices.

Configuring IDP Policy Push Timeout

IDP policies, due to their possibly large number of attack objects, may take a long time to upload and compile. The default timeout for IDP policy is 40 minutes, but you can set it higher if your policy uploads are timing out. Usually, this will only occur the first time a policy is pushed to a newly deployed Sensor.

To set the timeout to a higher value, edit the following file:

```
/usr/netscreen/DevSvr/var/devSvr.cfg
```

Change the following setting:

```
devSvrDirectiveHandler.idpPolicyPush.timeout 2400000
```

The setting is measured in milliseconds (1000's of a second). So, 2400000 milliseconds is equal to 40 minutes.

Updating Existing Security Policies

To install a new or modified policy on a managed device, from the toolbar, select **Devices > Configuration > Update Device Config**. If you changed the device configuration or assigned policy for a device, that device is automatically selected. Unselect any devices you do not want to update.

You can also enable session rematch for policy installations on managed devices running ScreenOS 5.1 and later. Session rematch enables NSM to preserve the existing sessions that are being tracked by the installed security policy during the policy update procedure. At the end of the update, NSM restores all valid sessions on the managed device and deletes all invalid sessions (a session is considered valid when the From Zone, Source, To Zone, Destination, and Service of the traffic is the same before and after the new policy installation).

You enable session rematch when you update devices (from the menu bar, select **Devices > Configuration > Update Device Config**). To enable session rematch from the Update Devices dialog box, select **Options**, then select **Rematch**, session treatment when modifying a policy rule, then click **OK**.



NOTE: You can also enable/disable session rematch in the system-wide device update settings. To configure, from the menu bar, select **Tools > Preferences > Device Update**. The system-wide setting (enabled or disabled) becomes the default setting for all device updates, but you can change the setting as needed for each individual update.

After you have selected the devices you want to update (and configured session rematch, if desired), click **OK** to begin the update process. The Job Manager dialog box appears and displays the progress of the policy installation. As the update is performed, the main display area of the Job Manager dialog box displays the CLI commands that the management system is sending to the physical device. In some cases, you might see that the policy is unset, then reset on the device.

NSM does not need to reset the policy when:

- The security policy you are installing does not exist on the physical device. The update installs the security policy on the device.
- The security policy you are installing already exists on the physical device. The update modifies the policy on the physical device, without resetting the policy.

NSM **must reset the policy** when the security policy you are installing already exists on the physical device, but an object within the policy has changed in NSM. The update first unsets the current policy on the device, deletes the old object, adds the new changed object, then installs the entire security policy again on the physical device.



NOTE: Additionally, NSM must reset the policy during an import when the security policy exists on the device, but does not exist in the management system.

After the update has completed, close the Job Manager window. The rules in the policy become active on the devices you selected in the Install On column of the rule. To see the exact rules that were applied to a specific device, in Device Manager, right-click a device and select **Policy > View Pending Device Policy**.

Updating Only the IDP Rulebases on ISG Devices

On ISG devices with IDP, you can elect to push only the IDP rulebases, not the entire policy.

To push only the IDP rulebases, not the firewall or multicast rulebases, select the **Update IDP Rulebase Only** check box in the Update Device Options dialog box.

The IDP-on-ISG rulebases are as follows:

- IDP
- Backdoor
- Exempt

Managing Rules and Policies

Managing rules and policies for multiple security devices can seem daunting at first. Take some time to carefully design your policies to make them efficient.

- [Helpful Tips on page 559](#)
- [Selecting Rules on page 559](#)
- [Editing Rule Order on page 560](#)
- [Using Cut, Copy, and Paste on Rules on page 560](#)
- [Deleting a Rule on page 562](#)
- [Disabling a Rule on page 562](#)
- [Using Rule Groups on page 562](#)
- [Reimporting Devices and Security Policies on page 562](#)
- [Merging Policies on page 563](#)
- [Importing SRX Series Devices That Contain Inactive Policies on page 564](#)
- [Exporting Policies on page 565](#)

Helpful Tips

Some helpful tips about managing your rules and policies:

- Because a device can have only one security policy installed at a time, you must include all rules for that device in one policy.
- The Policies navigation tree lists security policies alphabetically. You can create (or import) an unlimited number of security policies.
- Each security policy contains a default firewall rulebase (Zone); you can add other rulebases (Global, Multicast, IDP, Exempt, Backdoor) to create additional rules.
- Each rulebase can contain one or more rules, up to 40,000 max for the security policy. The top rule in the rulebase is rule 1, and second rule is rule 2, and so on. To combine rules for easier management within the Zone rulebase, you can create rule groups.
- Each rule group can contain one or more rules, up to 40,000 max for the security policy. Rules within a rule group follow the rulebase numbering sequence.
- The IDP, Exempt, or Backdoor rulebases are not included when you:
 - Merge two policies into a single policy
 - Import a security policy from an existing IDP-capable security device
- You cannot disable an entire security policy or a rulebase. You can, however, disable individual rules; for details, see [“Disabling a Rule” on page 562](#).
- When you reimport a device that was previously managed by NSM, you must manually reassign a policy to it. For information about reimporting issues, see [“Reimporting Devices and Security Policies” on page 562](#).

Selecting Rules

To select a single rule, click anywhere in the rule. The following sections explain these rule functions:

- [Editing Rule Order on page 560](#)
- [Using Cut, Copy, and Paste on Rules on page 560](#)
- [Using Cut, Copy, and Paste on Rule Fields on page 560](#)
- [Deleting a Rule on page 562](#)
- [Disabling a Rule on page 562](#)
- [Using Rule Groups on page 562](#)

Editing Rule Order

To change the order of rules in a policy, right-click the No. Column (the first column) of a rule and select **Move Rule Up** or **Move Rule Down**.

Using Cut, Copy, and Paste on Rules

To quickly create multiple rules that use the same basic information, copy and paste the rule, then change the parameters in each copied rule to make the rule unique (this is especially useful for rules that contain detailed rule options such as attack protection).



NOTE: When you cut and paste a rule, your preferred ID is retained. However, when you copy and paste a rule, a new ID is created.

To cut and paste a rule, right-click inside the No. column (the first column) of the rule and select **Edit > Cut**. Next, select a rule that is above or below the position you want to paste the cut rule into, then select **Edit > Paste > <above> <below>**.

Using Cut, Copy, and Paste on Rule Fields

You can cut, copy, and paste a column field in a rule to other column fields that have the same context. When you cut or copy a field, you can perform multiple paste operations. Cut, copy, and paste operations are available for all column fields except rule ID column fields. The cut operation is not available for generic column fields such as Notification and Action.

To perform cut, copy, and paste operations on rule fields:

1. From the policy rule you want to cut or copy from, right-click in the field and select **Edit > Cut** or **Edit > Copy**. When an element in the field is cut, it is replaced by either “any” or “default”, depending on the field.
2. From the policy rule field you want to update, right-click in the field and select **Edit > Paste**.

The field value you cut or copied is added in the field that received the paste operation. If an element is pasted into a field that specifies “any,” then “any” is deleted.

Cut, copy, and paste operations are not supported for the following policies:

- VPN policy rules
- Central rules on a regional server regular policy.

The following limitations apply to cut, copy and paste operations:

- Cut operations on fields of predefined policies are disallowed because predefined policies cannot be edited; however, copy operations are supported.
- There is no undo operation for cut, copy, or paste.
- You can copy but cannot cut fields that specify “any” or “default.”
- Cut, copy, and paste operations are supported on negate source or destination addresses, but only the address is cut or copied, not the negation. When pasting into a field, the existing negate attribute in the field takes precedence.

Dragging and Dropping Objects

Use the pull-down menu in the “Shared Objects For Policy” pane to easily select and add shared objects, including address, service, Global MIP, Global VIP, attack, device, VLAN, and custom field objects, to your security policies. Select the object and drag it into the appropriate policy column. When you drag objects beyond the visible rows or columns, the scroll bar will move horizontally or vertically, if there are more rules or columns available into which an object can be dropped.

From the main Address Tree and Service Tree, you can drag Address and Service objects into and out of groups.

Drag and drop support is also available in configuration dialogs for the following:

- Source and Destination columns of Zone-based and Global Firewall rulebases
- Source, Destination, and Attacks columns of IDP rulebase
- Source, Destination, and Application columns of APE rulebase
- Source, Destination, and Attacks columns of Exempt rulebase
- Source and Destination columns of Backdoor rulebase
- Source and Destination columns of Network Honeypot rulebase
- Source and Destination columns of Traffic Anomalies rulebase
- Source and Destination columns of SYN Protector rulebase
- Source and Destination columns of Permitted Object entries



NOTE: You cannot drag an object into a column that is not appropriate for that object. For example, you cannot drop a service object into the “Install On” column; you cannot drop a standalone IDP device into the “Install On” column for a zone-based firewall rulebase. Dragging and dropping objects is also not supported on any predefined IDP policy.

Deleting a Rule

To delete a rule, right-click inside the No. column (the first column) of the rule and select **Delete**. You can also delete a rule group; however, deleting the rule group also deletes all rules within the rule group.

Disabling a Rule

To disable a rule, right-click inside the No. column (the first column) of the rule and select **Disable**. The rule remains in the rulebase, but a gray diagonal stripe indicates that it has been disabled. While the rule is disabled, NSM does not install the rule on any devices.

To enable a rule, right-click inside the No. column (the first column) of the rule and select **Disable** again to clear the checkbox. You can disable rule groups using the same method.

Using Rule Groups

To create a rule group, select the rules you want to include in the group, then right-click and select create rule group. Enter a name and description for the rule group, then click **OK**.

Combining rules into a rule group can help you better manage rules. For example, you might want to create rule group for:

- VPN rules or VPN links
- Rules that manage traffic from a specific zone or interface on the security device
- Rules for a specific device or device group
- Rules that provide attack or AV protection
- Rules that manage VoIP traffic with GTP objects

You can add, edit, and delete rule groups; however, deleting a rule group also deletes all rules within that group. If necessary, you can also ungroup a rule group.

You can create multiple rule groups (40,000 rules maximum in a security policy). NSM supports one level of rule groups; you cannot create a rule group within a rule group.

Reimporting Devices and Security Policies

Occasionally, you might need to delete and then again add a security device to NSM. After you reimport the device configuration for a device that was previously managed by NSM:

- If you made no changes to the device policies using the WebUI or CLI, when you reimport the device, NSM does not create a new security policy.
- If you made changes to the devices policies using the WebUI or CLI, when you reimport the device, NSM creates a new security policy.

You must manually reassign a policy to a reimported device. For example, if you reimport a previously-managed security device, you might want to first merge the imported policy with a more comprehensive policy, then assign the comprehensive policy to the device.



NOTE: Importing the running configuration from a device completely overwrites all configuration information stored within NSM for that device. To help avoid accidental configuration overwriting, when you attempt to import a configuration from a currently managed security device, NSM prompts you for confirmation.

Merging Policies

When you import policies from a single managed device, those policies appear in NSM as rules in a new policy. Each device policy is imported as a single rule, and the rules make up the policy that exists on the device.



NOTE: In the ScreenOS WebUI and CLI, a security policy is a single statement that defines a source, destination, zone, direction, and service. In NSM, those same statements are known as rules, and a security policy is a collection of rules.

To simplify policy management and maintenance, you can merge two policies into a single security policy. To merge two policies, select a source policy and a target policy:

- The source policy contains the rules that you want to merge into another policy (in the UI, this is the From Policy).
- The target policy receives the rules from the source policy (in the UI, this is the To Policy).

NSM copies the rules from the source policy and pastes them above, below, or inline with the rules in the target policy. When placing rules inline, be aware of the intra-policy dependence of both policies. Because rule order is important (rules are executed top-down), rules can be dependent on other rules. If you rearrange the order of dependent rules by inserting merged rules, the security device changes the way it handles the packets. If you are unsure if you have intra-policy dependence in your rules, it's best to merge rules above or below the existing rules.

After creating a single security policy that contains both source and target rules, NSM also identifies rules that contain similar values in the source, destination, service, and install on columns, then collapses those rules into a single rule. NSM does not collapse rules that contain different zones, or rules that refer to unique VPNs.

By default, NSM also updates the device policy pointers to reference the new merged policy (the device policy pointer indicates which security policy is assigned to a device). When configuring Policy Merge settings, you can edit this option to keep the device policy pointers for both the source and target policies.

You can merge any two security policies. To access the Policy Merge tool, select the **Policies**, then use the menu bar to select **Tools > Policy Merge**. See the *NSM Online Help* for details.



NOTE: You can merge rules from 5.0 and later devices that use the deny action into rules from 5.1 and later devices that use the reject action, provided that the source, destination, source, and service are the same for the rules.

Policy A contains the rules as shown in [Figure 94 on page 564](#).

Figure 94: Security Policy A Rules (Before Policy Merge)

No.	ID	Match					Action	Install On
		From Zone	Source	To Zone	Destination	Service		
1	1	trust	1.1.1.1	untrust	2.2.2.2	FTP	permit	Boston
2	3	trust	3.3.3.3	untrust	4.4.4.4	HTTP	permit	Paris
3	2	trust	5.5.5.5	untrust	6.6.6.6	ICMP-ANY TELNET	permit	Bozeman

Policy B contains the rules as shown in [Figure 95 on page 564](#).

Figure 95: Security Policy B Rules (Before Policy Merge)

No.	ID	Match					Action	Install On
		From Zone	Source	To Zone	Destination	Service		
1	1	trust	1.1.1.1	untrust	2.2.2.2	FTP	permit	Bozeman Chicago
2	3	trust	3.3.3.3	untrust	4.4.4.4	HTTP HTTPS	permit	Boston Seoul
3	2	trust	5.5.5.5	untrust	6.6.6.6	ICMP-ANY TELNET	permit	Tokyo

To merge Policy A (from policy) with Policy B (to policy), from the file menu, select **Tools > Policy Merge Tool** and configure the merge.

NSM copies all rules from Policy A and pastes them above the rules in Policy B. Next, NSM merges the matching values in the columns to create a single, simplified policy (Policy C), shown in [Figure 96 on page 564](#).

Figure 96: Security Policy Rules (Merged from Policy A and Policy B)

No.	ID	Match					Action	Install On
		From Zone	Source	To Zone	Destination	Service		
1	1	trust	1.1.1.1	untrust	2.2.2.2	FTP	permit	Bozeman Chicago
2	3	trust	3.3.3.3	untrust	4.4.4.4	HTTP HTTPS	permit	Boston Seoul
3	2	trust	5.5.5.5	untrust	6.6.6.6	ICMP-ANY TELNET	permit	Tokyo

Importing SRX Series Devices That Contain Inactive Policies

When NSM imports an SRX Series device that contains an inactive policy, the inactive policy configuration is imported into the NSM device configuration table (device obj), as occurs in the case of in-device policy management. In addition, the inactive policies are

not displayed on the UI when the device is in central policy manager mode. All shared objects that are used in the inactive policies are imported into their respective shared objects table. Consequently, if the shared object (that is used in the inactive policy) is edited or deleted, the changes are updated to the device on next update.



NOTE: Re-importing an SRX Series device that contains inactive policies does not create duplicates of the shared objects that are included in the inactive policies in NSM.

Exporting Policies

You can export a security policy rulebase to an HTML file.

To export a security policy, select **File > Export Policy**. (You can also use the button or Alt-E.) In the **Export Policy** dialog box, select from the following options:

- Select all rulebases
- Expand rule groups
- Show expanded view
- Print filter condition
- Link all shared object details
- Run in background

Click the **Browse** button to select a default export directory for all future exports. Click **Export** to export the file. You can choose to export the policy as a background process by selecting the **Run in background** option. To check progress during a policy export, click **Export Policy Status** and view the completed percentage in a dialog

Each export creates a new directory. The default directory name is **<policyname>_YYMMDD_HHMMSS**. The export process puts each rulebase in a separate HTML file in that directory.

Use an HTML browser to view the exported file. Expanded views may make the output too wide for a standard printer. Shared objects like Address, Service, Install on, Attack columns and so on, appear as links. Click on a link to view details about the selected object in a new HTML page.

To export an expanded view of the Zone based Firewall Rules from a security policy, select a policy from **Policies**. Then select **File > Export Policy** from the menu bar. In the dialog box, select **Zone based Firewall Rules**. Select **Show Expanded View**. Browse to an export directory and click **Select Export Directory**. Click **Export**.

NSM creates a new subdirectory for each export. To view the policy, point your browser at file zone.html in the created subdirectory.

Automatic Policy Versioning

NSM policies are objects that reference other objects in the database. Those objects, in turn, reference other objects. When a version is created for a “versioning” you must also version all directly and indirectly referenced objects in the database to maintain referential integrity.

To simplify the process, you can generate a database snapshot, which is a virtual copy of the whole database at a specific point in time. However, you cannot change the database snapshot. At a later time when changes are made, you update an “object version” of the database.

This section explains how to set NSM for automatic policy versioning, create a new policy version, and view existing versions.

Setting NSM to Automatic Policy Versioning

This section explains how to use the GUI to make NSM default to automatic policy versioning.

To set the NSM default to policy versioning:

1. In the NSM GUI, select **Tools > Preferences**.
2. Under Object Versioning, check **Policy**.



NOTE: You can use this window to set NSM to automatically create versions after changes for devices and shared objects

3. Click **OK**.

NSM is set to automatically create a new version after each database change to the policy objects.

Viewing Existing Policy Versions

This section explains how to use the GUI to view existing policy versions.



NOTE: Some policies or shared objects must exist in NSM before you can view them.

To view existing policy versions:

1. In the NSM GUI, right-click on a policy.
2. In the NSM GUI, select **View Versions**.

The Version History window displays the version history for the selected policy. You can use this window to create a new version or work with existing versions. When you set NSM up for automatic policy versioning, a new version is created each time you save changes to this object.

3. Click **Close** to end your viewing session.

Creating a New Policy Version

This section explains how to use the GUI to create a new policy version.

Some policies or shared objects must exist before you can create a version. See [“Setting NSM to Automatic Policy Versioning” on page 566](#).

To create a new policy version:

1. In the NSM GUI, right-click on a policy.
2. Select **View Versions**.
3. Click **Create Version**.
4. Optionally, fill in Version Comments and click **OK**.

The new version appears in the list in the Version History window.

5. Click **Close** to save the changes.

Using a Filter to Search for a Policy Version

This section explains how to use a filter to search for a policy version.

To use a filter to search for a version:

1. In the NSM GUI, right-click on a policy.
2. In the popup menu, select **View Versions**.

The Version History window appears.

3. In the window, select the version and click **Filter/Search**.

The Version Filter Definition dialog box appears.

4. In the dialog box, fill in the filter criteria:
 - Comment Contains — Enter a comment on which to filter the version search.
 - Created After — Use the up and down arrows to select the date and time.
 - Created Before — Use the up and down arrows to select the date and time

5. Press **OK**.

The Version History now displays existing versions that meet the criteria.

6. When you are finished, press **Close**.

Editing Comments for an Existing Policy Version

This section explains how to edit comments for an existing policy version.

To edit comments for an existing version

1. In the NSM GUI, right-click on a policy.
2. In the popup menu, select **View Versions**.

The Version History window appears.

3. In the Version History window, select an existing version and press **Edit Comments**.

The Database Version dialog box displays the Version Comments.

4. Change the comments and click **OK**.

The Version History window displays the new comments.

5. When you are finished, click **Close**.

Comparing Two Versions

This section explains how to compare two existing versions.

When you import a device with assigned policy objects, and there is a difference between the existing and imported objects, NSM creates a new policy object that has two versions. The older version being the original and the newer version being the modified object. You can then use the compare versions tool to find out the changes to the object.

To compare two versions:

1. In the NSM GUI, right-click on a policy.
2. In the popup menu, select **View Versions**.

The Version History window appears.

3. Select two versions in the window.
4. Click **Compare** to view the differences.
5. When you are finished, click **Close**.

Restore an Older Version

This section explains how to restore an older version.

To restore an older version:

1. In the NSM GUI, right-click on a policy.

2. In the popup menu, select **View Versions**.

The Version History window appears.

3. Select an older (non-current) version in the window.

4. Click **Restore** to make the selected version current.

The Rollback Security Policy window appears.

5. Select an earlier version in the window and click **Next**.

A Diff window appears comparing the old and current version.

6. View the differences and click **Next**.

The Object Editor appears.

7. Make any necessary changes and click **Finish**.

The Version History shows the results of the version restore operation.

Viewing, Editing, Filtering, and Sorting Database Versions

This section explains how to view, edit, filter, and sort database versions.



NOTE: Beginning with the NSM 2008.1 release, the term “domain version” is replaced by “database version.”

More than one database version must exist before you can sort them.

To view, edit, filter, and sort versions

1. In the NSM GUI, select **Tools > Database Versions**.

All versions are listed in the popup Database Versions window.

2. To view, edit, filter, and sort versions:

- View — Select a listed version and click **Open**.

The popup window disappears and the selected version is displayed on the right in the main GUI window.

- Edit — Select a listed version and click **Edit Comments**.

In the popup Database Version dialog, enter changes in the Version Comments field and click **OK**.

The change appears in the Database Version list.

- Filter — Select a listed version and click **Filter/Search**.

In the popup dialog Database Version Filters, enter appropriate values in the fields listed below. Click **OK** to set the filter. You can search for existing filter settings by viewing the current settings in the Filter/Search fields.

- From DB Snapshot Id — Click the up arrow to increment the starting (from) database ID number in the applicable range. Click the down arrow to decrement the starting number. (default = none)
 - To DB Snapshot Id — Click the up arrow to increment the final (to) database ID number in the applicable range. Click the down arrow to decrement the final number. (default = none)
 - Comments Contains — You can enter partial text from the version comments in this field.
 - Create After — Click the up arrow the increment the start date for the applicable range. Click the down arrow to decrement the date.(default = none)
 - Created Before — Click the up arrow the increment the end date for the applicable range. Click the down arrow to decrement the date.(default = none)
 - Associated Object Type — If the database version is created as the result of a device update, this the pull-down menu for this filed shows all devices associated with that job. If the database version is create as a results of saving an object (such as a policy), the pull-down menu names that object. (default = Any)
- Sort — In the Database Versions window, click on any column header to sort the column by ascending or descending order.



NOTE: In large environments that has multiple administrators managing the devices, the NSM database grows very quickly due to domain versions or database versions. Include the following script in the guiSvr.cfg file to disable the domain version and improve system performance.

```
guiSvrManager.domainVersionDisabled yes
```

Displaying the Differences Between Database Versions

This section explains how to display the difference between two versions.

To display the differences:

1. In the NSM GUI, select **Tools > Database Versions**.
2. In the popup menu, select two databases.
3. Click **Compare**.

The Comparing — Database Snapshots window appears. The window lists changes for all user visible objects from older database versions to the current version. In the left pane, under Snapshot Diff Tree, the coloring over items indicates that the listed database version (SOSS, IDP-1, JS-1, and so on) is changed (yellow), added (green), or deleted (pink).

4. Click on a listed database to display detailed information about that version.

For more information about the types of displayed data, see [“Introduction to Network and Security Manager” on page 3](#) and the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

5. When you are finished reviewing data about the different versions, click **Close** on the window to return to the main GUI page.

Update Device with an Older Database Version

This section explains how to use the NSM GUI to update a single device with an older database snapshot. This action returns the device configuration to the earlier configuration associated with the older database.

To update the device:

1. Select **Configure > Device Manager > Devices**.
2. Under the Device Tree tab, right-click on a listed device.
3. In the popup menu, select **View Versions**.

The Version History window appears.

4. Select the older database version in the window and click **Update Device**.
5. When the Update Device prompt appears, click **Yes** to change the version.

The Job Information window appears indicating “Successful Completion.”

6. Click **Close** in the Job Information and Version History windows to end the operation.

If you right-click on the device, the Version History window lists the newly updated Current version and the previous version.

Pre and Post Rules

In NSM, a policy supports many kinds of rulebases. Each rulebase is an ordered list of rules. Prerule and postrule lists are also ordered lists of rules that are defined from the Central Manager at the global domain and subdomain levels as well as on regional servers in standalone NSM installations. You can define and apply rules for each rulebase type.

When you update a device, device-specific policy configurations are generated for the device. This creates rulebases by applying the following rules in the following order (from first to last):

- Prerules
- Policy rulebase rules
- Postrules

The prerules and postrules feature provides a policy definition at a domain level that can be applied to all devices within the specific domain and all subdomains. Users can define two sets of rules for any rulebase type that can be applied as prerules and postrules for any device of the given domain and subdomains.



NOTE: The Central Manager attack database version must match the regional server attack database version to push prerules and postrules.

Prerules and postrules are two sets of rules of any rulebase type that can be created for any domain. Configuration of pre/post rules are located in the main navigational tree under Policy Manager called Central Manager Policies. Domain Administrators can edit domain level policies from this option.

Prerules apply before any rules of a rulebase are applied to a device and post rules apply after any rules of a rulebase are applied to a device. Prerules and postrules in the integrated view are not editable. There is only one instance of pre/post rules for a specific domain.

Domain hierarchy is used when applying pre/post rules to subdomains. Within any subdomain, global domain pre rules take precedence over subdomain pre rules, which take precedence over Security policy specific rules. Similarly, Security policy rules take precedence over subdomain post rules, which take precedence over global domain post rules.



NOTE: You cannot push a pre/post rule from the central manager to a regional server.

All features of security policies are available for prerules and postrules.

- Import device command—Imports all rules into the security policy that is created for the device.
- Config summary—displays the prerules and postrules.
- View device pending policy—Displays the policy being pushed to a device including prerules and postrules from current and parent domains.
- Validate policy—Validates policy rules.
- View domain rules—When checked, any predefined or custom policy displays the prerules and postrules above and below the policy rules. These rules are displayed in a different color and not editable.

prerules and postrules can include rulegroups. The firewall rulebase for prerules and postrules cannot contain VPN rules or VPN links.

When the regional server pushes a rulebase to a device that is not contained within the regular policy, a warning message is displayed in the Job Manager window notifying the user that a rulebase was pushed that is not contained within the regular policy.

Rule Application Sequence

Since prerules and postrules are defined at the Central Manager, global, and subdomain levels, NSM imposes a rule application precedence. When all prerules and postrules are defined, the application order of rules in a rulebase are applied in the following order (from first to last):

- Central Manager pre rules
- Global domain pre rules
- Subdomain prerules
- Specific rulebase rules the device uses
- Subdomain postrules
- Global domain postrules
- Central Manager postrules



NOTE: By default, the prerules and postrules are limited to 250 rules each. The default range of prerules and postrules for each level is as follows:

Level	Prerule and postrule	Range
Central Manager	Prerule	999750 to 999999
	Postrule	999500 to 999749
Global domain	Prerule	999250 to 999499
	Postrule	999000 to 999249
Subdomain	Prerule	998750 to 998999
	Postrule	998500 to 998749

ScreenOS Devices

ScreenOS devices require rules to have unique IDs. Rules pushed to devices are the merged result of prerules and postrules based on pre/post policy and local policy from the device. Enforcing uniqueness at the single policy level is not sufficient.

With the Central Manager prerules and postrules, NSM enforces the uniqueness of a device rule's preferred ID server-wide. Therefore, when an administrator adds a domain level pre/post rule either from the regional server or from the Central Manager server pushing prerules and postrules to the regional server, the regional server generates a server-wide unique preferred ID for the new rule. There is a preset ID range for firewall rulebases.

Validation of prerules and postrules

In Central Manager servers, prerules and postrules are validated the same way as rules validated in NSM policy manager. Central Manager pushes prerules and postrules to the regional server and fills mapping tables with polymorphic objects. (See [“Polymorphic Objects” on page 576](#) for more details.) Invalid prerules and postrules in the regional server are removed when the policy is pushed to a device during the device update operation.

Install-On Column for prerules and postrules

In 2007.2 NSM Policy Manager, the Install-On column is the mechanism to specify which devices use a particular rule. While configuring a pre/post rule in Central Manager, rule application is applied at regional server level. The Install-On column, in this case, accepts only the Regional Server object or ANY as legal entries. When a Central Manager pushes a pre/post rule to a regional server, content in this column specifies which rule is pushed to which regional server.

Managing prerules and postrules

To manage post/pre rules, Central Manager administrators can:

- Add prerules and postrules
- Push prerules and postrules to Regional Server
- Modify prerules and postrules
- Delete prerules and postrules

Add prerules and postrules

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client.

To add a pre/post rule:

1. In the main navigation tree, select **Policy Manager > Central Manager Policies**.
2. Select either **Central Manager Pre Rules** or **Central Manager Post Rules**.
3. Click the Add icon in the toolbar and select **Add Rule**.
4. Select a regional server object for the rule's Install On column, as necessary.

Prerules and postrules can be added at the subdomain, global, or central manager level. Prerules and postrules use the precedence of central manager, global and then subdomain when applied to a policy.

Push prerules and postrules to Regional Server

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client, and a pre/post rule has been added.

To push a pre/post rule:

1. In the main navigation tree, select **Policy Manager > Central Manager Policies**.
2. Select either **Central Manager Pre Rules** or **Central Manager Post Rules**.
3. Select **Tools > Update Regional Servers**.
4. Select the regional servers to which you want to push prerules and postrules.

Central Manager Administrator monitors progress from the Job Manager.

Prerules and postrules and their referenced shared objects are replicated in the regional servers managed by Central Manager. The status and time of the prerules and postrules push is clearly marked when an administrator is logged onto a regional server.

Modify prerules and postrules

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client, and a pre/post rule has been pushed to a regional server.

To modify a pre/post rule:

1. In the main navigation tree, select **Policy Manager > Central Manager Policies**.
2. Select either **Central Manager Pre Rules** or **Central Manager Post Rules**.

3. Right-click the rule you want to modify and select **Copy, Paste, or Cut**. If you select **Paste**, you have additional options to paste the rule before or after another rule.

A modified pre/post rule replaces the existing pre/post rule on the regional server. Associated shared objects, if they are new, are replicated in the regional server.

Delete prerules and postrules

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client, and a pre/post rule has been pushed to a regional server.

To delete a pre/post rule:

1. In the main navigation tree, select **Policy Manager > Central Manager Policies**.
2. Select either **Central Manager Pre Rules** or **Central Manager Post Rules**.
3. Right-click the rule you want to modify and select **Delete**.

Associated shared objects (if they are not polymorphic objects), in the regional server, are also deleted from the regional server.

Polymorphic Objects

The Policy Manager uses shared objects (such as address, zone, and attack) when defining various components of a policy rule. Polymorphic objects are objects that can be defined at the Central Manager or regional server level. Polymorphic objects can be used as place holders for values that will be defined in a different context (in a regional server domain or subdomain, for instance).

Prerules and postrules are defined at the Central Manager level or regional server level and can use shared objects that are defined by regional administrators. To provide regional server administrators the capability of customizing Central Manager prerules and postrules, shared objects defined in Central Manager are flexible and can be customized by regional administrators, creating polymorphic objects.

Customizing Polymorphic Objects

Each polymorphic object contains a mapping table. Each entry of the mapping table has an attribute of domain, device, and a concrete shared object reference of the same type. You can customize a shared object by adding, deleting, or modifying an entry in the mapping table.

The regional server administrator can customize polymorphic objects by adding local, concrete shared objects to it. The mapping table shows only the current domain's entries. Therefore, if an administrator is in the global domain, no subdomain entries are visible.

This section contains the following topics:

- [Access Control of Polymorphic Object on page 577](#)
- [Validation of Polymorphic Object on page 577](#)
- [Supported Polymorphic Object Categories on page 577](#)

Access Control of Polymorphic Object

Table 50 on page 577 defines accessibility of polymorphic objects in different servers.

Table 50: Polymorphic Objects

	Polymorphic objects created and used in Central Manager Server	Polymorphic objects created in a Central Manager, but used in a regional server	Polymorphic objects created and used in a regional server
Create	Yes	No	Yes
Customize mapping table	No	Yes	Yes
Change name/color and other attributes	Yes	No	Yes
Delete	Yes	Yes if not referenced by central rules.	Yes

Validation of Polymorphic Object

When an administrator first creates a polymorphic object, the customization state is set to pending. The validation routine generates a warning/error message when encountering the pending state. Validation of a polymorphic object is triggered if an object is used by a rule and that rule is edited or viewed in the policy manager.

Supported Polymorphic Object Categories

Polymorphic objects are in the same category as concrete objects of the same nature. The shared object type attribute includes a new value for polymorphic objects of a specific category. The following objects categories can have polymorphic type: Address, Service, Zone, Global NAT and Routing Instance.

Zone is a polymorphic shared object at the Central Manager level. After a global pre/post rule is pushed, zone objects resolve into names in prerules and postrules. A vsys zone can only be supported with a polymorphic zone. Administrators must map every vsys manually with a vsys zone name.

Global MIP/DIP/VIP objects in current NSM are intrinsically polymorphic. For all other categories of shared objects at Central Manager level, only concrete shared objects are supported.

Manage Polymorphic Objects

You can create polymorphic objects at the Central Manager level or the regional server level. When polymorphic objects are created in the Central Manager they are pushed to one or more regional servers where they are available to be populated with real values.

The workflow for using polymorphic objects is:

- “Create a Polymorphic Object” on page 578

- [“Add a Polymorphic Object to a Pre/Post Rule” on page 578](#)
- [“Map a Polymorphic Object to a Real Value” on page 579](#)

Create a Polymorphic Object

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client.

To create a polymorphic object:

1. In the main navigation tree, select **Object Manager**.
2. Select one of the following objects that can be polymorphic objects:
 - Address
 - Service
 - Zone
 - Global NAT
 - Routing Instance
3. Click the Add icon in the toolbar and select **Polymorphic Address** to open the Add Polymorphic Address dialog box.
4. Enter the following information for the new polymorphic address, then click **OK**:
 - Name
 - Color (optional)
 - IP version- IPv4 or IPv6
 - Comment (optional)

NSM adds the polymorphic address object to the address tree.

Add a Polymorphic Object to a Pre/Post Rule

This procedure assumes that a Central Manager administrator is logged onto a Central Manager client.

To create a polymorphic object:

1. In the main navigation tree, select **Policy Manager > Central Manager Policies**.
2. Select either **Central Manager Pre Rules** or **Central Manager Post Rules**.
3. Click the Add icon in the toolbar and select **Polymorphic Address**.

Polymorphic objects have a color-coding of green and are ready to be used by Central Manager prerules and postrules.

4. Drag the polymorphic object from the Shared Objects section to the Destination column.

Map a Polymorphic Object to a Real Value

The following procedure assumes that a Central Manager administrator is logged onto a Central Manager client and a regional server object has been created.

To map a polymorphic object to a real value:

1. In the toolbar, click the **Login to Regional Server** drop down list.
2. Select one or more of the available regional servers and click **OK**. If you are prompted to save any changes, select the appropriate option.

The Job Manager provides a status window for pushing the polymorphic object to the selected regional servers.

3. In the main navigation tree of the regional server, select **Object Manager > Address Object** to show the polymorphic address objects pushed to this regional server.
4. Double-click the object you want to map to a real value.
5. Click the Add icon in the toolbar to open the New Address Map Entry dialog box.
6. Enter the values of the address object you want to map and click **OK** to map real values to the polymorphic object you pushed to the regional server.

Mapping Polymorphic Objects Before Importing or Updating Affected Devices

If any polymorphic objects (zone, address, or service) are contained in any pre/post rules defined by the Central Manager or regional server, mappings for all the polymorphic objects referenced on the regional server must be defined on the regional server before you import or update the affected devices. If an error message is returned on import or update indicating that a mapping for a polymorphic object was not defined, you can define a mapping for the polymorphic object listed in the error message, and import or update the device again.

CHAPTER 10

Configuring Voice Policies

Border Signaling Gateway (BSG) transaction policies determine how Junos OS handles VoIP signals. You can create transaction policies by defining rules which can contain an ordered list of terms. Every entry in a rulebase is a rule. Multiple transaction rules can be grouped together in an ordered list and named a rule set. You can associate rules or rule sets with service points in a gateway from the device configuration editor. A transaction rule set corresponds to the new-transaction-policy-set in the command-line interface (CLI) nomenclature.

While configuring BSG transaction rules, Policy Manager only displays shared objects that are applicable to transaction policies. You can copy, paste, drag and drop any of these shared objects into the transaction rule.

Juniper Networks M Series and MX Series routers running Junos 9.5 and later can be managed in two modes: Central Policy management (CPM) and In-Device management. You can create, view, edit and delete BSG transaction rulebases on M Series and MX Series routers only when you use the Central Policy management mode.

This chapter contains the following sections:

- [Adding a BSG Transaction Rulebase on page 581](#)
- [Adding Rules to the BSG Transaction Rulebase on page 582](#)

Adding a BSG Transaction Rulebase

You must add a BSG Transaction rulebase to a policy before you can add rules to it. To add a BSG Transaction rulebase:

1. In the navigation tree, select **Policies**.
2. Select the policy to which you want to add the BSG rulebase.
3. In the upper right corner of the **Policy**, click (+) and select **Voice Policies > Add BSG Transaction Rulebase**. The **BSG Transactions** tab appears in the policy.

Adding Rules to the BSG Transaction Rulebase

To add rules to a BSG Transaction rulebase:

1. Select **Policies** in the navigation tree.
2. Select a policy to which you want to add a BSG transaction rulebase.
3. Click (+) at the upper right of the **Policy** window.
4. Select **Voice Policies > Add BSG Transaction Rulebase**. The **BSG Transactions** tab appears in the policy window.
5. Click (+) at the upper right of the **Terms and Policies** window. Select **Add Policy** to add a new policy. Select **Add Term** to add a term to a rule in the policy.
6. Configure the rules in the policy. Right-click on each column to **Add**, **Edit**, **Filter**, **Delete**, or **Revert to template or default value**. Under the **From** header, set the following conditions to be matched:
 - **Src address** Enter the source address of the SIP request.
 - **Method** Select the method of request: invite, message, options, publish, refer, subscribe.
 - **Request URIs** The uniform resource identifier of the request source. Enter a regular expression.
 - **Contacts** Enter a regular expression.
7. Select the desired action for the rule under the **Then** header. The actions are: .
 - **Accept** Accept the traffic and send it to its destination.
 - **Reject** Do not accept the traffic and return a rejection message. You can log or sample rejected traffic.
 - **Trace** Trace messages are accepted by this policy.
 - **Route** Add or edit the route for the incoming request. You can set the egress service point or the next hop for the request.
8. View the admission controller applied to the policy in the **Admission Controller** column.
9. Right-click on the **Install On** column to select a target device on which to install the transaction policy.

10. View the applicable shared objects in the drop-down list in the **Shared Objects for Policy** section of the window. You can add, edit, delete and search for shared objects such as BSG Service Points and Admission Controllers.
11. Add, delete, edit and search for policy sets in **Policy Sets** section to the right of the policy window.

While creating transaction policies, the following restrictions apply:

- You can assign transaction rules to a service point in the device editor. For every association, NSM defines a rule in the service point's gateway and refers to it in the service point. If the policy assignment changes, you must reassociate the rules with the service points in the device editor.
- BSG Transaction rule merge and validation is not supported.
- Since NSM does not collect traffic logs for M Series and MX Series routers, the **Jump to Policy** option is not applicable from log reports.
- Admission-controller settings are dropped from the policies pushed to devices running Junos OS Releases earlier than 9.5.



NOTE: NSM 2009.1 and later releases support BSG transactions in the Policy Manager for M Series and MX Series only. On other devices, BSG transactions can be configured from the device editor.

The BSG transaction rulebase is excluded when you update devices that do not support transaction rules. NSM displays an appropriate warning during the update. NSM also provides an **Install On** option at the term, rule, and rule set levels. If a rule or a rule set is not applicable to the device being updated, NSM skips that rule or rule set.

Configuring Junos NAT Policies

Network address translation (NAT) policies determine how the network address information within packet headers gets translated. Either or both source and destination addresses in a packet may be translated. The translation can include IP addresses as well as port numbers.

The types of NAT policies that are supported on Juniper Networks devices are: Source NAT policy, Destination NAT policy, and Static NAT policy.

To configure a NAT policy on your network:

1. Select a policy.

2. Add a NAT rulebase to this policy.

A rulebase determines the overall direction of the traffic to be processed and consists of rule sets.

3. Add a rule set to this NAT rulebase.

A rule set consists of a general set of matching conditions for traffic. If the traffic matches these conditions, then that traffic is selected for NAT. A rule set can contain multiple rules.

4. Add a rule to this rule set.

Rules are part of a NAT rule set and specify the traffic to be matched and the action to be performed when the traffic matches the rule.

5. Push this policy on a device that performs the NAT.

This chapter contains the following sections:

- [Source NAT Policy on page 586](#)
- [Destination NAT Policy on page 590](#)
- [Static NAT Policy on page 593](#)

Source NAT Policy

Source NAT policy is used to allow hosts with private IP addresses to access a public network through the translation of the source IP address within a packet leaving the Juniper Networks device. For more information on source NAT, see <http://www.juniper.net/techpubs/software/junos-security/junos-security10.0/junos-security-swconfig-security/source-nat-config-overview-section.html#source-nat-config-overview-section>.

To configure a source NAT policy, the first step is to add a source NAT rulebase to this policy. For more information on adding a source NAT rulebase to a policy, see [“Adding a Source NAT Rulebase” on page 586](#).

Adding a Source NAT Rulebase

To add a source NAT rulebase:

1. From **Policy Manager**, select **Policies**.
2. Select the policy to which you want to add a source NAT rulebase. The **Security Policy** window appears.

If no policies are listed, then you must create a new policy and proceed to add a rulebase to it.
3. Click (+) at the upper right of the **Security Policy** window.
4. Select **NAT Policies > Add Source NAT Rulebase**. The **Source NAT** tab appears in the policy.

The next step is to add rule sets to this rulebase. For more information on adding a source NAT rule sets to the rulebase, see [“Adding a Rule Set to the Source NAT Rulebase” on page 586](#).

Adding a Rule Set to the Source NAT Rulebase

To add a rule set to the source NAT rulebase:

1. Click (+) at the upper left corner of the **Source NAT** tab.
2. Select **Add Rule Set** to add a new rule set. The **New Rule Set** dialog box appears.

Here, you must specify a unique name for the rule set and set the direction of the traffic to be processed by specifying the source and destination, as follows:
 - a. Enter a name in the **Name** field.
 - b. Select the source and destination from the **Source** and **Destination** drop-down lists, respectively. Here, you are identifying the traffic flowing from a specific source to a specific destination that must be address translated.

Source and destination can be one of the following:

- **Routing Instance**—Select the routing instance from the list.

In general, the list displays the routing instances configured within a specific device or just the shared routing instances depending on whether the **Select From Device** check box is selected (default) or not and can have the following values:

- The default routing instance (**default**), which ships with the device. You can use this routing instance, if you do not wish to configure anything new.
- Other routing instances, if you have added them previously. To add a new routing instance, use **Object Manager > Routing Instance Objects**.

- **Zone**—Select the zone from the list.

In general, the list displays the zones configured within a specific device or just the shared zones depending on whether the **Select From Device** check box is selected (default) or not and can have the following values:

- The default zone (**junos-global**), which ships with the device. You can use this zone, if you do not wish to configure anything new.
- Other zones, if you have added them previously. To add a new zone, use **Object Manager > Zone Objects**.

- **Interface**—Select the interface(s) from the list. The interfaces are listed only if you have imported the device to NSM. Hence, for modeled devices, no interfaces are listed.

c. Click **OK**.

A rule set with the specified name gets created and is displayed in the **Security Policy** window.

The next step is to add rules to the rule set. For more information, see [“Adding a Rule to a Source NAT Rule Set” on page 587](#).

Adding a Rule to a Source NAT Rule Set

To add a new rule to a rule set:

1. From the **Source NAT** tab, select the rule set to which you want to add the rule.
2. Click **(+)** at the upper left corner of the **Source NAT** tab.

3. Select **Add Rule** to add a new rule to the selected rule set. The **New Rule** dialog box appears.

Here, you must specify a unique name for the rule and set the conditions and the action to be performed when the traffic matches these conditions, as follows:

- a. Enter a name, which uniquely identifies the rule within the rule set.
- b. Select the source address from the list. This address represents the public IP address through which the traffic leaves the private network.
- c. Select the destination address from the list. The addresses listed here represent the hosts, which are in the public network.
- d. If using Port Address Translation (PAT), specify a port range (between 1024 and 65535) in the **Low** and **High** fields. When PAT is used, multiple hosts can share the same IP address. For more information on PAT, see <http://www.juniper.net/techpubs/software/junos-security/junos-security10.0/junos-security-swconfig-security/id-11012.html#id-11012>.
- e. Specify one of the following actions:
 - **Off**—Do not perform source NAT.
 - **Pool**—Use the specified user-defined address pool to perform source NAT. The values are listed only if have configured the pools, previously. For configuring a pool, see “[Configuring Source NAT Objects](#)” on page 455.

If you wish to configure persistent NAT, see “[Editing a Source NAT Rule or Rule Set](#)” on page 588.
 - **Interface**—Use the egress interface’s IP address to perform source NAT.
- f. Click **OK**.

A rule with the specified name gets created and is displayed in the **Security Policy** window.

Editing a Source NAT Rule or Rule Set

From the **Security Policy** window, select a row and then right-click on a column to edit the values. Depending on whether you have selected a rule or a rule set, you may get some or all of the following actions to perform:

- Under the **Name** header:
 - **Add Rule**—Enables you to add rules to the rule set from the **New Rule** dialog box. Specify the values and click **OK**.
 - **Add Source**—Enables you to view and modify the source that you set previously.
- Under the **From** header:

- **Zone/RJ/Interface > View/Modify Source**—Enables you to view and modify the source that you set previously.
- **Src Address:**
 - **Edit**—Enables you to cut, copy, and paste the values that are within this field.
 - **Add Src address**—Enables you to add additional sources.
 - **Edit Address**—Enables you to edit the address of the selected source.
 - **Delete Address**—Enables you to delete a source.
- Under the **To** header:
 - **Zone/RJ/Interface > View/Modify Destination**—Enables you to view and modify the destination that you set previously.
 - **Dest Address:**
 - **Edit**—Enables you to cut, copy, and paste the values that are within this field.
 - **Add Dest address**—Enables you to add additional destinations.
 - **Edit Address**—Enables you to edit the address of the selected destination.
 - **Delete Address**—Enables you to delete a destination.
 - **Dest Port:**
 - **Edit**—Enables you to copy and paste the values that are within this field.
 - **Edit Destination Port**—Enables you to edit the port range.
- Under the **Action** header: Select a row and right-click to perform the following actions:
 - **Edit Action**—Enables you to edit the action set previously. If you select a pool as the action, then you get the options to configure persistent NAT, as follows:
 - Select the **Persistent NAT** check box to implement this functionality.
 - **Inactivity Timeout**—Time, in seconds, that the persistent NAT binding remains in the device's memory when all the sessions of the binding entry are gone. When the configured time-out is reached, the binding is removed from memory.
 - **Max Session Number**—Maximum number of sessions with which a persistent NAT binding can be associated.
 - **Permit**—Can have one of the following values:
 - **Any remote host**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address.
 - **Target host**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address.
 - **Target host port**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address.

For more information on persistent NAT, see

<http://www.juniper.net/techpubs/software/junos-security/junos-security10.0/junos-security-swconfig-security/understand-persistent-nat-section.html#understand-persistent-nat-section>.

- Right-click on the **Install On** column to select a target device on which to install or remove the NAT policy.
- View the applicable shared objects in the drop-down list in the **Shared Objects for Policy** section of the window. You can add, edit, delete and search for shared objects, which are applicable to the specific NAT rulebase.

Destination NAT Policy

Destination NAT policy is used to allow hosts from public network to communicate with private network through the translation of the destination IP address within a packet that is entering the Juniper Networks device. For more information on destination NAT, see <http://www.juniper.net/techpubs/software/junos-security/junos-security10.0/junos-security-swconfig-security/jd0e90828.html#jd0e90837>.

To configure a destination NAT policy, the first step is to add a destination NAT rulebase to this policy. For more information on adding a destination NAT rulebase to a policy, see “[Adding a Destination NAT Rulebase](#)” on page 590.

Adding a Destination NAT Rulebase

To add a destination NAT rulebase:

1. From **Policy Manager**, select **Policies**.
2. Select the policy to which you want to add a destination NAT rulebase. The **Security Policy** window appears.

If no policies are listed, then you must create a new policy and proceed to add a rulebase to it.

3. Click (+) at the upper right of the **Security Policy** window.

Select **NAT Policies > Add Destination NAT Rulebase**. The **Destination NAT** tab appears in the policy.

The next step is to add rule sets to this rulebase. For more information on adding a destination NAT rule sets to the rulebase, see “[Adding a Rule Set to a Destination NAT Rulebase](#)” on page 590.

Adding a Rule Set to a Destination NAT Rulebase

To add a rule set to the destination NAT rulebase:

1. Click (+) at the upper left corner of the **Destination NAT** tab.

2. Select **Add Rule Set** to add a new rule set. The **New Rule Set** dialog box appears.

Here, you must specify a unique name for the rule set and specify the source through which the traffic enters the private network, as follows:

- a. Enter a name in the **Name** field.
- b. Select the source from the **Source** drop-down list.

Source can be one of the following:

- **Routing Instance**—Select the routing instance from the list.

In general, the list displays the routing instances configured within a specific device or just the shared routing instances depending on whether the **Select From Device** check box is selected (default) or not and can have the following values:

- The default routing instance (**default**), which ships with the device. You can use this routing instance, if you do not wish to configure anything new.
- Other routing instances, if you have added them previously. To add a new routing instance, use **Object Manager > Routing Instance Objects**.

- **Zone**—Select the zone from the list.

In general, the list displays the zones configured within a specific device or just the shared zones depending on whether the **Select From Device** check box is selected (default) or not and can have the following values:

- The default zone (**junos-global**), which ships with the device. You can use this zone, if you do not wish to configure anything new.
- Other zones, if you have added them previously. To add a new zone, use **Object Manager > Zone Objects**.

- **Interface**—Select the interface(s) from the list. The interfaces are listed only if you have imported the device to NSM. Hence, for modeled devices, no interfaces are listed.

- c. Click **OK**.

A rule set with the specified name gets created and is displayed in the **Security Policy** window.

The next step is to add rules to the rule set. For more information, see [“Adding a Rule to a Destination NAT Rule Set” on page 591](#).

Adding a Rule to a Destination NAT Rule Set

To add a new rule to a rule set:

1. From the **Destination NAT** tab, select the rule set to which you want to add the rule.

2. Click (+) at the upper left corner of the **Destination NAT** tab.
3. Select **Add Rule** to add a new rule to the selected rule set. The **New Rule** dialog box appears.

Here, you must specify a unique name for the rule and set the conditions and the action to be performed when the traffic matches these conditions, as follows:

- a. Enter a name, which uniquely identifies the rule within the rule set.
- b. Select the source address from the list. The addresses listed here represent the hosts that are in the public network.
- c. Select the destination address from the list. This address represents the public IP address through which the traffic enters the private network. You can also use multicast address group objects as the destination.
- d. Select a destination port. This is the port through which the traffic enters the private network.
- e. Specify one of the following actions:
 - **Off**—Do not perform destination NAT.
 - **Pool**—Use the specified user-defined address pool to perform destination NAT. The values are listed only if have configured the pools, previously. For configuring a pool, see [“Configuring Destination NAT Objects” on page 459](#).
- f. Click **OK**.

A rule with the specified name gets created and is displayed in the **Security Policy** window.

Editing a Destination NAT Rule or Rule Set

From the **Security Policy** window, select a row and then right-click on a column to edit the values. Depending on whether you have selected a rule or a rule set, you may get some or all of the following actions to perform:

- Under the **Name** header:
 - **Add Rule**—Enables you to add rules to the rule set from the **New Rule** dialog box. Specify the values and click **OK**.
 - **Add Destination**—Enables you to view and modify the destination that you set previously.
- Under the **From** header > **Zone/RJ/Interface** > **View/Modify Source**—Enables you to view and modify the source that you set previously.
- Under the **Match** header:
 - **Src Address**:

- **Edit**—Enables you to cut, copy, and paste the values that are within this field.
- **Add Src address**—Enables you to add additional sources.
- **Edit Address**—Enables you to edit the address of the selected source.
- **Delete Address**—Enables you to delete a source.
- **Dest Address > Select Destination Address**—Enables you to select a destination.
- **Dest Port:**
 - **Edit**—Enables you to copy, and paste the values that are within this field.
 - **Edit Destination Port**—Enables you to edit the destination port.
- Under the **Action** header, select a row and right-click to perform the following action:
 - **Edit Action**—Enables you to edit the action set previously.
- Right-click on the **Install On** column to select a target device on which to install or remove the NAT policy.
- View the applicable shared objects in the drop-down list in the **Shared Objects for Policy** section of the window. You can add, edit, delete and search for shared objects, which are applicable to the specific NAT rulebase.

Static NAT Policy

Static NAT facilitates address translation in both directions (incoming and outgoing) and the characteristics are:

- A unique public IP address must be allocated to each private IP address
- Address pool is not required as there is one to one mapping of the public to private IP addresses

Adding a Static NAT Rulebase

To add a static NAT rulebase:

1. From **Policy Manager**, select **Policies**.
2. Select the policy to which you want to add a static NAT rulebase. The **Security Policy** window appears.

If no policies are listed, then you must create a new policy and proceed to add a rulebase to it.

3. Click (+) at the upper right of the **Security Policy** window.

Select **NAT Policies > Add Static NAT Rulebase**. The **Static NAT** tab appears in the policy.

The next step is to add rule sets to this rulebase. For more information on adding a static NAT rule sets to the rulebase, see [“Adding a Rule Set to a Static NAT Rulebase” on page 594](#).

Adding a Rule Set to a Static NAT Rulebase

To add a rule set to the static NAT rulebase:

1. Click **(+)** at the upper left corner of the **Static NAT** tab.
2. Select **Add Rule Set** to add a new rule set. The **New Rule Set** dialog box appears.
Here, you must specify a unique name for the rule set and the source from where the traffic originates, as follows:
 - a. Enter a name in the **Name** field.
 - b. Select the source from the **Source** drop-down list.
Source can be one of the following:
 - **Routing Instance**—Select the routing instance from the list.
In general, the list displays the routing instances configured within a specific device or just the shared routing instances depending on whether the **Select From Device** check box is selected (default) or not and can have the following values:
 - The default routing instance (**default**), which ships with the device. You can use this routing instance, if you do not wish to configure anything new.
 - Other routing instances, if you have added them previously. To add a new routing instance, use **Object Manager > Routing Instance Objects**.
 - **Zone**—Select the zone from the list.
In general, the list displays the zones configured within a specific device or just the shared zones depending on whether the **Select From Device** check box is selected (default) or not and can have the following values:
 - The default zone (**junos-global**), which ships with the device. You can use this zone, if you do not wish to configure anything new.
 - Other zones, if you have added them previously. To add a new zone, use **Object Manager > Zone Objects**.
 - **Interface**—Select the interface(s) from the list. The interfaces are listed only if you have imported the device to NSM. Hence, for modeled devices, no interfaces are listed.
 - c. Click **OK**.

A rule set with the specified name gets created and is displayed in the **Security Policy** window.

The next step is to add rules to the rule set. For more information, see [“Adding a Rule to a Static NAT Rule Set” on page 595](#).

Adding a Rule to a Static NAT Rule Set

To add a new rule to a rule set:

1. From the **Static NAT** tab, select the rule set to which you want to add the rule.
2. Click (+) at the upper left corner of the **Static NAT** tab.
3. Select **Add Rule** to add a new rule to the selected rule set. The **New Rule** dialog box appears.

Here, you must specify a unique name for the rule and set the conditions and the action to be performed when the traffic matches these conditions, as follows:

- a. Enter a name, which uniquely identifies the rule within the rule set.
- b. Enter a destination address. The traffic from the source is routed to this destination.
- c. Enter an address prefix. Enter the IP address to which the source traffic must be translated to. As static NAT supports one to one mapping, if your source consists of a number of hosts, then make sure that you enter an equal number of public IP addresses in this field.
- d. Enter the routing instance. Enter the routing instance to which the IP addresses in the address prefix field (see previous row) are bound to. This field is optional.
- e. Click **OK**.

A rule with the specified name gets created and is displayed in the **Security Policy** window.

Editing a Static NAT Rule/Rule Set

From the **Security Policy** window, select a row and then right-click on a column to edit the values. You can perform the following actions:

- Under the **Name** header:
 - **Add Rule**—Enables you to add rules to the rule set. Right-clicking opens the **New Rule** dialog box. Specify the name, destination address, address prefix, and routing instance. Once you are satisfied with the values, click **OK**.
 - **Add Source**—Enables you to view and modify the source that you set previously.
- Under the **Zone/RJ/Interface** header > **View/Modify Source**—Enables you to view and modify the source that you set previously.
- Under the **Dest Addr** header > **Select Destination Address**—Edit the destination address.

- Under the **Action** header>Select a row and then right-click to > **Edit Action**—Enables you to edit the action set previously.
- Right-click on the **Install On** column to select a target device on which to install the NAT policy.
- View the applicable shared objects in the drop-down list in the **Shared Objects for Policy** section of the window. You can add, edit, delete and search for shared objects, which are applicable to the specific NAT rulebase.

Configuring VPNs

VPNs route private data through a public Internet. Like normal Internet traffic, data in a VPN is routed from source to destination using public Internet networking equipment. Unlike normal traffic, however, the source and destination use a Security Association (SA) pair to create a secure, private tunnel through which the data traverses the Internet. A tunnel has a defined start point and end point, (usually an IP address), and is a private connection through which the data can move freely. By encrypting and authenticating the data while in the tunnel, you can ensure the security and integrity of the data.

VPNs can also connect widely distributed networks to make separate networks appear as a single wide area network (WAN). VPNs replace costly Point-to-Point Protocol (PPP) and Frame Relay connections that require dedicated lines (and sometimes even satellites) between your private networks.

This chapter discusses the concepts involved in creating secure tunnels between devices, details the differences between VPN types, helps you determine the best VPN for your network, and guides you through creating and configuring your chosen VPN.



NOTE: For step-by-step instructions on creating VPNs, see the *NSM Online Help* topic “VPNs”.

- [About VPNs on page 598](#)
- [Planning for Your VPN on page 599](#)
- [Preparing VPN Components on page 608](#)
- [Creating VPNs with VPN Manager on page 614](#)
- [VPN Manager Examples on page 632](#)
- [Creating Device-Level VPNs on page 647](#)
- [Device-Level VPN Examples on page 662](#)
- [Auto-Connect Virtual Private Network on page 671](#)
- [IVE VPN Monitoring on page 673](#)

About VPNs

With Network and Security Manager (NSM), you can use basic networking principles and your Juniper Networks security devices to create VPNs that connect your headquarters with your branch offices and your remote users with your protected networks.

NSM supports tunnel and transport modes for AutoKey IKE, Manual Key, L2TP, and L2TP-over-AutoKey IKE VPNs in policy or route-based configurations. You can create the VPN at the system-level or device-level:

- System-Level VPN (VPN Manager)—Design a system level VPN and automatically set up connections, tunnels, and rules for all devices in the VPN.
- Device-Level VPN (Device Manager)—Manually configure VPN information for each security device, then add VPN rules to a security policy to create a policy-based VPN or configure routes on each security device to create a route-based VPNs.



NOTE: Each VPN that a device belongs to reduces the maximum number of templates by one. This includes VPNs configured in VPN Manager and VPNs configured at the device-level. You can apply a maximum of 63 templates to a single device.

Creating System-Level VPNs with VPN Manager

For AutoKey IKE and L2TP VPNs, create the VPN at the system-level using VPN Manager. VPN Manager supports:

- AutoKey IKE VPNs—In policy-based or route-based modes. You can also create a Mixed-Mode VPN to connect policy-based VPN members to route-based VPNs members in a single VPN.
- L2TP-over-AutoKey IKE RAS VPNs and L2TP RAS VPNs—Can connect and authenticate multiple L2TP remote access services (RAS) users and protected resources with or without encryption.
- Reusable VPN Components—Create objects to represent your protected resources, CA certificates and CRLs, custom IKE proposals, and NAT configurations, then use these objects in multiple VPNs.
- Compact and Expanded Views—Choose the Compact (default) or Expanded view to create your VPN. Both views offer the same configuration options.
- Autogenerated Tunnels—Create tunnel interfaces on each route-based VPNs member automatically. Use the device tunnel summary to review all autogenerated tunnels in the VPN.
- Autogenerated VPN Rules—Create all VPN rules with a single click. NSM automatically generates the rules between each policy-based VPN member. You can review these

rules, configure additional rule options (such as traffic shaping, attack protection, and logging), then insert the rules into a security policy.

- **Autogenerated VPN Routes**—Automatically add virtual router information using the VPN Manager for each device based on the routing type. Specify a routing type of topology to autogenerate a route for all VPN members based on the configured routing type (static or dynamic). This information changes the tunnel interface data and virtual router data for each device.

To view all VPNs created with VPN Manager, select VPN Manager in the navigation tree. A list of saved VPNs appears in the main display area in table format. You can add and delete VPNs from this view.

VPN Manager does not support Manual Key VPNs; to create a Manual Key VPN in NSM, you must create the VPN at the device-level in Device Manager.

Creating Device-Level VPNs in Device Manager

For Manual Key VPNs, create the VPN at the device-level by manually configuring VPN information for each security device.

After you have configured the VPN on each security device in the VPN, add VPN rules to a security policy to create the VPN tunnel (for policy-based VPNs) or to control traffic through the tunnel (for route-based VPNs).

You can also create AutoKey IKE, L2TP, and L2TP-over-AutoKey IKE VPNs at the device-level.

Supported VPN Configurations

NSM supports all possible VPN configurations that are supported by the CLI and Juniper Networks ScreenOS WebUI, including:

- **NAT-Traversal**—Because NAT obscures the IP address in some IPSec packet headers, VPN nodes cannot receive VPN traffic that passes through an external NAT device. To enable VPN traffic to traverse a NAT device, you can use NAT Traversal (NAT-T) to encapsulate the VPN packets in UDP. If a VPN node with NAT-T enabled detects an external NAT device, it checks every VPN packet to determine if NAT-T is necessary.
- **XAuth**—To authenticate remote access services (RAS) users, use XAuth to assign users an authentication token (such as SecureID) and to make TCP/IP settings (IP address, DNS server, and WINS server) for the peer gateway.

Planning for Your VPN

NSM offers you maximum flexibility for creating a VPN. You can choose your topology, authentication level, and creation method. Because you have so many choices, it's a good idea to determine what your needs are before you create the VPN so you can make the right decisions for your network.

These decisions include:

- VPN Topology—What do you want to connect? How many devices? How do you want these devices to communicate? Will you have users as VPN members?
- Data Protection—How much security do you need? Do you need encryption, authentication, or both? Is security more or less important than performance?
- Tunnel Type—Do you want an always-on connection or traffic-based connection?
- VPN Manager or Device-Level—How do you want to create the VPN? Maintain the VPN?

The following sections provide information to help you make these decisions.

Determining Your VPN Members and Topology

You can use a VPN to connect:

- Security devices—Create a VPN between two or more security devices to establish secure communication between separate networks.
- Network components—Create a VPN between a two or more network components to establish secure communication between specific machines.
- Remote users—Create a VPN between a user and a security device to enable secure access to protected networks.



NOTE: In NSM, remote users are known as remote access service (RAS) users.

Each device, component, and RAS user in a VPN is considered a VPN node. The VPN connects each node to other nodes using a VPN tunnel. VPN tunnel termination points are the end points of the tunnel; traffic enters and departs the VPN tunnel through these end points. Each tunnel has two termination points: a source and destination, which are the source and destination zones on security device.

Using Network Address Translation (NAT)

Network Address Translation (NAT) maps private IP addresses to public, Internet-routeable IP addresses. Because your security device is also a NAT server, you can use private, unregistered IP addresses for your internal network, minimizing the number of registered IP addresses you must buy and use.

If you enable NAT, when an internal system connects to the Internet, the security device translates the unregistered IP address in the outbound data packets to the registered address of the security device. The security device also relays responses back to the original system. Additionally, because your internal systems do not have a valid Internet IP address, your systems are invisible to the outside Internet, meaning that attackers cannot discover the IP addresses in use on your network.

Site-to-Site

Site-to-site VPNs are the most common type of VPN. Typically, each remote site is an individual security device or RAS user that connects to a central security device.

- Advantages—Simple, easy to configure.
- Disadvantages—The central security device is a single point of failure.

Use a site-to-site VPN to connect remote networks to a single, central network inexpensively.

Hub and Spoke

In a hub and spoke VPN, multiple security devices (spokes) communicate through a central device (the hub).

- Advantages—Can connect several devices and users. Hub and spoke VPNs are easy to maintain because you only need to reconfigure the spoke and the hub device, which save you administration and resource costs. If you have smaller security devices with limited tunnel capacity, you can use hub and spoke VPNs to increase the number of available tunnels.
- Disadvantages—The hub is a single point of failure; however, you can use NSRP for redundancy.

A hub acts as a concentrator for the other VPN members, but does not necessarily have resources that are available to other members. In fact, you can specify a security device that is not a VPN member to act as the hub: If you include the hub in the VPN, the hub device can send and receive traffic from all spokes; if you do not include the hub, the hub device routes traffic between spokes.

Use a hub and spoke topology when you want to route VPN traffic through a VPN member that does not contain protected resources.

Dual Hubs and Spokes

In VPNs running ScreenOS 6.3 and later, you can use Next Hop Resolution Protocol (NHRP) combined with IPSEC, to establish secure tunnels dynamically between spokes, achieving a scalable and easy-to-use VPN solution for distributed enterprises. You can configure a permanent static tunnel between dual hubs, while dynamic tunnels are set up and configured between spokes through the NHRP protocol, whenever the need arises.

In order to set up this system, the hub can be a device running ScreenOS 6.3 or earlier, but the spokes must be ScreenOS 6.3 or later devices.

Full Mesh

In a full mesh VPN, all VPN member can communicate with all other VPN members.

- Advantages—Because a full mesh configuration uses redundant IPSec tunnels, traffic continues to flow even if a node fails.
- Disadvantages—When you add a member to the VPN, you must reconfigure all devices.

Use a full mesh VPN when you need to ensure that every VPN member can communicate with every other VPN member.

Creating Redundancy

To ensure stable, continuous VPN connection, use redundant gateways to create multiple tunnels between resources. If a tunnel fails, the management system automatically reroutes traffic. Redundant gateways use NSRP to determine the tunnel status.

Protecting Data in the VPN

To protect traffic as it passes over the Internet, you can create a secure tunnel between devices using a tunneling protocol. Each device in the VPN uses the tunneling protocol to establish a secure data path, enabling traffic between the devices to flow securely from source to destination. NSM provides two tunneling protocols, IPSec and L2TP, as detailed in the following sections.

Using IPSec

IPSec is a suite of related protocols that tunnel data between devices and cryptographically secure communications at the network layer. Each device in the VPN has the same IPSec configuration, enabling traffic between the devices to flow securely from source to destination.

Because IPSec functions at the network layer, it protects all data generated by any application or protocol that uses IP. Network layer encryption protects data generated by all protocols at the upper layers of the protocol stack. It also protects all data throughout the entire journey of the packet. Data is encrypted at the source and remains encrypted until reaching its destination. Intermediate systems that transmit the packet (like routers and switches on the Internet) do not need to decrypt the packet to route it, and do not need to support IPSec.

When you create your VPN in NSM, you can use one or more IPSec services to establish the tunnel and protect your data. Typically, VPNs use encryption and authentication services to enable basic security between devices; however, for critical data paths, using certificates can greatly enhance the security of the VPN. NSM supports the following IPSec data protection services for VPNs.

Using Authentication

To authenticate the data in the VPN tunnel, you can use the AH protocol, pre-shared secrets, or certificates:

- **Authentication Header (AH)**—AH authenticates the integrity and authenticity of data in the VPN. You can authenticate packets using Message Digest version 5 (MD5), Secure Hash Algorithm-1 (SHA-1), or Hash-based Message Authentication Code (HMAC).
- **Preshared Secret**—NSM generates an ephemeral secret, distributes the secret to each VPN node, then authenticates the VPN data using MD5 or SHA hash algorithms against the secret.
- **Certificates**—IKE uses a trusted authority on the client as the certificate server. For details on using certificates, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Authentication only authenticates the data; it does not encrypt the data in the VPN. To ensure privacy, you must encrypt the data using ESP.

Using Encapsulating Security Payload (ESP)

ESP encrypts the data in the VPN with DES, Triple DES, or AES symmetric encryption. When the encrypted data arrives at the destination, the receiving device uses a *key* to decrypt the data. For additional security, you can encrypt the keys that decrypt the data using Diffie-Hellman asymmetric encryption. ESP can also authenticate data in the VPN using MD5 and SHA-1 algorithms. You can use ESP to encrypt, authenticate, or encrypt and authenticate data depending on your security requirements.



NOTE: We strongly recommend that you do not use null AH with ESP.

Because ESP uses keys to encrypt and decrypt data, each VPN node must have the correct key to send and receive VPN data through the VPN tunnel.

You can manually configure a key for each VPN node, or use a key exchange protocol to automate key generation and distribution:

- **Manual Key IKE**—In a manual key VPN, you specify the encryption algorithm, authentication algorithm, and the Security Parameter Index (SPI) for each VPN node. Because all security parameters are static and consistent, VPN nodes can send and receive data automatically, without negotiation.
- **Autokey IKE**—In an AutoKey IKE VPN, you can use the Internet Key Exchange (IKE) protocol to generate and distribute encryption keys and authentication algorithms to all VPN nodes. IKE automatically generates new encryption keys for the traffic on the network, and automatically replaces those keys when they expire. Because IKE generates keys automatically, you can give each key a short life span, making it expire before it can be broken. By also exchanging authentication algorithms, IKE can confirm that the communication in the VPN tunnel is secure.

Because all security parameters are dynamically assigned, VPN nodes must negotiate the exact set of security parameters that will be used to send and receive data to other VPN nodes. To enable negotiations, each VPN node contains a list of proposals; each proposal is a set of encryption keys and authentication algorithms. When a VPN node attempts to send data through the VPN tunnel, IKE compares the proposals from each VPN node and selects a proposal that is common to both nodes. If IKE cannot find a proposal that exists on both nodes, the connection is not established.

IKE negotiations include two phases:

- In Phase 1, two members establish a secure and authenticated communication channel.
- In Phase 2, two members negotiate Security Associations for services (such as IPSec) that require key material and/or parameters.

VPN nodes must use the same authentication and encryption algorithms to establish communication.

- **Replay protection**—In a replay attack, an attacker intercepts a series of legitimate packets and uses them to create a denial-of-service (DoS) against the packet destination or to gain entry to trusted networks. Replay protection enables your security devices to inspect every IPSec packet to see if the packet has been received before—if packets arrive outside a specified sequence range, the security device rejects them.

Using L2TP

Layer 2 Tunneling Protocol (L2TP) is another tunneling protocol used to transmit data securely across the Internet. Because L2TP can transport Point to Point Protocol (PPP) frames over IP, it is often used to:

- Establish PPP connections (Example: authenticate ADSL services using PPP for users with an ISP at the opposite side of a Telco IP/ATM network)
- Transmit non-IP protocols (Example: bridge Novell and other network protocols)

PPP can send IP datagrams over a serial link, and is often used to enable dial-up users to connect to their ISP and to the Internet. PPP authenticates username and password, and assigns parameters such as IP address, IP gateway, and DNS. PPP can also tunnel non-IP traffic across a serial link, such as Novell IPX or Appletalk.

PPP is also useful because it can carry non-IP traffic and authenticate connections to RADIUS servers. However, because PPP is not an IP protocol, Internet routers and switches cannot route PPP packets. To route PPP packets, you use L2TP, which encapsulates PPP packet inside an Internet routeable, UDP packet. L2TP VPNs supports remote access service users using Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) authentication.

Using L2TP Over AutoKey IKE

L2TP only transmits packets; for encryption, authentication, or other data protection services, you must further encapsulate the L2TP packet using AutoKey IKE.

Choosing a VPN Tunnel Type

You can configure three types of VPN tunnels with NSM:

- **Policy-based VPNs**—The VPN tunnel is created and maintained only during the transfer of network traffic that matches a VPN rule, and is torn down when the connection ends. Use policy-based VPNs when you want to encrypt and authenticate certain types of traffic between two VPN members.
- **Route-based VPNs**—The VPN tunnel is created when the route is defined and is maintained continuously. Use route-based VPNs when you want to encrypt and authenticate all traffic between two VPN members. You cannot add RAS users in a routing-mode VPN.
- **Mixed-mode VPNs**—Connects policy-based VPNs to route-based VPNs in a mixed-mode VPN. You cannot add RAS users in a mixed-mode VPN.

The following sections detail Policy-based and Route-based VPN types.

About Policy-Based VPNs

A policy-based VPN tunnels traffic between two security devices or between one security device and a remote user. Each time a security device detects traffic that matches the from zone, source, to zone, destination, and service in the VPN rule, it creates the VPN tunnel to encrypt, authenticate, and send the data to the specified destination. When no traffic matches the VPN rule, the firewall tears down the VPN tunnel.

To create a policy-based VPN, use NSM to configure a policy based on the network components you want to protect, including protected resources, then push the configuration to the security devices. The security devices use the configuration to create the VPN tunnel. A protected resource is a combination of a network component and a service; protected resources in a VPN can communicate with other protected resources using the specified services. In a VPN rule, you add protected resources as the source and destination IP addresses.

Policy-based VPNs can use any of the supported data protection methods. Use policy-based VPNs when you want to enable Remote Access Services (RAS). You can add users to the VPN just as you add devices, enabling user access to all resources within the VPN.

About Route-Based VPNs

Like a policy-based VPN, a route-based VPN tunnels traffic between two security devices or between one security device and a remote user. However, a route-based VPN automatically tunnels all traffic between two termination points, without regard for the type of traffic. Because the tunnel is an always-on connection between two network points, the security device views the tunnel as a static network resource through which to route traffic.

To create the termination points of the tunnel, you designate an interface on the security device as a tunnel interface, then define a static route or use a dynamic routing protocol (BGP, OSPF) between all tunnel interfaces in the VPN. The tunnel interface, just like a physical interface, maintains state to enable dynamic routing protocols to make route decisions. When using VPN Manager to create your route-based VPNs, the tunnel interfaces are automatically created for you.

VPN Checklist

After you have carefully considered your VPN requirements, create a VPN checklist to help you determine the VPN components you need to create. You might also want to create a network diagram of your topology that includes protected resources, VPN members, their IP addresses and gateways, and the type of tunnel between them.

Define Members and Topology

What do you want to connect?

- Devices
- Network Components/Protected Resources

- Remote Access Service (RAS) Users
- Extranet Devices

How do you want to connect the VPN members?

- Site to Site
- Hub and Spoke
- Dual Hubs and Spokes
- Full Mesh

You might want to create a network diagram to map out your VPN visually, with IP addresses, to help you configure your topology.

Define VPN Type: Policy-Based, Route-Based, or Mixed-Mode

What type of traffic do you want to protect?

- Use a policy-based VPN to encrypt and authenticate certain types of traffic between two network nodes.
- Use a route-based VPNs to encrypt and authenticate all traffic between two network nodes.
- Use a mixed-mode VPN to encrypt and authenticate traffic between policy-based and route-based VPNs nodes.

Define Security Protocol (Encryption and Authentication)

How do you want to protect the VPN traffic?

- Autokey IKE
- L2TP
- L2TP over AutoKey IKE
- Manual Key (you cannot use VPN Manager to create a Manual Key VPN)

You must also decide if you want to use certificates to authenticate communication between the VPN members.

Define Method: VPN Manager or Device-Level?

How do you want to create the tunnel? Using VPN Manager or configuring each device?

Using VPN Manager

When adding a VPN using the VPN Manager, you enter the VPN members, gateways, IKE properties, and VPN topology, then autogenerate the VPN rules that create the VPN. You can inspect the VPN rules and override any VPN property before sending the VPN configuration to your devices.

Choose the VPN type that best matches your VPN requirements:

- Autokey IKE VPN—Use to authenticate and encrypt traffic between devices and/or protected resources. An Autokey IKE VPN supports:
 - Mixed-mode VPNs (policy-based members and route-based members)
 - Policy-based VPNs
 - Route-based VPNs
 - ESP and AH Authentication
 - ESP AutoKey IKE Encryption
 - IP traffic
 - Tunnels between devices (routing-based) and protected resources (policy-based)
- Autokey IKE RAS VPN—Use to authenticate and encrypt traffic between remote users and protected resources. An Autokey IKE RAS VPN supports:
 - Policy-based VPNs
 - ESP and AH Authentication
 - ESP AutoKey IKE Encryption
 - IP traffic
 - Remote access users
- L2TP RAS VPN—Use to authenticate (but not encrypt) PPP or other non-IP traffic between RAS users and protected resources. An L2TP RAS VPN supports:
 - Policy-based VPNs
 - AH Authentication
 - PPP or other non-IP traffic
 - Remote access users
- L2TP over Autokey IKE RAS VPN—Use to authenticate and encrypt PPP traffic between remote users and protected resources. An L2TP over Autokey IKE RAS VPN supports:
 - Policy-based VPNs
 - ESP and AH Authentication
 - ESP AutoKey IKE Encryption
 - PPP or other non-IP traffic
 - Remote access users

Creating Device-Level VPNs

You can create the following VPN types:

- AutoKey IKE VPN
- Manual Key IKE VPN

- L2TP VPN
- Redundant Site-Site VPN

Preparing VPN Components

After you have determined how you want to configure your VPN, you can begin preparing the VPN components necessary to create the VPN. A VPN combines device-level components (such as devices, zones, and routes) with network-level components (authentication, users, and NAT) to create a secure system of communication. Before you can create a VPN, you must first configure the components that comprise the VPN.

Each VPN type has basic, required, and optional components:

- “Preparing Basic VPN Components” on page 608
- “Preparing Required Policy-Based VPN Components” on page 608
- “Configuring Required Routing-Based VPN Components” on page 611
- “Configuring Optional VPN Components” on page 613

For mixed-mode VPNs, you must configure all basic and required policy- and route-based components.



NOTE: For step-by-step instructions on creating VPNs, see the *NSM Online Help* topic “VPNs” .

Preparing Basic VPN Components

To create any type of VPN, ensure that all security devices you want to use in the VPN are managed by NSM and configured correctly.

- **Devices**—Add the security devices you want to include in the VPN to NSM, ensuring that all devices are in the same domain. If you need to add a device to a VPN in a different domain, you must add the device as an extranet device in the domain that contains the VPN, then add the extranet device to the VPN.
- **Zones**—Configure each security device with at least two zones (trust and untrust); each zone must contain at least one interface (physical or virtual).

Preparing Required Policy-Based VPN Components

A policy-based VPN requires several components:

- Address objects
- Protected resources
- NAT objects
- User objects

The following sections detail how to configure each component; after you have created a component, you can use it to create your VPN.

Configuring Address Objects

You must create address objects to represent your network components in the UI. For details on creating and configuring address objects.

Configuring Protected Resources

You should determine your protected resources first to help you identify the devices you need to include in the VPN. After you know what you want to protect, you can use VPN Manager or manually configure your security devices to create the VPN. A protected resource object represents the network components (address objects) and services (service objects) you want to protect and the security device that protects them.

The address specifies secured destination, the service specifies the type of traffic to be tunneled, and the device specifies where the VPN terminates (typically an outgoing interface in untrust zone). In a VPN rule, protected resources are the source and destination IP addresses.

When creating protected resources:

- To protect multiple network components that are accessible by the same security device, add the address objects that represent those network components to the protected resource object.
- To protect a single network component that is accessible by multiple security devices, add multiple devices to the protected resource object. You must configure each device to be a part of the VPN.
- To manage different services for the same network component, create multiple protected resource objects that use the same address object and security device but specify a different service object.
- If you change the security device that protects a resource, NSM removes the previous security device from all affected VPNs and adds the new security device. However, NSM does not configure the VPN topology for the new security device—you must reconfigure the topology to include the new device manually.

For more details on creating protected resources.

Configuring Shared NAT Objects

For VPNs that support policy-based NAT, you must create one or more shared NAT objects. A shared NAT object contains references to device-specific NAT objects, enabling multiple devices to share a single object.

First, create a device-specific NAT object by editing the device configuration of each security device member. Then, create a global NAT object that includes the device-specific NAT objects. In the Object Manager, create a single shared NAT object to represent similar device-specific NAT objects (for example, a global DIP represents multiple device-specific DIPs). Use the global NAT object in your VPN; when you install the VPN on a device, that device automatically replaces the shared NAT object with its device-specific NAT object.

For details on shared NAT objects.

Configuring Remote Access Service (RAS) Users

For VPNs that support RAS users, you must create a user object to represent each user. NSM supports two types of users:

- **Local Users**—A local user has an account on the security device that guards the protected resources in the VPN. When a local user attempts to connect to a protected resource, the security device authenticates the user.
- **External Users**—An external user has an account on RADIUS or SecureID Authentication Server. When an external user attempts to connect to a protected resource, the security device forwards the request to the authentication server for authentication.

Authenticating RAS Users

You can authenticate/encrypt a RAS user using one or more of the following protocols:

- **XAuth**—Uses IPSec ESP and a username and password for authentication. XAuth RAS users must authenticate with a username and password when they connect to the VPN tunnel.
- **AutoKey IKE**—Uses IPSec ESP and AH for encryption and authentication. AutoKey IKE users have a unique IKE ID that NSM uses to identify and authenticate the user during IKE Phase I negotiations. To simplify RAS management for large numbers of AutoKey IKE users, you can also create AutoKey IKE groups that use a shared Group IKE ID.



NOTE: We strongly recommend that you do not use null AH with ESP.

- **L2TP**—Uses Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) for authentication (password sent in the clear).
- **Manual Key IKE**—Uses IPSec ESP and AH for encryption and authentication. Because manual key users are device-specific, you create them in the security device configuration, not in the Object Manager. For details on creating manual key users, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

NSM allows certificate with DC in certificate DN to be used for dialup user IKE ID selection.

When you use certificate DN as dialup user IKE ID, the following takes place:

- On the device sever, a partial or whole DN is associated with a VPN configuration.
- On the client side, the certificate DN is sent as IKE ID for the server to match the VPN configuration based on the content of DN.

The server DN configuration can contain a container part and a wildcard part as follows:

- The container part contains a continuous section of the DN; for example, "OU=a,O=b". Any DN containing all specified elements in correct order are accepted.

- Up to seven wildcards can be specified, one for each of the following element: CN, OU, O, L, ST, C, Email.

NSM needs to support DC container type when using ASN1-DN to create IKE ID or a group of IKE ID that enables multiple, concurrent connections to the same VPN tunnel. During Phase 1 negotiations, IKE first attempts to make an exact match between the RAS IKE ID and peer gateway IKE ID.

If no match is found, IKE then attempts to make a partial match between the RAS IKE ID and Group IKE ID. When selecting this type, you must enter a container identity or a wildcard ID (CN, OU, O, L, ST, C, Email).

NSM devices authenticate a RAS IKE user's ID if the values in the RAS IKE user's ASN1-DN identity fields exactly match the values in the group IKE user's ASN1-DN identity fields. The container ID type supports multiple entries for each identity field (for example, "ou=eng,ou=sw,ou=screensos"). The ordering of the values in the identity fields of the two ASN1-DN strings must be identical. In this IKE ID matching part, we need to allow DC element to be matched.

NSM also supports DC in wildcard when using ASN1-DN to create IKE ID or a group of Wildcard ID.

NSM devices authenticate a RAS IKE user's ID if the values in the RAS IKE user's ASN1-DN identity fields match those in the group IKE user's ASN1-DN identity fields. The wildcard ID supports only one value per identity field (for example, "ou=eng" or "ou=sw", but not "ou=eng,ou=sw"). The ordering of the identity fields in the two ASN1-DN strings are inconsequential. In this IKE ID matching part, we need to support DC as a wildcard element.

Configuring Group IKE IDs

If your VPN includes multiple remote users, it can be impractical to create an IKE ID and VPN rule for each. Instead, you can use a Group IKE ID to authenticate multiple users in a single VPN rule. In the security device configuration VPN settings, create a VPN Group and specify the maximum number of concurrent connections that the group supports (cannot exceed the maximum number of allowed Phase 1 SAs or the maximum number of VPN tunnels allowed on the Juniper Networks security device platform).

For details on group IKE IDs, see the *Juniper Networks ScreenOS 5.x Concepts and Examples Guide*.

Configuring Required Routing-Based VPN Components

A route-based VPNs requires two components:

- Tunnel Interface or Zone
- Route (Static or Dynamic)

The following sections detail how to configure each required component.

For VPNs created with VPN Manager, you create the VPN first to autogenerate the tunnel interfaces, then create the routes on the device itself using those tunnel interfaces. For

VPNs created at the device level, you can create the tunnel interfaces and routes before or after configuring the VPN.

Configuring Tunnel Interfaces and Tunnel Zones

A VPN requires a physical or virtual interface on the security device, and each security device supports a specific number of physical and virtual interfaces. To support multiple VPNs on a device, you might want to create tunnel interfaces and tunnel zones to increase the number of available interfaces on the device.



NOTE: VPN Manager automatically creates the necessary tunnel interfaces for route-based VPNs. For device-level VPNs, you can create the tunnel interfaces before or after creating the VPN.

If you do not need to do address translation (NAT), use unnumbered.

- **Tunnel Interfaces**—A tunnel interface handles VPN traffic between the VPN tunnel and the protected resources. You can create numbered tunnel interfaces that use unique IP addresses and netmasks, or unnumbered tunnel interfaces that do not have their own IP address and netmask (unnumbered tunnel interface borrows the IP address of the default interface of the security zone).
- **Tunnel Zones**—A tunnel zone is a logical construction that includes one or more numbered tunnel interfaces. You must bind the VPN tunnel to the tunnel zone (not the numbered tunnel interfaces); the VPN tunnel uses the default interface for the tunnel zone. In a policy-based VPN, you can link:
 - A single VPN tunnel to multiple tunnel interfaces
 - Multiple VPN tunnels to a single tunnel interface

For details on tunnel interfaces and tunnel zones, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Configuring Static and Dynamic Routes

A security device must know the path, or route, between each protected resource or security device in the VPN before it can forward packets from the source network to the destination network on the other side of the tunnel. To specify the route, you can use static routes, which define a specific, unchanging path between two VPN nodes, or dynamic routes, which define an algorithm that dynamically determines the best path between two VPN nodes.



NOTE: If you are using VPN Manager to create the route-based VPNs, you create the routes after autogenerating the VPN. If you are creating a device-level VPN, you can create the routes after configuring the tunnel interfaces.

To create a static route, you must manually create a route for each tunnel on each device. For VPNs with more than just a few devices, Juniper Networks highly recommends using a dynamic routing protocol to automatically determine the best route for VPN traffic:

To route between different networks over the Internet, use Border Gateway Protocol (BGP); to route within the same network, use Open Shortest Path First (OSPF). For details on creating routes, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Configuring Optional VPN Components

In any type of VPN, you can also use three optional components:

- Authentication Server
- Certificate and Certificate Revocation List objects
- PKI Defaults

The following sections explain how to configure each optional component; after you have created the component, you can use it to create your VPN.

Creating Authentication Servers

To externally authenticate VPN traffic for XAuth and L2TP, you must create an authentication server object to use in your VPN.

Creating Certificate Objects

To authenticate external devices, use a Group IKE ID to authenticate multiple RAS users, or provide additional authentication for the security devices in your VPN, you must obtain and install a digital certificate on each VPN member. A digital certificate is an electronic means for verifying identity through the word of a trusted third party, known as a Certificate Authority (CA). The CA is a trusted partner of the VPN member using the digital certificate as well as the member receiving it.

The CA also issues certificates, often with a set time limit. If you do not renew the certificate before the time limit is reached, the CA considers the certificate inactive. A VPN member attempting to use an expired certificate is immediately detected (and rejected) by the CA.

To use certificates in your VPN, you must configure:

- Local Certificate—Use a local certificate for each security device that is a VPN member.
- Certificate Authority (CA) Object—Use a CA object to obtain a local and CA certificate.
- Certificate Revocation List (CRL) Object—Use a CRL object to ensure that expired certificates are not accepted; a CRL is optional.

Configuring Local Certificates

A local certificate validates the identity of the security device in a VPN tunnel connection. To get a local certificate for a device, you must prompt the device to generate a certificate request (includes public/private key pair request) using the Generate Certificate Request

directive. In response, the device provides certificate request that includes the encrypted public key for the device. Using this encrypted public key, you can contact a independent CA (or use your own internal CA, if available) to obtain a local device certificate file (a .cer file).

You must install this local certificate file on the managed device using NSM before you can use certificates to validate that device in your VPN. Because the local certificate is device-specific, you must use a unique local certificate for each device.

You can also use SCEP to configure the device to automatically obtain local certificate (and a CA certificate) from the CA directly. For details on local certificates, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Configuring CA Objects

A CA certificate validates the identity of the CA that issued the local device certificate. You can obtain a CA certificate file (.cer) from the CA that issued the local certification, then use this file to create a Certificate Authority object.

You must install this CA certificate on the managed device using NSM before you can use certificate to validate that device in your VPN. Because the CA certificate is an object, however, you can use the same CA for multiple devices, as long as those devices use local certificates that were issued by that CA.

You can also use SCEP to configure the device to automatically obtain a CA certificate at the same time it receives the local certificate.

Configuring CRL Objects

A Certificate Revocation List (CRL) identifies invalid certificates. You can obtain a CRL file (.crl) from the CA that issued the local certification and CA certificate for the device, then use this file to create a Certificate Revocation object.

You must install the CRL on the managed device using NSM before you can use a CRL to check for revoked certificates in your VPN. Because the CRL is an object, however, you can use the same CRL for multiple devices, as long as those devices use local and CA certificates that were issued by that CA.

After you have received a CRL list, you can use the CRL object in your VPN. For details on configuring a certificate revocation list object

Creating PKI Defaults

You can configure default PKI settings for each security device that define how that device handles certificates. When configuring a VPN that includes the device, you can use these default settings. For details on PKI defaults, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Creating VPNs with VPN Manager

Configuring a VPN using VPN Manager is an eight stage process:

- “Adding the VPN” on page 615
- “Configuring Members” on page 616 (policy-based, RAS users, routing-based)
- “Configuring Topology” on page 620 (AutoKey IKE only)
- “Configuring Gateways” on page 622
- “Configuring IKE” on page 626
- “Autogenerating VPN Rules” on page 628
- “Configuring Overrides” on page 629
- “Adding the VPN Link” on page 631



NOTE: For an L2TP RAS VPN, you do not need to configure gateways or IKE.

The following sections detail each step.



NOTE: For step-by-step instructions on creating VPNs, see the *NSM Online Help* topic “VPNs”.

Adding the VPN

1. From the menu bar, click **VPN Manager > New** and select the VPN type:
 - AutoKey IKE VPN—Use to connect devices and/or protected resources. An AutoKey IKE VPN supports mixed-mode, policy-based, and routing-based VPNs, but does not support RAS users.
 - AutoKey IKE RAS VPN—Use to connect IKE RAS users and protected resources. An Autokey IKE RAS VPN supports policy-based VPNs and IKE RAS users, but does not support routing-based VPNs, mixed-mode VPNs, or L2TP RAS users.
 - L2TP RAS VPN—Use to connect L2TP RAS users and protected resources without encryption.
 - L2TP over AutoKey IKE RAS VPN—Use to connect L2TP RAS users and protected resources. An L2TP over AutoKey IKE RAS VPN supports policy-based VPNs and L2TP RAS users, but does not support routing-based or mixed-mode VPNs.
2. Enter a name for the VPN, then specify the general properties for the VPN:
 - Enable—Use this option to enable/disable the VPN. If you disable the VPN, the autogenerated VPN rules, VPN member gateways, and other device configuration settings are not installed on your managed devices.
 - Termination Point—Select the Default Zone for the VPN Termination Point. Typically, the default zone is untrust. When you configure the topology for the VPN, you can select a unique termination point for each VPN member.

- View Properties—Configure the VPN components that the VPN Manager displays for the VPN:
 - Type (AutoKey IKE VPN Only). Select the components you want to configure for the VPN: Route-based components, Policy-based components, or both. By default, VPN Manager displays all Route- and Policy-based components for an AutoKey IKE VPN.
 - Dial Backup. When enabled, VPN Manager displays the dial backup option for route-based components (dial backup is supported only on NetScreen-5GT devices running ScreenOS 5.1 and later).
3. Click **OK** to save the VPN and return to VPN Manager.

Configuring Members

The second step in configuring your VPN is to add members to the VPN. Depending on the type of VPN you are creating, you can add protected resources, security devices, and/or RAS users as VPN members.

Adding Policy-Based Members

In policy-based configuration area, you can add protected resources to the VPN. Click **Protected Resources** link and select the predefined Protected Resources you want to include in the VPN.

After you have added the protected resources, you can configure NAT and/or L2TP settings on the security device that protects each resource:

- For L2TP RAS VPNs and L2TP over AutoKey IKE VPN protected resources, you must configure L2TP settings.
- For all protected resources, you can configure policy-based NAT. Use policy-based NAT to translate private source IP addresses to Internet-routeable IP addresses. Configuring NAT is optional; if you do not use NAT on your network, you do not need to configure NAT for the VPN.

The following sections detail how to configure NAT and L2TP.

Configuring NAT

Below the Protected Resources window, select **NAT** to display the protecting security devices for each protected resource. Select the device for which you want to configure NAT. Enable NAT and specify the following values (you cannot edit the name of the device or the zone that contains the protected resource).

- Configure Incoming DIP—You can enable the security device to use a Dynamic IP pool for incoming VPN traffic. For each incoming VPN packet, the device translates the destination address into a IP address that is selected from the DIP pool.
- Interface for Incoming DIP. Select the interface that receives traffic addressed to Dynamic IP addresses.

- Incoming Global DIP. Select the Global DIP object that represents range of IP addresses available to the security device. (This DIP pool must include IP addresses that are routeable on your internal network.)

For details on configuring DIP objects.

- Configure Tunnel Interface and Zone—You can bind the VPN tunnel to a tunnel interface or tunnel zone to increase the number of available interfaces in the security device.



NOTE: If the security device is running ScreenOS 5.x and configured in transparent mode, you can only configure the zone (the interface does not appear).

To use a tunnel interface and/or tunnel zone in your VPN, you must first create the tunnel interface or zone on the device; for details, see [“Configuring Tunnel Interfaces and Tunnel Zones” on page 612](#) and the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

- Tunnel Zone. Select a preconfigured tunnel zone on the security devices to bind the VPN tunnel directly to the tunnel zone. The tunnel zone must include one or more numbered tunnel interfaces; when the security devices route VPN traffic to the tunnel zone, the traffic uses one or more of the tunnel interfaces to reach the protected resources.
- Tunnel Interface. Select a preconfigured tunnel interface on the security devices to bind the VPN tunnel to the tunnel interface. The security devices route all VPN traffic through the tunnel interface to the protected resources.
- Configure MIP, VIP, and Outgoing DIP
 - Enable MIP. Enable MIP to use a mapped IP address for the interface.
 - Global MIP. Select the global MIP object that represents the mapped IP address you want to use for the interface.
 - Global VIP. Select the global VIP object that represents the virtual IP address you want to use for the interface.
 - Global DIP (Outgoing). You can enable the security device to use a Dynamic IP pool for outgoing VPN traffic. For each outgoing VPN packet, the device translates the source address into a IP address selected from the DIP pool. Select the Global DIP object that represents range of IP addresses available to the security device. (This DIP pool must include IP address that are routeable on the Internet.)

Configuring L2TP

For L2TP RAS VPNs and L2TP over AutoKey IKE VPN protected resources, you must configure L2TP settings.

To connect to an L2TP VPN tunnel, the L2TP RAS user uses the IP address and WINS/DNS information assigned by the user’s ISP. However, when the L2TP RAS user sends VPN traffic through the tunnel, the security device assigns a new IP address and WINS/DNS information that enables the traffic to reach the destination network.

Below the Protected Resources pane, select **L2TP/NAT** to display the protecting security devices for each protected resource. (If you are configuring an AutoKey IKE VPN or AutoKey IKE RAS VPN, this option does not appear.) Select the device for which you want to configure L2TP. In the L2TP tab, specify the following values (you cannot edit the name of the device).

- **Host Name**—Enter the name of the L2TP host.
- **Keep Alive**—The number of seconds a VPN member waits between sending hello packets to an L2TP RAS user.
- **Peer IP**—Enter the IP address of the L2TP peer.
- **Secret**—Enter the shared secret that authenticates communication in the L2TP tunnel.
- **Remote Settings**—Select the remote settings object that represents the DNS and WINS servers assigned to L2TP RAS users after they have connected to the tunnel.
- **IP Pool Name**—Select the IP pool object that represents the available IP addresses that can be assigned to L2TP RAS users after they have connected to the tunnel.
- **Auth Server**—Because the L2TP must authenticate L2TP users, use custom settings to associate those users with a specific Authentication Server. You can also configure the device to query the remote settings object for DNS and WINS information for those users.

To use the default authentication server for L2TP users, add the users to the device first.



NOTE: When configuring a VPN that includes RAS users, if you added the user as a L2TP local user and assigned an IP pool and remote settings object on a specific device in the VPN, those settings override the settings defined in the VPN.

Adding RAS Users

In the Remote User area, you can add RAS users to the VPN. (When configuring an AutoKey IKE VPN, this area does not appear.) Click the **Users** link to display the user selection dialog box, then click the Edit icon to select the predefined RAS users or user groups you want to include in the VPN. For details on creating RAS users and groups.

Defining a Default Gateway

You can include a single RAS user in multiple VPNs. To specify this VPN as the default entry point for all RAS users listed in the VPN, enable Use as Default Gateway.

Adding Routing-Based Members

In the routing-based configuration area, you can add routing-based members to the VPN. (When configuring an AutoKey IKE RAS VPN, an L2TP RAS VPN, or an L2TP over AutoKey IKE RAS VPN, this area does not appear.) A routing-based VPN member is a security device that will route traffic (statically or dynamically) through a tunnel interface to one or more VPN members.

VPN Manager automatically creates the necessary tunnel interfaces for each route-based VPNs member. However, after VPN Manager autogenerates the VPN tunnels, you must configure static or dynamic routes on the security devices to route traffic through these tunnel interfaces. For details on creating routes, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Click the security devices link to display the route-based member selection dialog box.

- **Configure Tunnel Interface Settings**—Select a Primary Zone, Secondary Zone, and physical source interface for each security device. The selected zone passes VPN traffic through the selected interface on the security device.
 - The Zone settings apply to all route-based members selected in the members window.
 - If the Primary Zone is not defined or available on the security device, VPN traffic automatically uses the Secondary Zone.
 - The Physical Source Interface is the default physical interface on the device that transmits VPN traffic.
- **Configure Tunnel Options**—ScreenOS 5.x and later devices support additional functionality for handling VPN tunnels:
 - To use a single tunnel interface on each device for VPN traffic, enable **Generate Single Tunnel Interface** for 5.x devices. When enabled, the security device uses the route table and the next-hop tunnel binding table to link a specific destination to one of a number of VPN tunnels bound to the same tunnel interface. By mapping the next-hop gateway IP address specified in the route table entry to a specific VPN tunnel in the NHTB table, the device can use one tunnel interface for all VPN traffic through the device. This option is enabled by default.
 - To create entries in the Next Hop Tunnel Binding (NHTB) table, enable **Generate NHTB entries** for 5.x devices. When this option is selected, VPN Manager autogenerates NHTB entries for each VPN tunnel.



NOTE: If you are using a single interface for all VPN traffic on the device but you do not select this option, you must manually add the NHTB routes in the NHTB table, or configure BGP to automatically create the entries for you. This option is disabled by default.

- Select **Dial Backup** to enable NetScreen-5GT security devices to use the serial port as a backup termination point for the VPN tunnel. When this option is enabled, VPN Manager automatically generates the termination point for the serial interface during VPN creation (you do not need to select the serial interface manually when configuring Termination Points).
- **Configure Members**—Click the Add icon to select the predefined security devices you want to include in the VPN. After you have added the device to the VPN, you can double-click the device and configure overrides for the default tunnel interface zone,

the physical source interface. For devices running ScreenOS 5.x and later, you can also enable/disable single tunnel interface and NHTB entries.

After VPN Manager generates the tunnel interfaces, you must configure static or dynamic routes on each VPN member to route traffic to other VPN members.

Configuring Topology

In the general configuration area, you can define the topology and/or termination points of the VPN:

- The topology of the VPN determines how VPN members *logically* connect to each other. The topology is the communication path that VPN traffic must take to reach a VPN member.
- The termination points of the VPN determine how VPN members *physically* connect to each other. A termination point is the interface on each VPN member that sends and receives VPN traffic to and from the VPN tunnel.



NOTE: If you change the security device that protects a resource, NSM removes the previous security device from all affected VPNs and adds the new security device. However, NSM does not configure the VPN topology for the new security device—you must reconfigure the topology to include the new device manually.

For AutoKey IKE VPNs, you must define the topology for the VPN. Each VPN member is a node that has specific connection capabilities, and the topology describes the logical connections between those nodes.

A node can be:

- Hub—A hub can connect to a branch or main.
- Main—A main can connect to a hub, branch, or another main. When configuring a VPN that uses multiple mains, you can select to mesh all mains (all mains can communicate with each other) or disable all main meshing.
- Branch—A branch can connect to a hub or a main. Branches can send and receive VPN traffic to and from a hub or a main device, but cannot communicate directly with other branches unless in a dual hub setup.

Additionally, you can use a *supernet* to reduce the number of rules required for the hub device in a policy-based VPN. A supernet is an address object group containing the network address objects that represent the source and destination points of the VPN. Use a supernet when the hub device supports a small number of rules.

Configuring Common VPN Topologies

You can use VPN Manager to configure the following common VPN topologies:

- **Hub and Spoke**—Select a device to act as the hub; this device connects VPN members and enables them to communicate. Next, select the VPN members to be the spokes. You are not required to use a VPN member as a hub:

- If do not select a VPN member as the VPN hub, the hub routes VPN traffic from one branch to another.
- If you do select a VPN member as the VPN hub, the hub routes VPN traffic from itself and all connected branches.

Each spoke can send and receive VPN traffic to and from the hub, but cannot communicate directly with other spokes.

- **Dual Hub and Spoke**—You can select a device to act as a backup hub, and enable the spokes to communicate with each other by making the following settings.
 - Assign a VPN and gateway. Edit the Topology settings from **VPN Manager > VPNs > AutoKey IKE VPN >, New or Edit > Topology**. Select **Enable Auto-Connect VPN**. You can assign the device to be used as backup, from the drop-down list in the **Backup Hub** field. After selection, the backup hub is added to the General Configuration list.
 - Edit the vrouter on the spoke device, and assign the ACVPN-Dynamic and NHS IP Address. You can set these parameters from **VPN Manager > VPNs > AutoKey IKE VPN > VPN > Device Tunnel Summary > Edit Router > Dynamic Routing Protocol > NHRP > Parameters**. You cannot make this setting on a hub device. The ACVPN-Dynamic and the ACVPN-Profile settings are mutually exclusive, so if a device is already set as a Hub, then you cannot set it as a Spoke or vice versa.
 - Assign NHRP redistribution rules. You can make this setting from the **VPN Manager > VPNs > AutoKey IKE VPN > VPN > Device Tunnel Summary > Edit Router > Dynamic Routing Protocol > NHRP > Redistribution Rules**.
 - Add the NHRP option to the OSPF, BGP, and RIP redistribution rules. You can make these settings from:
 - **VPN Manager > VPNs > AutoKey IKE VPN > VPN > Device Tunnel Summary > Edit Router > Dynamic Routing Protocol > OSPF > Redistribution Rules**.
 - **VPN Manager > VPNs > AutoKey IKE VPN > VPN > Device Tunnel Summary > Edit Router > Dynamic Routing Protocol > BGP > Redistribution Rules**.
 - **VPN Manager > VPNs > AutoKey IKE VPN > VPN > Device Tunnel Summary > Edit Router > Dynamic Routing Protocol > RIP > Redistribution Rules**.
 - Set the routing on the tunnel interface from “**VPN Manager > VPNs > AutoKey IKE VPN > VPN > Device Tunnel Summary > Edit Interface > General Properties**”. Select **Routing to ACVPN-Dynamic**.



NOTE: You can enable the dual hub feature only if the Spoke device runs ScreenOS 6.3 or later. The Hub device could run ScreenOS 6.3 or an older version.

- **Main and Branch**—Main and branch topologies combine the flexibility of hub and spoke with the redundancy of full mesh. Because you can select multiple mains, each branch has an alternate tunnel to use if one main fails. To create a main and branch:
 - Select the devices to act at mains; these devices can communicate with all other VPN members.
 - Select remaining devices as branches; these devices communicate with all mains.
- **Full Mesh**—Select all VPN members to act as mains. All members can communicate with any other VPN member. Do not select a hub.
- **Site to Site**—Select both VPN members as mains. Each member can communicate with the other VPN member. Do not select a hub.

Defining Termination Points

You must define the termination interface for each security device in the VPN. The Termination Points tab displays the default termination points for the VPN. A termination point is the interface on a security device that sends and receives VPN traffic to and from the VPN tunnel, and is typically in the Untrust zone. Each VPN member (the security devices included as routing-based members and/or as protected resources for policy-based members) has a default termination interface.



NOTE: You do not need to select the serial interface on a NetScreen-5GT security device to enable dial backup for the VPN tunnel. If you have enabled Dial Backup for the device in the Route-Based Configuration area, VPN Manager automatically generates the termination point for the serial interface during VPN creation.

To override the default termination interface, right-click the VPN member, select Edit, and select a new termination interface for the device.

Configuring Gateways

To configure the gateways for VPN, click the **Gateway Parameters** link.

Configuring Gateway Properties

In the Properties tab, specify the following gateway values.

Selecting a Mode

The mode determines how Phase 1 negotiations occur. Select the mode that meets your VPN requirements:

- **Main mode**—The IKE identity of each node is protected. Each node sends three two-way messages (six messages total); the first two messages negotiate encryption and authentication algorithms that protect subsequent messages, including the IKE identity exchange between the nodes. Depending on the speed of your network connection and the encryption and authentication algorithms you use, main mode negotiations can take a long time to complete. Use Main mode when security is more important.

- **Aggressive mode**—The IKE identity of each node is not protected. The initiating node sends two messages and the receiving node sends one (three messages total); all messages are sent in the clear, including the IKE identity exchange between the nodes. Because Aggressive mode is typically faster but less secure than Main mode, use Aggressive mode when speed is more important than security.

For RAS VPNs, you must use the Aggressive mode; for VPNs that do not include RAS users, select the mode that meets your requirements.

Configuring Heartbeats

Use heartbeats to enable redundant gateways.

- **Hello**—Enter the number of seconds the security devices wait between sending hello pulses.
- **Reconnect**—Enter the maximum number of seconds the security devices wait for a reply to the hello pulse.
- **Threshold**—Enter the number of seconds that the security devices wait before attempting to reconnect.

Configuring NAT Traversal

Because NAT obscures the IP address in some IPSec packet headers, VPN nodes cannot receive VPN traffic that passes through an external NAT device. To enable VPN traffic to traverse a NAT device, you can use NAT Traversal (NAT-T) to encapsulate the VPN packets in UDP. If a VPN node with NAT-T enabled detects an external NAT device, it checks every VPN packet to determine if NAT-T is necessary.

Because checking every packet impacts VPN performance, you should only use NAT Traversal for remote users that must connect to the VPN over an external NAT device. You do not need to enable NAT-T for your internal security device nodes that use NAT; each VPN node knows the correct address translations for VPN traffic and does not need to encapsulate the traffic.

To use NAT-T, enable NAT-Traversal and specify:

- **UDP Checksum**—A 2-byte value (calculated from the UDP header, footer, and other UDP message fields) that verifies packet integrity. You **must** enable this option for NAT devices that require UDP checksum verification; however, most NAT devices (including security devices) do not require it.
- **Keep alive Frequency**—The number of seconds a VPN node waits between sending empty UDP packets through the NAT device. A NAT device keeps translated IP addresses active only during traffic flow, and invalidates unused IP addresses. To ensure that the VPN tunnel remains open, you can configure the VPN node to send empty "keep alive" packets through the NAT device.

Configuring XAuth

Use the XAuth protocol to authenticate RAS users with an authentication token (such as SecureID) and to make TCP/IP settings (IP address, DNS server, and WINS server) for the peer gateway.

- **Default Server**—To use the default XAuthentication server for the device. To change or assign a default XAuthentication server, edit the VPN settings in the security device configuration.
- **XAuth Server**—Use when the remote gateway is a security device that you want to assign TCP/IP settings.
 - **Auth Server Name.** Select a preconfigured authentication server object.
 - **Allowed Authentication Type.** Select Generic or Challenge Handshake Authentication Protocol (CHAP) (password is sent in the clear) to authenticate the remote gateway.
 - **Query Remote Setting.** Enable this option to query the remote settings object for DNS and WINS information.



NOTE: When configuring a VPN that includes RAS users, if you added the user as a L2TP or XAuth local user and assigned a remote settings object on a specific device in the VPN, those settings override the settings defined in the VPN.

- **XAuth Client**—Use when the remote gateway is a RAS user that you want to authenticate.
 - **Allowed Authentication Type.** Select Any or CHAP.
 - **User Name and Password.** Enter the user name and password that the RAS user must provide for authentication.



NOTE: All passwords handled by NSM are case-sensitive.

- **Bypass Authentication** to permit VPN traffic from VPN members to pass unauthenticated by the XAuth server.

Configuring Gateway Security

Determine the authentication mechanisms you want the VPN nodes to use for IKE Phase I negotiations. You can use a preshared key or certificates for authentication.

Preshared Key/Certificate

For Phase 1, select a Preshared Key Information or PKI Information:

- **Preshared Key**—Use if your VPN includes security devices and/or RAS users. VPN nodes use the preshared key during Phase 1 negotiations to authenticate each other; because each node knows the key in advance, negotiations use fewer messages and are quicker.
 - To generate a random key, enter a value for the seed, then click **Generate Key**. NSM uses the seed value to generate a random key, which is used to authenticate VPN members.



NOTE: Using a random key can generate a key in excess of 255 characters, which exceeds ScreenOS limits and might not be accepted by the security device during update. To reduce the key size, shorten the autogenerated key value by deleting characters.

- To use a predefined value for the key, enter a value for the Preshared Key.
- PKI—Use if your VPN includes extranet devices or you require the additional security provided by certificates (PKI uses certificates for VPN member authentication). For details on creating and managing certificates.

For Phase 1, select a proposal or proposal set. You can select from predefined or user-defined proposals:

- To use a predefined proposal set, select one of the following:
 - Basic (*nopfs-esp-des-sha*, *nopfs-esp-des-md5*)
 - Compatible (*nopfs-esp-3des-sha*, *nopfs-esp-3des-md5*, *nopfs-esp-des-sha*, *nopfs-esp-des-md5*)
 - Standard (*gs-esp-3des-sha*, *gs-esp-aes128-sha*)



NOTE: You cannot use a predefined proposal set with certificates—you must select a user-defined proposal or change the authentication method to Preshared Key.

- To use a user-defined proposal, select a single proposal from the list of predefined and custom IKE Phase 1 Proposals. For details on custom IKE proposals.

If your VPN includes only security devices, you can specify one predefined or custom proposal that NSM propagates to all nodes in the VPN. If your VPN includes extranet devices, you should use multiple proposals to increase security and ensure compatibility.

Preshared Secrets

You can use the same preshared secret for all nodes in the VPN, or create a unique preshared secret for communication from a specific node to another node.

Configuring IKE IDs

Every VPN node has a unique identification number, known as an IKE ID. During Phase 1 negotiations, the IKE protocol uses the IKE ID to authenticate the VPN member.

VPN Manager automatically creates the default IKE ID for you, based on the policy- or route-based members and RAS users, so you do not need to configure this option. However, if you do not want to use the default IKE ID, you can select a different IKE ID type and configure an IKE ID for each VPN gateway.

The IKE ID tab displays all security devices included as routing-based members and/or as protected resources for policy-based members. For each device, select the IKE ID type and enter the ID value:

- **ASN1-DN**—Abstract Syntax Notation, version 1 is a data representation format that is non-platform specific; Distinguished Name is the name of the computer. Use ASN1-DN to create a Group IKE ID that enables multiple, concurrent connections to the same VPN tunnel; use a Group IKE ID to make configuring and maintaining your VPN quicker and easier.

For details on how Group IKE IDs work, see [“Configuring Group IKE IDs” on page 611](#). For details on determining the ASN1-DN container and wildcard values for Group IKE IDs, see the *Juniper Networks ScreenOS 5.x Concepts and Examples Guide*.

- **FQDN**—Use a Fully Qualified Domain Name when the gateway is a dynamic IP address. FQDN is a name that identifies (qualifies) a computer to the DNS protocol using the computer name and the domain name, for example, server1.colorado.mycompany.com.
- **IP Address**—Use an IP address when the gateway has a static IP address.
- **U-FQDN**—Use a User Fully Qualified Domain Name when the gateway is a dynamic IP address, such as a RAS user. A U-FQDN is an e-mail address. For example: user1@mycompany.com.

Configuring IKE

To configure the IKE properties and Phase 2 Proposals for the VPN, click the **IKE Parameters** link. Because L2TP RAS VPNs do not support encryption, you do not need to configure IKE properties for L2TP RAS VPNs.

IKE Properties

Configure the IKE properties:

- **Idle Time to Disable SA**—Configure the number of minutes before a session that has no traffic automatically disables the SA.
- **Replay Protection**—In a replay attack, an attacker intercepts a series of legitimate packets and uses them to create a denial-of-service (DoS) against the packet destination or to gain entry to trusted networks. If replay protection is enabled, your security devices inspect every IPSec packet to see if the packet has been received before—if packets arrive outside a specified sequence range, the security device rejects them.
- **IPSec Mode**—Configure the mode:
 - Use tunnel mode for IPSec. Before an IP packet enters the VPN tunnel, NSM encapsulates the packet in the payload of another IP packet and attaches a new IP header. This new IP packet can be authenticated, encrypted, or both.

- Use transport mode for L2TP-over-AutoKey IKE VPNs. NSM does not encapsulate the IP packet, meaning that the original IP header must remain in plaintext. However, the original IP packet can be authenticated, and the payload can be encrypted.
- Do not set Fragment Bit in the Outer Header—The Fragment Bit controls how the IP packet is fragmented when traveling across networks.
 - Clear. Use this option to enable IP packets to be fragmented.
 - Set. Use this option to ensure that IP packets are not fragmented.
 - Copy. Select to use the same option as specified in the internal IP header of the original packet.

Monitor

You can enable VPN Monitor and configure the monitoring parameters for the device. Monitoring is off by default. To enable the VPN Monitor in Realtime Monitor to display statistics for the VPN tunnel, configure the following:

- VPN Monitor—When enabled, the security devices in the VPN send ICMP echo requests (pings) through the tunnel at specified intervals (configurable in seconds) to monitor network connectivity (each device uses the IP address of the local outgoing interface as the source address and the IP address of the remote gateway as the destination address). If the ping activity indicates that the VPN monitoring status has changed, the device triggers an SNMP trap; the VPN Monitor (in RealTime Monitor) tracks these SNMP statistics for VPN traffic in the tunnel and displays the tunnel status.
- Rekey—When enabled, the security devices in the VPN regenerate the IKE key after a failed VPN tunnel attempts to reestablish itself. When disabled, each device monitors the tunnel only when the VPN passes user-generated traffic (instead of using device-generated ICMP echo requests). Use the rekey option to:
 - Enable dynamic routing protocols to learn routes and transmit messages through the tunnel.
 - Automatically populate the next-hop tunnel binding table (NHTB table) and the route table when multiple VPN tunnels are bound to a single tunnel interface.

For details on VPN monitoring at the device level, see the *Juniper Networks ScreenOS 5.x Concepts and Examples Guide*.

Differentiated Services Code Point Mark

If you want to set the Differentiated Services Code Point (DSCP) field of the IPSec IPv4 header to a specified value for each route-based VPN at the Phase2 configuration level, devices running ScreenOS 6.1 and later allow you to on both ASIC and non-ASIC platforms.

ScreenOS 6.1 and later support the DSCP value configuration for tunnel mode ESP packets only.

You cannot configure the DSCP setting if:

- The IPSec mode is transport.

- The IPSec Mode is tunnel but the binding interface is not a tunnel interface.

You can set the following DSCP Marks in the AutoKey IKE Parameters page:

- **DSCP Marking** — You can select either enable or disable. If the selected IPSec mode is transport, this option is automatically disabled.
- **DSCP Value** — Set the DSCP value in the range of 0–63. Mouse over the field to see the range of allowed values.

Configuring Security Level

For Phase 2 negotiations, select a proposal or proposal set. You can select from predefined or user-defined proposals:

- To use a predefined proposal set, select one of the following:
 - Basic (*nopfs-esp-des-sha, nopfs-esp-des-md5*)
 - Compatible (*nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, nopfs-esp-des-md5*)
 - Standard (*gs-esp-3des-sha, gs-esp-aes128-sha*)
- To use a user-defined proposal, select a single proposal from the list of predefined and custom IKE Phase 2 Proposals. For details on custom IKE proposals.

If your VPN includes only security devices, you can specify one predefined or custom proposal that NSM propagates to all nodes in the VPN. If your VPN includes extranet devices, you should use multiple proposals to increase security and ensure compatibility.

Autogenerating VPN Rules

When you have completed configuring the policy- and route-based VPNs members, the topology (if necessary) and termination points, and the IKE (if necessary) and gateway parameters for the VPN, you are ready to autogenerate the VPN.

During autogeneration, NSM generates the VPN rules that control traffic between policy-based VPN members, and edits the device configuration (gateways, security parameters, and so on) of each VPN member to support the VPN.

Autogeneration *does not*:

- Insert the VPN rules into a security policy. After you have reviewed the VPN rules and made any necessary overrides, you must manually insert the VPN rules (known as a *VPN link*) into a security policy. For details, see [“Adding the VPN Link” on page 631](#).
- Install the new VPN rules or edited device configurations on the managed devices in the VPN. After you have inserted the VPN link into a security policy, you can install that policy on your devices using the Updated directive.
- Create static or dynamic routes for route-based VPNs.

To autogenerate the VPN, click **Save**.

Configuring Overrides

The override area enables you configure individual settings for each VPN rules (for policy-based and mixed-mode VPNs) and each VPN member. Each change you make to the autogenerated rules or VPN member configuration is known as an *override* to the VPN settings.

You might need to override the VPN settings to:

- Configure additional security for specific tunnels.
- Configure additional authentication between specific VPN members.
- Configure unique monitoring or reporting options for specific VPN members or VPN tunnels.
- Configure unique IKE IDs for each VPN member.

Editing Policy Rules

For policy-based and mixed-mode VPNs, NSM automatically generates the VPN rules to control traffic between VPN members. To view these autogenerated rules, click the **Policy Rules** link in the Overrides area; the rules appear in a separate NSM window, using the same row and column format as in the Security Policies.



NOTE: Policy rules do not appear for route-based VPNs.

Changing Rule Position

The position of the rules indicates the order that they apply to traffic. To change the position of a rule, you can:

- Right-click the rule and select **Move Rule Up** or **Move Rule Down**, or
- Right-click the rule and select **Change Rule Position**. In the New Position dialog box, enter a new rule number for this rule. (The rule number is the first column in the policy table.)

Filtering Rules

You can also filter the VPN rules by zones using the Zone Filter in the upper right-hand corner of the VPN rule window. Select a zone in From Zone and/or the To Zone to order the rules as desired.

To save this rule order, click **Apply**.

Configuring Rule Options

You can configure rule options for each rule, including traffic shaping, logging, antivirus and attack objects, and protection actions. For details on configuring these options.

Editing Device Configuration

For all VPNs, you can edit the device configuration for each VPN member. The device configuration displays the interfaces, gateways, and other VPN configuration information for each individual device.

Overriding Interfaces

For route-based and mixed-mode VPNs, this displays the tunnel interfaces and virtual routers configured on the VPN member. To override the general properties and dynamic routing protocols for each tunnel interface, right-click the tunnel interface and configure the settings.



NOTE: The changes you make to a Virtual Router in the Overrides area apply to the device configuration, not just the VPN configuration. When you change a VR setting in VPN manager, that change is saved and applied to the device when you save and apply the VPN. Similarly, when you change a VR setting for the device configuration in Device Manager, that change is reflected in the VPN configurations that includes the device.

For policy-based VPNs, no tunnel interfaces appear.

Overriding AutoKey IKE VPN Settings

For VPNs that use AutoKey IKE, this displays the VPN name, remote gateway, and IPSec Mode for each tunnel in the VPN. To override the general properties, security, binding/proxyID, and monitoring option for each VPN tunnel, right-click the VPN name and configure the settings as desired.

Overriding Gateways

For all VPNs, this displays the gateway name, gateway mode, IP address, and IKE phase I proposals for each VPN gateway. To override the general properties, security, and IKE ID/XAuth options for each gateway, right-click the gateway name and configure the settings as desired.

Overriding VPN Groups

For all VPNs, this displays VPN groups.

Overriding L2TP Settings

For L2TP VPNs, this displays L2TP information for each VPN member. To edit this information, right-click **L2TP entry** and configure the settings as desired.

Viewing the Device Tunnel Summary

For route-based and mixed-mode VPNs, you can view the VPN tunnels between each route-based member, including the source and peer devices, the tunnel interface, zone, and physical interface.



NOTE: The device tunnel summary does not appear for policy-based VPNs.

You cannot edit the device tunnels from this view; to make overrides to the VPN tunnels, edit the interface configuration for each device.

Adding the VPN Link

After you have reviewed the autogenerated information and made any desired overrides to the VPN, you must update your managed devices to activate the VPN. By default, the VPN you created in VPN manager is installed as the first rule in the security policy for each managed device. However, the security policy does not display the VPN.

You can manually add a VPN link to your security policy; a VPN link creates a link between the security policy and VPN (the link points to the VPN rules that exist in the VPN in VPN Manager). You might want to add a VPN link so you can reposition it elsewhere in the policy, or to make the VPN viewable in your policy.

To create a VPN link, in security policies, select an existing security policy (or create a new security policy), then right-click and select **Add VPN** link. Select the VPN name and click **OK** to add the link to the policy. By default, the link appears at the top of the policy, but you can move the VPN link anywhere in the policy, just as you would a firewall rule.

If you make changes to the VPN or create overrides, the VPN link automatically updates to reflect those edits.

Editing VPNs

To edit a VPN created with VPN Manager:

1. In the navigation tree, select VPNs. A table listing all configured VPNs appears in the main display area.
2. Right-click the VPN you want to edit and select **Edit**. The expanded VPN view dialog box appears.
3. Make the necessary changes, then click **OK** to apply your changes.

To revert any changes you have made to the VPN, right-click the VPN name in the navigation tree and select **Revert Changes**.

Editing VPN Protected Resources

To edit a protected resource in the VPN, right-click the protected resource and select **Edit Protected Resource**. Make your changes, then click **OK** to save your changes.

If you make changes to a protected resource object that is used in a VPN, NSM automatically generates new configuration and propagates your changes to all affected security devices. If you change the security device that protects a resource, NSM removes the previous security device from all affected VPNs and adds the new security device.

However, NSM does not configure the VPN topology for the new security device—you must reconfigure the topology to include the new device manually.

Editing Users

To edit a user object in the VPN, right-click the user and select **Edit Remote User**. Make your changes, then click **OK** to save your changes.

Editing the VPN Configuration

To add or delete a member, edit any VPN parameter, or reconfigure the VPN topology, select the VPN and click **OK**. Make your changes, then click **Save** to regenerate the VPN.



NOTE: After you click **Save**, you cannot revert your changes to a VPN.

Editing VPN Overrides

If you add, edit, or delete an override, the VPN link automatically updates the autogenerated rules to reflect those edits.

VPN Manager Examples

This section provides examples of common VPN types:

- “[Example: Configuring an Autokey IKE, Policy-Based Site-to-Site VPN](#)” on page 632
- “[Example: Configuring an Autokey IKE RAS, Policy-Based VPN](#)” on page 637
- “[Example: Configuring an Autokey IKE, Route-Based Site-to-Site VPN](#)” on page 640
- “[Example: Configuring XAuth Authentication with External User Group](#)” on page 643

The following sections provide step-by-step instructions on creating each VPN type.



NOTE: For examples on creating a Manual Key VPN, see “[Device-Level VPN Examples](#)” on page 662.

Example: Configuring an Autokey IKE, Policy-Based Site-to-Site VPN

An AutoKey IKE VPN connects protected resources using AutoKey IKE. Use this VPN type to connect and control traffic between two security devices. In this example, an AutoKey IKE tunnel using a pair of certificates (one at each end of the tunnel) provides the secure connection between the Tokyo and Paris offices. For the Phase 1 and 2 security levels, you specify the Phase 1 proposal as `rsa-g2-3des-sha` and select the predefined “Compatible” set of proposals for Phase 2. It is assumed that both participants already have RSA certificates and are using Entrust as the certificate authority (CA). All zones are in the `trust-vr`.

1. Configure security devices.
 - a. Configure the Tokyo device with the following interfaces:

- Ethernet1 is the Trust IP (10.1.1.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (1.1.1.1/24) in the Untrust zone.
- b. Configure the Paris device with the following interfaces:
- Ethernet1 is the Trust IP (10.2.2.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (2.2.2.2/24) in the Untrust zone.
2. Create the address objects that you will use to create Protected Resources (for details on creating or editing address objects,. If you imported a security device, the address book objects configured on that device are automatically imported as Address objects into the NSM UI.
- a. Add the Tokyo Trust LAN (10.1.1.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
- For Name, enter **Tokyo Trust LAN**.
 - For IP Address/Netmask, enter **10.1.1.0/24**.
 - For Color, select **magenta**.
 - For Comment, enter **Tokyo Trust Zone**.
- b. Add the Paris Trust LAN (10.2.2.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
- For Name, enter **Paris Trust LAN**.
 - For IP Address/Netmask, enter **10.2.2.0/24**.
 - For Color, select **magenta**.
 - For Comment, enter **Paris Trust Zone**.
3. Create the Tokyo Protected Resources object. In Protected Resources (under VPN Manager), click the Add icon. Configure as shown in [Figure 97 on page 634](#), then click **OK**:

Figure 97: Create Tokyo Protected Resource Object for AutoKey IKE VPN

The screenshot shows the 'Protected Resource' dialog box with the following configuration:

- Name: Tokyo Protected Resources
- Color: green
- Service Object: any
- Server/Client: Both
- Network Object: Tokyo Trust LAN
- Comment: (empty)

Below the fields is a table with two columns: 'Security Gateway Device' and 'Zone'.

Security Gateway Device	Zone
Tokyo	trust

At the bottom are buttons for OK, Cancel, and Apply.

4. Create the Paris Protected Resources object. In Protected Resources (under VPN Manager), click the Add icon. Configure, then click **OK**:

Figure 98: Create Paris Protected Resource Object for AutoKey IKE VPN

The screenshot shows the 'Protected Resource' dialog box with the following configuration:

- Name: Paris Protected Resources
- Color: green
- Service Object: any
- Server/Client: Both
- Network Object: Paris Trust LAN
- Comment: (empty)

Below the fields is a table with two columns: 'Security Gateway Device' and 'Zone'.

Security Gateway Device	Zone
Paris	trust

At the bottom are buttons for OK, Cancel, and Apply.

5. Create the VPN. In the navigation tree, double-click **VPN Manager**, then right-click **VPNs** and select **AutoKey IKE VPN**. The New AutoKey IKE VPN dialog box appears. Configure the General VPN Properties:
 - a. In Name, enter **Tokyo-Paris Policy-Based VPN**.
 - b. Select **Enable**.
 - c. In Termination Point, select **Untrust**.

- d. For VPN Type, select **Policy-Based**.
 - e. Click **OK** to save the VPN and return to VPN Manager.
 - f. In VPN Manager, select the **Tokyo-Paris Policy-Based VPN**. The VPN appears in the main display area.
6. Configure the policy-based members:
 - a. Select the **Protected Resources** link to display the Protected Resources list.
 - b. Select the **Paris Protected Resources** and the **Tokyo Protected Resources**.
 - c. Click **OK** to return to the main display area.
 7. Configure the VPN topology:
 - a. Select the **Topology** link to display the Topology dialog box.
 - b. Click the Add icon to display the Topology configuration dialog box. Configure the following:
 - For Hub and Supernet, leave the default of none.
 - Enable Mesh Main(s).
 - In the Mains window, select the Paris and Tokyo security devices.
 - c. Click **OK** to return to the Topology dialog box, then click **OK** to return to the main display area.
 8. Configure the termination points of the VPN:
 - a. Click the **Termination Points** link. The Termination Points dialog box appears.
 - b. Confirm that both Paris and Tokyo devices use a Termination Interface of ethernet3.
 - c. Click **OK** to return to the main display area.
 9. Configure the VPN gateway:
 - Click the **Gateway Parameters** link. The Properties tab appears. Leave all defaults and click the Security tab.
 - In the Security tab, configure the PKI Information and Phase 1 Proposals as shown in [Figure 99 on page 636](#).

Figure 99: Configure Gateway Parameters for AutoKey IKE VPN

- Click **Save** to save your configuration changes to the VPN.

To view the autogenerated rules, click the **Policy Rules** link in the Overrides section. VPN Manager generates the rules as shown in [Figure 100 on page 636](#).

Figure 100: View Autogenerated Rules for AutoKey IKE VPN

No.	Match					Action	Install On
	From Zone	Source	To Zone	Destination	Service		
1	trust	Paris Trust LAN	untrust	Tokyo Trust LAN	any	tunnel	Paris
2	untrust	Paris Trust LAN	trust	Tokyo Trust LAN	any	tunnel	Tokyo
3	untrust	Tokyo Trust LAN	trust	Paris Trust LAN	any	tunnel	Paris
4	trust	Tokyo Trust LAN	untrust	Paris Trust LAN	any	tunnel	Tokyo

- Add the VPN Link. You must create a VPN link between the Zone rulebase in a security policy and the VPN Manager autogenerated rules. You create this link by inserting a VPN link in the zone rulebase; this link points to the VPN rules that exist in the VPN Manager.

- In Security Policies, select an existing security policy (or create a new security policy). In the Zone rulebase, right-click and select **Add VPN link**.
- Select the **Tokyo-Paris Policy-Based VPN**, then click **OK** to add the link. By default, the link appears at the top of the rulebase, but you can move the VPN link anywhere in the rulebase, just as you would a firewall rule.

Example: Configuring an Autokey IKE RAS, Policy-Based VPN

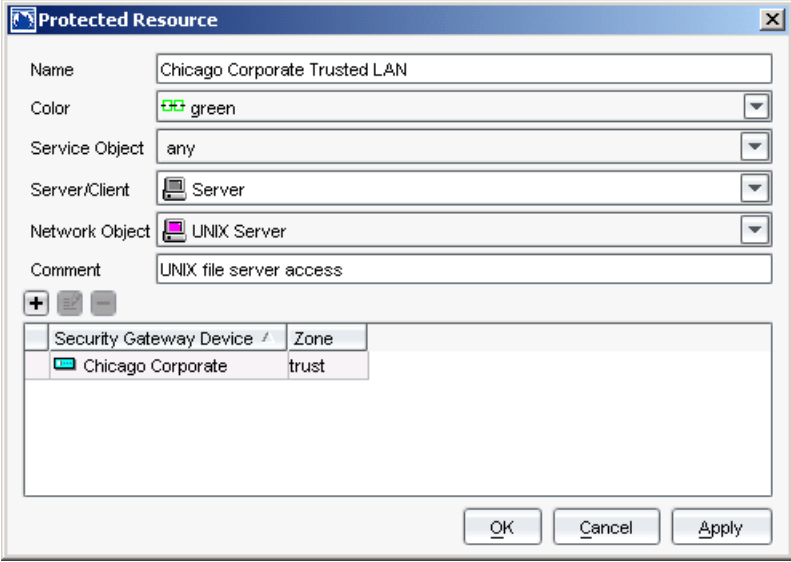
An AutoKey IKE RAS VPN connects RAS users and protected resources. In this example, Local Auth user Wendy (login name: reporter, password: Nd4syst4) wants to access resources on the UNIX server at the corporate site.

To accommodate Wendy, create an AutoKey IKE tunnel using a preshared key to provide the secure communication channel between IKE user Wendy and the UNIX server, which is protected by the Chicago Corporate security device.

The tunnel uses ESP with 3DES encryption and SHA-1 authentication. For the Phase 1 and 2 security levels, specify the Phase 1 proposal as pre-g2-3des-sha and select the predefined "Compatible" set of proposals for Phase 2.

1. Add the Chicago Corporate device and configure the following interfaces:
 - Ethernet1 is the Trust IP (10.1.1.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (1.1.1.1/24) in the Untrust zone.
2. Create the address objects that you will use to create Protected Resources (for details on creating or editing address objects).
 - a. Add the Chicago Corporate Trusted LAN (10.1.1.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
 - For Name, enter **Chicago Corporate Trust LAN**.
 - For IP Address/Netmask, enter **10.2.1.0/24**.
 - For Color, select **magenta**.
 - For Comments, enter **Chicago Trusted Network**.
 - b. Add the UNIX Server (10.1.1.5) as a host address object. In Address Objects, click the Add icon and select **Host**. Configure the following, then click **OK**:
 - For Name, enter **UNIX Server**.
 - For Color, select **magenta**.
 - For Comment, enter **UNIX file server, Chicago**.
 - Select IP and enter the IP Address **10.1.1.5**.
3. Create Chicago Corporate Trusted LAN Protected Resources to represent the destination point of the VPN. In Protected Resources (under VPN Manager), click the Add icon. Configure as shown in [Figure 101 on page 638](#), then click **OK**:

Figure 101: Add Chicago Protected Resource for AutoKey IKE RAS VPN



The **Protected Resource** dialog box is shown with the following configuration:

- Name:** Chicago Corporate Trusted LAN
- Color:** green
- Service Object:** any
- Server/Client:** Server
- Network Object:** UNIX Server
- Comment:** UNIX file server access

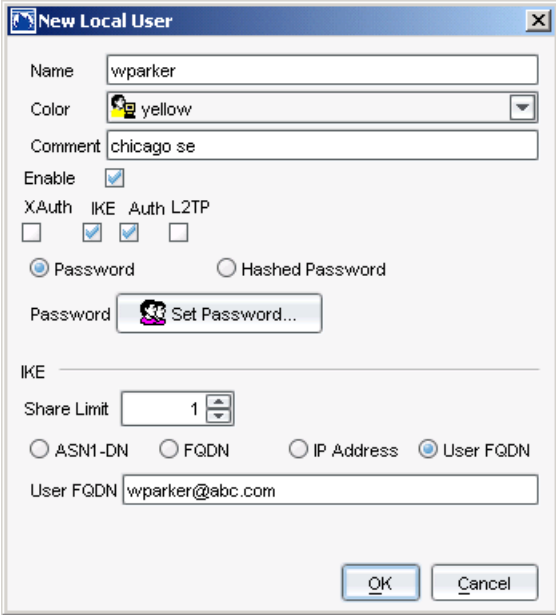
Below the fields is a table with two columns: **Security Gateway Device** and **Zone**.

Security Gateway Device	Zone
Chicago Corporate	trust

At the bottom are buttons for **OK**, **Cancel**, and **Apply**.

4. Create a local user object to represent Wendy. Local User objects are authenticated with the local NSM database.
5. In User Objects, select **Local User Objects**. In the main display area, click the Add icon and select **Local**. Configure, then click **OK**:

Figure 102: Add New Local User for AutoKey IKE RAS VPN



The **New Local User** dialog box is shown with the following configuration:

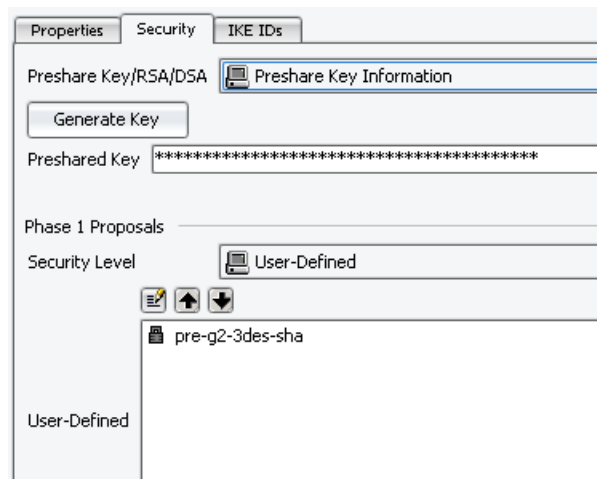
- Name:** wparker
- Color:** yellow
- Comment:** chicago se
- Enable:** ☒
- XAuth:** ☐ **IKE:** ☒ **Auth:** ☒ **L2TP:** ☐
- Password:** ☒ **Hashed Password:** ☐
- Password:** Set Password...
- IKE:**
 - Share Limit:** 1
 - ASN1-DN:** ☐ **FQDN:** ☐ **IP Address:** ☐ **User FQDN:** ☒
 - User FQDN:** wparker@abc.com

At the bottom are buttons for **OK** and **Cancel**.

6. Create the VPN. In the navigation tree, double-click **VPN Manager**, then right-click **VPNs** and select **AutoKey IKE RAS VPN**. The New AutoKey IKE RAS VPN dialog box appears. Configure as shown below:

- a. For Name, enter **UNIX Remote Access VPN**.
 - b. Select **Enable**.
 - c. In Termination Point, select **Untrust**.
 - d. Click **OK** to save the VPN and return to VPN Manager. In VPN Manager, select the **UNIX Remote Access VPN**.
7. Configure the policy-based members:
 - a. In the main display area, select the **Protected Resources** link.
 - b. In the Protected Resources list, select the Chicago **Corporate Trusted LAN**, then click **OK** to return to the main display area.
8. Configure the termination points of the VPN:
 - a. Click the **Termination Points** link. The Termination Points dialog box appears.
 - b. Configure Chicago Corporate to use ethernet3 as the termination point (this is the Untrust interface), then click **OK** to return to the main display area.
9. Configure the remote users for the VPN:
 - a. In the Remote Users section, click the **Users** link. The Remote User dialog box appears.
 - b. Select the user "wparker", then click **Save** to save your configuration changes to the VPN.
10. Configure the VPN gateway:
 - a. Click the **Gateway Parameters** link. The Properties tab appears. Leave all defaults and click the Security tab.
 - b. In the Security tab, enter the preshared key value (h1p8A24nG5), then click **Generate Key**.
 - c. For Phase 1 Proposals, select **User-Defined**, then click the Add/Edit icon to add the pre-g2-3des-sha proposal.

Figure 103: Configure Security for AutoKey IKE RAS VPN



- d. Click **Save** to save your configuration changes to the VPN.

To view the autogenerated rules, click the **Policy Rules** link in the Overrides section. VPN Manager generates the rules.

11. Add the VPN Link. You must create a VPN link between the security policy and the VPN Manager autogenerated rules. You create this link by inserting a VPN link in the security policy; this link points to the VPN rules that exist in the VPN Manager.
 - a. In Security Policies, select an existing security policy (or create a new security policy). Right-click and select **Add VPN** link.
 - b. Select the **UNIX Remote Access VPN**.
 - c. Click **OK** to add the link to the policy. By default, the link appears at the top of the policy, but you can move the VPN link anywhere in the policy, just as you would a firewall rule.

Example: Configuring an Autokey IKE, Route-Based Site-to-Site VPN

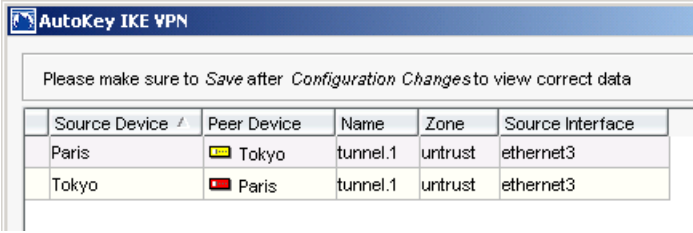
In this example, an AutoKey IKE VPN tunnel using a preshared key provides a secure connection between security devices protecting the Tokyo and Paris offices. The Untrust zone interface for both security devices use a static IP address. All security and tunnel zones are in the trust-vr. The preshared key is h1p8A24nG5. For the Phase 1 and 2 security levels, specify the Phase 1 proposal as pre-g2-3des-sha and the Phase 2 proposal as predefined compatible.

1. Configure the Tokyo device with the following interfaces:
 - Ethernet1 is the Trust IP (10.1.1.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (1.1.1.1/24).
2. Configure the Paris device with the following interfaces:

- Ethernet1 is the Trust IP (10.2.2.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (2.2.2.2/24) in the Untrust zone.
 - Create the address objects that you use for the VPN rule in the firewall rulebase (for details on creating VPN rules, see [“Adding VPN Rules” on page 661](#)).
3. Add the Tokyo Trust LAN (10.1.1.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
 4. Add the Tokyo and Paris security devices (for details on adding devices:
 - For Name, enter **Tokyo Trust LAN**.
 - For IP Address/Netmask, enter **10.1.1.0/24**.
 - For Color, select **magenta**.
 - For Comment, enter **Tokyo Trust Zone**.
 5. Add the Paris Trust LAN (10.2.2.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
 - For Name, enter **Paris Trust LAN**.
 - For IP Address/Netmask, enter **10.2.2.0/24**.
 - For Color, select **magenta**.
 - For Comment, enter **Paris Trust Zone**.
 - Create the VPN. In the navigation tree, double-click **VPN Manager**. Right-click **VPNs** and select **AutoKey IKE VPN**. The New AutoKey IKE VPN dialog box appears. Configure as shown below:
 6. In Name, enter **Tokyo-Paris Route-based VPNs**.
 7. Select **Enable**.
 8. In Termination Point, select **Untrust**.
 9. Click **OK** to save the VPN and return to VPN Manager. In VPN Manager, select the **Tokyo-Paris Route-based VPNs**.
 10. Configure the route-based members:
 - a. In the main display area, select the security device link (under Route-Based Configuration) to display the zone and tunnel options. Configure the default zone and tunnel options.
 - b. Click the Add icon to display available security devices. Select the Paris and Tokyo devices.
 - c. Click **OK** to add the members to the VPN.
 - d. Ensure that the route-based members are configured.
 - e. Click **OK** to save your settings and return to the main display area.
 - f. Configure the VPN topology:
 - g. Select the **Topology** link. The Topology dialog box appears.

- h. Click the Add icon to display the Topology configuration dialog box.
 - i. In the Mains window, select the Paris and Tokyo security devices.
 - j. Click **OK** to return to the Topology dialog box, then click **OK** to return to the main display area.
 - k. Configure the termination points of the VPN:
 - l. Click the **Termination Points** link. The Termination Points dialog box appears.
 - m. Confirm that both Paris and Tokyo devices use a Termination Interface of ethernet3, then click **OK** to return to the main display area.
 - n. Configure the VPN gateway:
 - o. Click the **Gateway Parameters** link. The Properties tab appears. Leave all defaults and click the Security tab.
 - p. In the Security tab, enter the preshared key value (h1p8A24nG5), then click **Generate Key**.
 - q. For Phase 1 Proposals, select **User-Defined**, then click the Add/Edit icon to add the pre-g2-3des-sha proposal.
11. Click **Save** to save your configuration changes to the VPN. Because this VPN is route-based, no rules are autogenerated. However, you can view the device tunnel summary to see all autogenerated tunnels between each security device in the VPN.

Figure 104: View Tunnel Summary for AutoKey IKE, RB Site-to-Site VPN



The screenshot shows a window titled "AutoKey IKE VPN". Below the title bar is a message: "Please make sure to Save after Configuration Changes to view correct data". Below this is a table with the following data:

Source Device	Peer Device	Name	Zone	Source Interface
Paris	Tokyo	tunnel.1	untrust	ethernet3
Tokyo	Paris	tunnel.1	untrust	ethernet3

A tunnel interface acts as a doorway to a VPN tunnel; traffic enters and exits a VPN tunnel via a tunnel interface. These tunnels are an "always-on" connection—the devices will route any traffic with an appropriate source and destination IP address through the VPN tunnel.

To control traffic through the tunnel, you must add firewall rules to the security policy that is installed on each VPN node.

Next, you must create the routes (in the route table of each device) that will connect the autogenerated tunnel interfaces and form the VPN tunnel (for details on creating routes, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*. You can use static or dynamic routes, however, this example details only the static route creation. For each device, you will create two routes using the trust virtual router (trust-vr):

- A route from 0.0.0.0/0 to eth3 in the untrust zone. This routes traffic from the trust zone through eth3 in the untrust zone, then to the next hop (default) gateway.
 - A route from the tunnel.1 interface (autogenerated by VPN Manager) to the untrust zone of the remote VPN node. This routes traffic destined for the remote VPN node through the tunnel.1 interface (where the packets are encapsulated), with a default next hop gateway of 0.0.0.0/0.
 - Configure the route on the Tokyo security device.
1. In Device Manager, double-click the device to open the device configuration dialog box. Select **Network** > **Virtual Router** to display the list of virtual routers on the device.
 2. Double-click the trust-vr route to open the vr for editing. In the virtual router dialog box, click **Routing Table**, then click the Add icon under destination-based Routing Table to add a new static route.
 3. Configure a route from the untrust interface to the gateway.
 4. Configure route from the trust zone to the tunnel interface.
 5. Click **OK** to save your changes to the virtual router, then click **OK** to save your changes to the Tokyo device.
 6. Configure the route on the Paris security device:
 7. In Device Manager, double-click the device to open the device configuration dialog box. Select **Network** > **Virtual Router** to display the list of virtual routers on the device.
 8. Double-click the trust-vr route to open the vr for editing. In the virtual router dialog box, click **Routing Table**, then click the Add icon under destination-based Routing Table to add a new static route.
 9. Configure a route from the untrust interface to the gateway.
 10. Configure route from the trust zone to the tunnel interface.
 11. Click **OK** to save your changes to the virtual router, then click **OK** to save your changes to the Paris device.

Example: Configuring XAuth Authentication with External User Group

In this example, you use a VPN to enable access for a group of resellers who require access to FTP servers in the corporate LAN. First, you must configure the RADIUS server using the custom port 4500 (default is 1645), then add an authentication server object in NSM to represent that server.

Next, to manage the users in this example, you define an external user group in two places: on the external RADIUS auth server and in NSM.

- On the RADIUS server, you populate the external user group with XAuth users, leaving the group unpopulated NSM. The RADIUS server authenticates the users during Phase 1 IKE negotiations.
- In NSM, you leave the external user group unpopulated, but you must define each user as a local user with an IKE ID, then create a group that includes those local users as

members. This IKE ID is used to authenticate the users during the Phase 2 IKE negotiations.

Additionally, you must add the security device and create an Address object to represent the FTP server, as well as a protected resource. After you have assembled all the VPN components, you are ready to create the VPN.

1. Configure the RADIUS Server. On the RADIUS server, load the Juniper Networks dictionary file and define Xauth user accounts. Use the Juniper Networks user group VSA to create the user group `xa_grp2` and apply it to the auth user accounts that you want to add to that group.



NOTE: For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation. If you are using a Microsoft IAS RADIUS server, there is no dictionary file to load. Instead, define the correct vendor-specific attributes (VSAs) on the server.

2. Add the authentication server object. In the main navigation tree, select **Object Manager > Authentication Servers** and click the Add icon. Configure the following, then click **OK**:
 - a. For name, enter **radius1**. Select a color and add a comment, if desired.
 - b. For Main Server, enter the IP **10.20.1.100**; for Primary Backup Server, enter IP **10.20.1.110**; for Secondary Backup Server, enter IP **10.20.1.120**.
 - c. For timeout, enter **30**.
 - d. Enable For XAuth Users.
 - e. For Server Type, select **RADIUS**.
1. Configure the RADIUS server:
 - For server port, select **4500** (default is 1645).
 - For secret, enter **A56hYY97kl**.
 - For retry timeout, enter **4**.
2. Add an External User Group (in NSM). In the Object Manager, select **User Objects > External User Groups**. Click the Add icon to display the New External User Group dialog box. Configure the following, then click **OK**.
 - For Name, enter **xa-grp2**.
 - For Color, select **yellow**.
 - For Comment, enter **Reseller Group RADIUS**.
 - Enable XAuth.
3. Add the local user object. In the Object Manager, select **User Objects > Local Users**. Click the Add icon and select User. The New Local User dialog box appears. Configure the following, then click **OK**.

- For Name, enter **jhansen**.
 - For Color, select **orange**.
 - For Comment, enter **reseller group**.
 - Select **Enable**, then select **IKE**.
 - For IKE settings, enable **User FQDN** and enter the e-mail address **jhansen@company.com**.
4. Add a Local User Group. In the Object Manager, select **User Objects > Local User Groups**. Click the Add icon to display the New Local User Group dialog box. Configure the following, then click **OK**:
 - For Name, enter **Reseller User Group**.
 - For Comment, enter **Reseller VPN XAuth RADIUS**.
 - For color, enter **green**.
 - Add jhansen as a member.
 5. Add a Network address object to represent the network used by Reseller group. In the Object Manager, select **Address Objects**, then click the Add icon and select **Network**. The New Network dialog box appears. Configure the following, then click **OK**:
 - For Name, enter **reseller1**.
 - For IP Address/Netmask, enter **10.2.2.0/24**.
 - For color, select **cyan**.
 - For Comment, enter **Reseller Group**.
 6. Add an address object to represent the FTP Server. In the Object Manager, select **Address Objects**, then click the Add icon and select **Host**. The New Host dialog box appears. Configure the following, then click **OK**:
 - For Name, enter **rsl-svr1**.
 - For Color, select **green**.
 - For Comment, enter **FTP Server**.
 - Select **IP**, then enter the IP Address **10.1.1.5**.
 7. Add a NetScreen-208 security device named "Bozeman." This is the device protects the FTP server. Configure the Bozeman device with the following interfaces:
 - Ethernet1 is the Trust IP (10.1.1.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (2.2.2.2/24) in the Untrust zone.
 8. Create a Protected Resource to represent the destination point of the VPN. In this example, the destination point is the FTP server in the trust zone of Bozeman. In

Protected Resources (under VPN Manager), click the Add icon. Configure the object, and then click **OK**:

9. Create the VPN. In the main navigation tree, select **VPN Manager > VPNs**. Click the Add icon and select **AutoKey IKE RAS VPN**. The New AutoKey IKE RAS VPN dialog box appears. Configure as shown below:
 - In Name, enter **Reseller Remote Access VPN**.
 - Select **Enable**.
 - In Termination Point, select **Untrust**.
 - Click **OK** to save the VPN and return to VPN Manager. The Reseller Remote Access VPN appears in the main display area.
1. Configure the policy-based members:
 - In the main display area, select the **Protected Resources** link.
 - In the Protected Resources list, select the rsl-svr1 protected resource, then click **OK**:
2. Configure the termination points of the VPN:
 - Click the **Termination Points** link. The Termination Points dialog box appears.
3. Configure the Bozeman device to use ethernet3 as the termination point (this is the Untrust zone interface).
4. Click **OK** to return to the main display area.
5. Configure the remote users for the VPN:
 - In the Remote Users section, click the **Users** link. The Remote User dialog box appears.
 - Select the Reseller local user group.
 - Click Save to **save** your configuration changes to the VPN.
6. Configure the VPN gateway:
 - Click the **Gateway Parameters** link. The Properties tab appears.
 - For Mode, select **Main**.
 - In the XAuth section, select XAuth Server and then select the radius1 authentication server for Auth Server Name. Later, after you have autogenerated the VPN rules and gateway, you can override this setting to include only the Reseller external user group.
 - In the Security tab, enter the preshared key value (netscreen4), then click **Generate Key**.
 - For Phase 1 Proposals, select User-Defined, then click the Add/Edit icon to add the pre-g2-3des-sha proposal.

- Click **OK** to save your changes to the gateway.
- Click **Save** to save your configuration changes to the VPN and autogenerate the policy rules.

To view the autogenerated rules, click the **Policy Rules** link in the Overrides section. VPN Manager generates the rules.

7. Configure Overrides. By default, the gateway attempts to authenticate all users using the specified authentication server (radius1). You must override the gateway security settings to enable the VPN to authenticate only the Reseller external user group:
 - In the overrides area, click the **Device Configuration** link.
 - In the navigation tree, double-click **Bozeman** and select **Gateway**. The autogenerated gateway for the Bozeman appears in the main display area.
 - Right-click the autogenerated gateway and select **Edit**. The Properties tab appears.
 - In the IKE IDs/XAuth tab, configure the XAuth area to authenticate only the Reseller external group.
 - For user, select **User Group**.
 - For User Group, select **xa-grp2**.
 - Click **OK** to save your overrides.
8. Add the VPN Link. You can create a VPN link between the security policy and the VPN Manager autogenerated rules. You create this link by inserting a VPN link in the security policy; this link points to the VPN rules that exist in the VPN Manager.
 - In Security Policies, select an existing security policy (or create a new security policy). Right-click and select **Add VPN** link.
 - Select the **Reseller Remote Access VPN**.
 - Click **OK** to add the link to the policy.

By default, the link appears at the top of the policy, but you can move the VPN link anywhere in the policy, just as you would a firewall rule.

Creating Device-Level VPNs

You can create four types of device-level VPNs:

- Use an **AutoKey IKE VPN** to connect devices and/or protected resources. An AutoKey IKE VPN supports mixed-mode, policy-based, and routing-based VPNs, but does not support RAS users. For details on each step, see [“Creating AutoKey IKE VPNs” on page 648](#).
- Use a **Manual Key IKE VPNs** to authenticate devices, protected resources, and RAS users in the VPN with manual keys. For details on each step, see [“Creating Manual Key VPNs” on page 656](#).

- Use an **L2TP RAS VPN** to connect L2TP RAS users and protected resources with authentication but without encryption. For details on each step, see [“Creating L2TP VPNs” on page 659](#).
- Use an **L2TP-over-AutoKey IKE RAS VPN** to connect L2TP RAS users and protected resources. An L2TP-over-AutoKey IKE RAS VPN supports policy-based VPNs and L2TP RAS users, but does not support routing-based VPNs. For details on each step, see [“Creating L2TP Over Autokey IKE VPNs” on page 660](#).

Supported Configurations

IKE VPNs support tunnel mode, and can be policy-based or route-based; however, route-based VPNs do not support RAS users.

L2TP VPNs support transport mode, and can be policy-based.

Creating AutoKey IKE VPNs

Creating device-level AutoKey IKE VPNs is a four stage process:

- Configure Gateway
- Configure Routes (Route-based only)
- Configure VPN on the Device
- Add VPN rules to security policy

IKEv2 and EAP Support

As part of the ScreenOS support, NSM allows you to configure IKEv2 features which include identity hiding, perfect forward secrecy, two phases, and cryptographic negotiation. The protocol redesign makes IKEv1 incompatible with IKEv2 even though they both use the UDP port (500 or 4500) for communication.

IKEv2 also supports Extensible Authentication Protocol (EAP). Using EAP, IKEv2 can leverage the existing authentication infrastructure and credential databases, because EAP allows users to choose a suitable method for existing credentials, and also facilitates separation of the IKEv2 responder (VPN gateway) from the EAP authentication endpoint (backend AAA server).

From the NSM UI, you can:

- Set the global account type to be authenticated by the authentication server:
 1. Navigate from **Object Manager > Authentication Servers**.
 2. Select **For IKEv2EAP users** from the Authentication Servers List.
 3. Click **OK**.
- Specify the self and peer authentication, and authentication methods for the IKEv2 gateway:

1. Navigate from **VPN Settings > Gateway > Gateway properties**.
2. Enter the required information in the **Authenticated by EAP** fields.
3. Enter the required information in the **Auth-method** fields. If you enabled IKEv1, then these fields are hidden. This setting is also hidden on devices that do not support IKEv2.
4. Click **OK**.
5. Navigate from **VPN Manager > VPNs > AutoKey IKE VPN > Gateway properties**
6. Select the authentication method for the device.
7. Click **OK**.

For details on configuring IKEv2 for VPNs, refer to the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

Configuring Gateways

A gateway is an interface on your security device that sends and receives traffic; a remote gateway is an interface on another device that handles traffic for that device. Each security device member has a remote gateway that it sends and receives VPN traffic to and from. To configure a gateway for a VPN member, you need to define the local gateway (the interface on the VPN member that handles VPN traffic) and the remote gateway (the interface on the other VPN member that handles VPN traffic). The interface can be physical or virtual.

- For remote gateways that use static IP addresses, specify the IP address or host name of the remote device.
- For remote gateways that use dynamic IP addresses, configure an IKE ID for the remote device.
- For remote gateways that are RAS users, specify a Local User object as a remote gateway to enable RAS user access.

To add a gateway to a security device, open the device configuration, select **VPN Settings**, and click the Add icon to display the New Gateway Dialog box. Configure the gateway as detailed in the following sections.

Properties

Enter a name for the new gateway, then specify the following gateway values:

- **Mode**—The mode determines how Phase 1 negotiations occur.
 - In *Main* mode, the IKE identity of each node is protected. Each node sends three two-way messages (six messages total); the first two messages negotiate encryption and authentication algorithms that protect subsequent messages, including the IKE identity exchange between the nodes. Depending on the speed of your network connection and the encryption and authentication algorithms you use, main mode negotiations can take a long time to complete. Use Main mode when security is more important.

- In *Aggressive* mode, the IKE identity of each node is not protected. The initiating node sends two messages and the receiving node sends one (three messages total); all messages are sent in the clear, including the IKE identity exchange between the nodes. Because Aggressive mode is typically faster but less secure than Main mode, use Aggressive mode when speed is more important than security. However, you must use Aggressive mode for VPNs that include RAS users.
- Remote Gateway—The remote gateway is the VPN gateway on the receiving VPN node, and can be an interface with a static or dynamic IP address, or local or external user object.
 - Static IP Address. For remote gateways that use a static IP address, enter the IP address and mask.
 - RAS User/Group. For remote gateways that are users, select the User object or User Group object that represents the RAS user.
 - Dynamic IP Address. For remote gateways that use a dynamic IP address, select dynamic IP address.
- Outgoing Interface—The outgoing interface (also known as the termination interface) is the interface on the security device that sends and receives VPN traffic. Typically, the outgoing interface is in the untrust zone.
- Heartbeats—Use heartbeats to enable redundant gateways. You can use the default or set your own thresholds:
 - Hello. Enter the number of seconds the security device waits between sending hello pulses.
 - Reconnect. Enter the maximum number of seconds the security device waits for a reply to the hello pulse.
 - Threshold. Enter the number of seconds that the security device waits before attempting to reconnect.
- NAT Traversal—Because NAT obscures the IP address in some IPSec packet headers, a VPN node cannot receive VPN traffic that passes through an external NAT device. To enable VPN traffic to traverse a NAT device, you can use NAT Traversal (NAT-T) to encapsulate the VPN packets in UDP. If a VPN node with NAT-T enabled detects an external NAT device, it checks every VPN packet to determine if NAT-T is necessary. Because checking every packet impacts VPN performance, you should only use NAT Traversal for remote users that must connect to the VPN over an external NAT device.

You do not need to enable NAT-T for your internal security device nodes that use NAT; each VPN node knows the correct address translations for VPN traffic and does not need to encapsulate the traffic.

To use NAT-T, enable NAT-Traversal and specify:

- UDP Checksum. A 2-byte value (calculated from the UDP header, footer, and other UDP message fields) that verifies packet integrity. You must enable this option for NAT devices that require UDP checksum verification; however, most NAT devices (including security devices) do not require it.

- **Keep alive Frequency.** The number of seconds a VPN node waits between sending empty UDP packets through the NAT device. A NAT device keeps translated IP addresses active only during traffic flow, and invalidates unused IP addresses. To ensure that the VPN tunnel remains open, you can configure the VPN node to send empty "keep alive" packets through the NAT device.

IKE IDs/XAuth

Every VPN member has a unique identification number, known as an IKE ID. During Phase 1 negotiations, the IKE protocol uses the ID to authenticate the VPN member. You must select and configure an ID type for the VPN members at each end of the tunnel. However, the ID type can be different for each member:

- **ASN1-DN**—Abstract Syntax Notation, version 1 is a data representation format that is non-platform specific; Distinguished Name is the name of the computer. Use ASN1-DN to create a Group ID that enables multiple RAS users to connect to the VPN tunnel concurrently.
 - At the peer ID, specify values for the Container Match and Wildcard Match.
 - At the local ID, specify the value.

Using a Group ID can make configuring and maintaining your VPN quicker and easier. For details on how Group IKE IDs work, see [“Configuring Group IKE IDs” on page 611](#). For details on determining the ASN1-DN container and wildcard values for Group IKE IDs, see the documentation for your version of ScreenOS.

- **FQDN**—Use a Fully Qualified Domain Name when the VPN member uses a dynamic IP address. FQDN is a name that identifies (qualifies) a computer to the DNS protocol using the computer name and the domain name; ex. server1.colorado.mycompany.com.
- **IP Address**—Use an IP address when the VPN member uses a static IP address.
- **U-FQDN**—Use a User Fully Qualified Domain Name when the VPN member uses a dynamic IP address (such as a RAS user). A U-FQDN is an e-mail address, such as user1@mycompany.com.

Use the XAuth protocol to authenticate RAS users with an authentication token (such as SecureID) and to make TCP/IP settings (IP address, DNS server, and WINS server) for the peer gateway.

- **Default Server**—Use the default server to use the default XAuthentication server for the device. To change or assign a default XAuthentication server, edit the VPN settings > Defaults > Xauth settings.
- **XAuth Server**—Use to specify the authentication server that assigns TCP/IP settings to the remote gateway.
 - **XAuth Server Name.** Select a preconfigured authentication server object. For details on creating authentication server objects.
 - **Allowed Authentication Type.** Select generic or Challenge Handshake Authentication Protocol (CHAP) (password is sent in the clear) to authenticate the remote gateway.

- Query Remote Setting. Enable this option to query the remote settings object for DNS and WINS information.
- Users and Groups. To authenticate XAuth RAS users using the authentication server, enable User or User Group and select a preconfigured user object.
- XAuth Client—Use when the remote gateway is a RAS user that you want to authenticate.
- Allowed Authentication Type. Select Any or Challenge Handshake Authentication Protocol (CHAP) for authentication (password is sent in the clear).
- User Name and Password. Enter the user name and password that the RAS user must provide for authentication.



NOTE: All passwords handled by NSM are case-sensitive.

- Bypass Authentication—Use to permit VPN traffic from this VPN member to pass unauthenticated by the Auth server.

Security

Select the authentication method you want to use in the VPN:

- Preshared Key—Use if your VPN includes security devices and/or RAS users. VPN nodes use the preshared key during Phase 1 negotiations to authenticate each other; because each node knows the key in advance, negotiations use fewer messages and are quicker.
- To generate a random key, enter a value for the seed, then click Generate Key. NSM uses the seed value to generate a random key, which is used to authenticate VPN members.



NOTE: Using a random key can generate a value in excess of 255 characters, which exceeds ScreenOS limits and might not be accepted by the security device during update. To reduce the key size, shorten the autogenerated key value by deleting characters.

- To use a predefined value for the key, enter a value for the Preshared Key.
- PKI—Use if your VPN includes extranet devices or you require the additional security provided by certificates (PKI uses certificates for VPN member authentication).

For Phase 1 negotiations, select a proposal or proposal set. You can select from predefined or user-defined proposals:

- To use a predefined proposal set, select one of the following:
 - Basic (*nopfs-esp-des-sha*, *nopfs-esp-des-md5*)
 - Compatible (*nopfs-esp-3des-sha*, *nopfs-esp-3des-md5*, *nopfs-esp-des-sha*, *nopfs-esp-des-md5*)

- Standard (*gs-esp-3des-sha*, *gs-esp-aes128-sha*)



NOTE: You cannot use a predefined proposal set with certificates—you must select a user-defined proposal or change the authentication method to Preshared Key.

- To use a user-defined proposal, select a single proposal from the list of predefined and custom IKE Phase 1 Proposals. For details on custom IKE proposals.

If your VPN includes only security devices, you can specify one predefined or custom proposal that NSM propagates to all nodes in the VPN. If your VPN includes extranet devices, you should use multiple proposals to increase security and ensure compatibility.

Configuring Routes (Route-based only)

For a routing-based VPN member, you must configure:

- Tunnel zone or tunnel interfaces on the member.
- Static or dynamic routes from the member to other VPN members.

VPN traffic flows through the tunnel zones or tunnel interfaces on the security device, and uses static or dynamic routes to reach other VPN members. You must create the tunnel zones and interfaces before configuring routes.

For details on configuring tunnel zones, tunnel interfaces, static routes, or dynamic routes, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

After you have configured the tunnel zone or interface on the security device, you must bind the VPN to that zone or interface to make the VPN functional, as described in the following section.

Configuring the VPN

When you configure the VPN, you are defining the gateway the security device uses to connect to the VPN, the IKE Phase 2 proposals used by that gateway, and how you want NSM to monitor the VPN tunnel.

For route-based VPNs, you are also binding the VPN to the tunnel interface or zone that sends and receives VPN traffic to and from the device.

Properties

Enter the following values:

- VPN name—Enter a name for the VPN.
- Remote Gateway—Select the gateway for the VPN.
- Idle Time to Disable SA—Configure the number of minutes before a session that has no traffic automatically disables the SA.
- Replay Protection—In a replay attack, an attacker intercepts a series of legitimate packets and uses them to create a denial-of-service (DoS) against the packet

destination or to gain entry to trusted networks. If replay protection is enabled, your security devices inspect every IPSec packet to see if the packet has been received before—if packets arrive outside a specified sequence range, the security device rejects them.

- **IPSec Mode—Configure the mode:**
 - Use tunnel mode for IPSec. Before an IP packet enters the VPN tunnel, NSM encapsulates the packet in the payload of another IP packet and attaches a new IP header. This new IP packet can be authenticated, encrypted, or both.
 - Use transport mode for L2TP-over-IPSec. NSM does not encapsulate the IP packet, meaning that the original IP header must remain in plaintext. However, the original IP packet can be authenticated, and the payload can be encrypted.
- **Do not set Fragment Bit in the Outer Header—**The Fragment Bit controls how the IP packet is fragmented when traveling across networks.
 - Clear. Use this option to enable IP packets to be fragmented.
 - Set. Use this option to ensure that IP packets are not fragmented.
 - Copy. Select to use the same option as specified in the internal IP header of the original packet.

Security

For Phase 2 negotiations, select a proposal or proposal set. You can select from predefined or user-defined proposals:

- To use a predefined proposal set, select one of the following:
 - Basic (*nopfs-esp-des-sha*, *nopfs-esp-des-md5*)
 - Compatible (*nopfs-esp-3des-sha*, *nopfs-esp-3des-md5*, *nopfs-esp-des-sha*, *nopfs-esp-des-md5*)
 - Standard (*gs-esp-3des-sha*, *gs-esp-aes128-sha*)
- To use a user-defined proposal, select a single proposal from the list of predefined and custom IKE Phase 2 Proposals.

If your VPN includes only security devices, you can specify one predefined or custom proposal that NSM propagates to all nodes in the VPN. If your VPN includes extranet devices, you should use multiple proposals to increase security and ensure compatibility.

Binding/ProxyID

You can bind the VPN tunnel to a tunnel interface or tunnel zone to increase the number of available interfaces in the security device. To use a tunnel interface and/or tunnel zone in your VPN, you must first create the tunnel interface or zone on the device; for details, see [“Configuring Tunnel Interfaces and Tunnel Zones” on page 612](#) and the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

- **None**—Select none when you do not want to bind the VPN tunnel to a tunnel interface or zone.
- **Tunnel Interface**—Select a preconfigured tunnel interface on the security device to bind the VPN tunnel to the tunnel interface. The security device routes all VPN traffic through the tunnel interface to the protected resources.
- **Tunnel Zone**—Select a preconfigured tunnel zone on the security device to bind the VPN tunnel directly to the tunnel zone. The tunnel zone must include one or more numbered tunnel interfaces; when the security device routes VPN traffic to the tunnel zone, the traffic uses one or more of the tunnel interfaces to reach the protected resources.
- **DSCP Marks** — ScreenOS 6.1 and later supports the DSCP value configuration for tunnel mode ESP packets only. You cannot configure the DSCP setting if the IPSec mode is tunnel mode but the binding interface is not a tunnel interface.

You can set the following **DSCP Marks** under the **Binding /Proxy** tab on the **AutoKey IKE Parameters** page:

- **DSCP Marking** — You can select either enable or disable. If the selected IPSec mode is transport, this option is automatically disabled.
- **DSCP Value** — Set the DSCP value in the range of 0–63. Mouse over the field to see the range of allowed values.

You can also enable proxy and configure the proxy parameters.

Monitor

You can enable VPN Monitor and configure the monitoring parameters for the device. Monitoring is off by default. To enable the VPN Monitor in Realtime Monitor to display statistics for the VPN tunnel, configure the following:

- **VPN Monitor**—When enabled, the device sends ICMP echo requests (pings) through the tunnel at specified intervals (configurable in seconds) to monitor network connectivity (the device uses the IP address of the local outgoing interface as the source address and the IP address of the remote gateway as the destination address). If the ping activity indicates that the VPN monitoring status has changed, the device triggers an SNMP trap; VPN Monitor (in RealTime Monitor) tracks these SNMP statistics for VPN traffic in the tunnel and displays the tunnel status.
- **Rekey**—When enabled, the device regenerates the IKE key after a failed VPN tunnel attempts to reestablish itself. When disabled, the device monitors the tunnel only when the VPN passes user-generated traffic (instead of using device-generated ICMP echo requests). Use the rekey option to:
 - Keep the VPN tunnel up even when traffic is not passing through
 - Monitor devices at the remote site.
 - Enable dynamic routing protocols to learn routes at a remote site and transmit messages through the tunnel.
 - Automatically populate the next-hop tunnel binding table (NHTB table) and the route table when multiple VPN tunnels are bound to a single tunnel interface.

- **Optimized**—When enabled, the device optimizes its VPN monitoring behavior as follows:
 - Considers incoming traffic in the VPN tunnel as ICMP echo replies. This reduces false alarms that might occur when traffic through the tunnel is heavy and the echo replies cannot get through.
 - Suppresses VPN monitoring pings when the tunnel passes both incoming and outgoing traffic. This can help reduce network traffic.
- **Source Interface and Destination IP**—Configure these options to use VPN Monitoring when the other end of the VPN tunnel is not a security device. Specify the source and destination IP addresses.

Adding a VPN Rule

After you have configured the VPN on each device you want to include in the VPN, you can add a VPN rule to a security policy:

- For policy-based VPNs, you must add a VPN rule to create the VPN tunnel.
- For route-based VPNs, the VPN tunnel is already in place. However, you might want to add a VPN rule to control traffic through the tunnel.

For details on adding and configuring a VPN rule in a security policy, see [“Adding VPN Rules” on page 661](#).

Creating Manual Key VPNs

Creating a device-level Manual Key VPN is a four stage process:

1. Configure XAuth Users
2. Configure Routes (Route-based only)
3. Configure VPN on Device
4. Add VPN rules to security policy

Adding XAuth Users

For VPNs that use IPSec manual key to provide remote access services, you must add an XAuth User to the security device. An XAuth User has an account on the security device that guards the protected resources in the VPN; when the user attempts to connect to a protected resource, the security device authenticates the user.

To add a XAuth User for a security device, in the security device configuration L2TP/XAuth/Local User, click the Add icon. Enter a name for the user, then specify:

- **User**—Select a preconfigured Local User object that is configured for XAuth.
- **Remote Setting**—Select a preconfigured Remote Settings object.
- **IP Pool**—Select a preconfigured IP Pool object.
- **Static IP**—Enter the static IP address of the Local User.

Configuring Routes (Route-based only)

For a routing-based VPN member, you must configure:

- Tunnel zone or tunnel interfaces on the member.
- Static or dynamic routes from the member to other VPN members.

VPN traffic flows through the tunnel zones or tunnel interfaces on the security device, and uses static or dynamic routes to reach other VPN members. You must create the tunnel zones and interfaces before configuring routes. For details on configuring tunnel zones, tunnel interfaces, and static or dynamic routes, see the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

After you have configured the tunnel zone or interface on the security device, you must bind the VPN to that zone or interface to make the VPN functional, as described in the following section.

Configuring the VPN

The following sections detail how to configure the VPN.

Properties

Enter the following values:

- VPN name—Enter a name for the VPN.
- Gateway—Enter a gateway for the VPN.
- Local SPI—The local Security Parameter Index.
- Remote SPI—The remote Security Parameter Index.
- Outgoing Interface—The outgoing interface is the interface on the security device that sends and receives VPN traffic. Typically, the outgoing interface is in the untrust zone.
- Do not set Fragment Bit in the Outer Header—The Fragment Bit controls how the IP packet is fragmented when traveling across networks.
 - Clear. Use this option to enable IP packets to be fragmented.
 - Set. Use this option to ensure that IP packets are not fragmented.
 - Copy. Select to use the same option as specified in the internal IP header of the original packet.
- IPSec Protocol—Specify the IPSec protocol and algorithm you want to use for data authentication and/or encryption. Because this information is static for each VPN member, they do not need to negotiate for communication.
 - AH. Use Authentication Header to authenticate the VPN traffic, but not encrypt the traffic. If you select AH, you must also specify the key or password that AH uses in the authentication algorithm.



NOTE: All passwords handled by NSM are case-sensitive.

- ESP. Use Encapsulating Security Payload to authenticate and encrypt the VPN traffic. If you select ESP, because ESP uses keys to encrypt and decrypt data, you must also specify the key or password that the VPN node uses to send and receive VPN data through the VPN tunnel.

Binding

You can bind the VPN tunnel to a tunnel interface or tunnel zone to increase the number of available interfaces in the security device. To use a tunnel interface and/or tunnel zone in your VPN, you must first create the tunnel interface or zone on the device; for details, see [“Configuring Tunnel Interfaces and Tunnel Zones” on page 612](#) and the *Network and Security Manager Configuring ScreenOS and IDP Devices Guide*.

- None—Select none when you do not want to bind the VPN tunnel to a tunnel interface or zone.
- Tunnel Interface—Select a preconfigured tunnel interface on the security device to bind the VPN tunnel to the tunnel interface. The security device routes all VPN traffic through the tunnel interface to the protected resources.
- Tunnel Zone—Select a preconfigured tunnel zone on the security device to bind the VPN tunnel directly to the tunnel zone. The tunnel zone must include one or more numbered tunnel interfaces; when the security device routes VPN traffic to the tunnel zone, the traffic uses one or more of the tunnel interfaces to reach the protected resources.

Monitor

You can enable VPN Monitor and configure the monitoring parameters for the device. Monitoring is off by default. To enable the VPN Monitor in Realtime Monitor to display statistics for the VPN tunnel, configure the following:

- VPN Monitor—When enabled, the device sends ICMP echo requests (pings) through the tunnel at specified intervals (configurable in seconds) to monitor network connectivity (the device uses the IP address of the local outgoing interface as the source address and the IP address of the remote gateway as the destination address). If the ping activity indicates that the VPN monitoring status has changed, the device triggers an SNMP trap; VPN Monitor (in RealTime Monitor) tracks these SNMP statistics for VPN traffic in the tunnel and displays the tunnel status.
- Rekey—When enabled, the device regenerates the IKE key after a failed VPN tunnel attempts to reestablish itself. When disabled, the device monitors the tunnel only when the VPN passes user-generated traffic (instead of using device-generated ICMP echo requests). Use the rekey option to:
 - Keep the VPN tunnel up even when traffic is not passing through.
 - Monitor devices at the remote site.

- Enable dynamic routing protocols to learn routes at a remote site and transmit messages through the tunnel.
- Automatically populate the next-hop tunnel binding table (NHTB table) and the route table when multiple VPN tunnels are bound to a single tunnel interface.
- Optimized—When enabled, the device optimizes its VPN monitoring behavior as follows:
 - Considers incoming traffic in the VPN tunnel as ICMP echo replies. This reduces false alarms that might occur when traffic through the tunnel is heavy and the echo replies cannot get through.
 - Suppresses VPN monitoring pings when the tunnel passes both incoming and outgoing traffic. This can help reduce network traffic.
- Source Interface and Destination IP—Configure these options to use VPN Monitoring when the other end of the VPN tunnel is not a security device. Specify the source and destination IP addresses.

Adding a VPN Rule

After you have configured the VPN on each device you want to include in the VPN, you can add a VPN rule to a security policy:

- For policy-based VPNs, you must add a VPN rule to create the VPN tunnel.
- For route-based VPNs, the VPN tunnel is already in place. However, you might want to add a VPN rule to control traffic through the tunnel.

For details on adding and configuring a VPN rule in a security policy, see [“Adding VPN Rules” on page 661](#).

Creating L2TP VPNs

Creating device-level L2TP VPN is a three stage process:

- [Adding L2TP Users on page 659](#)
- [Configuring L2TP on page 660](#)
- [Adding a VPN Rule on page 660](#)

Adding L2TP Users

For VPNs that use L2TP to provide remote access services, you must add an L2TP User to the security device. An L2TP User has an account on the security device that guards the protected resources in the VPN; when the user attempts to connect to a protected resource, the security device authenticates the user.

To add a L2TP User for a security device, in the security device configuration L2TP/XAuth/Local User, click the Add icon. Enter a name for the user, then specify:

- User—Select a preconfigured Local User object that is configured for L2TP.
- Remote Setting—Select a preconfigured Remote Settings object.

- IP Pool—Select a preconfigured IP Pool object.
- Static IP—Enter the static IP address of the Local User.

Configuring L2TP

To connect to an L2TP VPN tunnel, the L2TP RAS user uses the IP address and WINS/DNS information assigned by the user's ISP. However, when the L2TP RAS user sends VPN traffic through the tunnel, the security device assigns a new IP address and WINS/DNS information that enables the traffic to reach the destination network.

Enter a name for the L2TP VPN, then specify the following information:

- Host Name—Enter the name of the L2TP host.
- Outgoing Interface—The outgoing interface is the interface on the security device that sends and receives VPN traffic. Typically, the outgoing interface is in the untrust zone.
- Keep Alive—The number of seconds a VPN member waits between sending hello packets to an L2TP RAS user.
- Peer IP—Enter the IP address of the L2TP peer.
- Secret—Enter the shared secret that authenticates communication in the L2TP tunnel.
- Remote Settings—Select the preconfigured remote settings object that represents the DNS and WINS servers assigned to L2TP RAS users after they have connected to the tunnel.
- IP Pool Name—Select the preconfigured IP pool object that represents the available IP addresses that can be assigned to L2TP RAS users after they have connected to the tunnel.
- Auth Server
 - Use the default settings to use the default authentication server for the domain. To change or assign a domain authentication server, edit the domain settings.
 - Use custom settings to specify a preconfigured authentication server object to assign TCP/IP settings to the gateway and authenticate specific L2TP User or User Groups.

Adding a VPN Rule

After you have configured the VPN on each device you want to include in the VPN, you can add a VPN rule to a security policy:

- For policy-based VPNs, you must add a VPN rule to create the VPN tunnel.
- For route-based VPNs, the VPN tunnel is already in place. However, you might want to add a VPN rule to control traffic through the tunnel.

For details on adding VPN rules to a security policy, see [“Adding VPN Rules” on page 661](#).

Creating L2TP Over Autokey IKE VPNs

Creating a device-level L2TP-over-Autokey IKE VPN is a multi-stage process:

1. Add L2TP Users (see [“Adding L2TP Users” on page 659](#))
2. Configure L2TP Settings (see [“Configuring L2TP” on page 660](#))
3. Configure Peer Gateway (see [“Configuring Gateways” on page 649](#))
4. Configure Routes (Route-based only) (see [“Configuring Routes \(Route-based only\)” on page 653](#))
5. Add VPN to Device (see [“Configuring the VPN” on page 653](#))
6. Add VPN rules to security policy (see [“Adding a VPN Rule” on page 660](#))

Adding VPN Rules

To create a policy-based VPN or to add access policies to a route-based VPNs, you must add a VPN rule to a security policy for each device in the VPN.

Adding a VPN Rule is a three stage process:

- [Configuring the VPN on page 661](#)
- [Configuring the Security Policy on page 662](#)
- [Assign and Install the Security Policy on page 662](#)

Configuring the VPN

In Security Policies, select a predefined security policy (or create a new policy), and add a VPN rule. right-click in the Source Address, Destination Address, Action, or Install On column and select Configure VPN to display the Configure VPN dialog box.

1. Select the source security device that contains the termination interface for the VPN tunnel.
2. Select a VPN Type:
 - For IKE VPNs, select the VPN that you configured on the device.
 - For L2TP VPNs, you must also select the L2TP tunnel that you configured on the device.
3. Select the Protected Resources for the VPN:
 - If both VPN termination points are security devices, choose the protected resources that represent the network components you want to protect. You can also select a predefined Global MIP or VIP for the device.
 - If the source VPN termination point is a RAS user, select Source is Dialup and choose the Protected Resources behind the destination VPN termination point that represent the network components you want to protect on the remote network.
 - If the destination VPN termination point is a RAS user, select Destination is Dialup and choose the Protected Resources behind the source VPN termination point that represent the network components you want to protect on the local network.

Configuring the Security Policy

To configure the remaining columns for the VPN rule:

- From Zone—Select the zone on the source VPN member that contains the termination interface for the VPN tunnel.
- To Zone—Select the zone on the destination VPN member that contains the termination interface for the VPN tunnel.
- Service column—Select the services you want to permit in the VPN tunnel.

You do not need to configure the action—NSM automatically defines the action as tunnel. You can also configure traffic shaping, options, authentication, antivirus, or attack protection for the VPN Rule.

To deny a host, use a deny rule before the VPN rule.

Assign and Install the Security Policy

You must assign the security policy to the each VPN member and install the security policy on those devices before the VPN is active.

Device-Level VPN Examples

This section provides examples of the two device-level VPN types:

- [“Example: Configuring a Route-Based Site-to-Site VPN, Manual Key” on page 662](#)
- [“Example: Configuring a Policy-Based Site-to-Site VPN, Manual Key” on page 668](#)
- [“Example: Configuring a Policy-Based RAS VPN, L2TP” on page 670](#)

The following sections provide step-by-step instructions on creating each type of device-level VPN.



NOTE: For examples on creating other VPN types using VPN Manager, see [“VPN Manager Examples” on page 632](#).

Example: Configuring a Route-Based Site-to-Site VPN, Manual Key

In this example, a Manual Key tunnel provides a secure communication channel between offices in Tokyo and Paris. The Trust zones at each site are in NAT mode. The Trust and Untrust security zones are in the trust-vr routing domain, and the Untrust zone interface (ethernet3) serves as the outgoing interface for the VPN tunnel.

To set up the tunnel, you must configure the security devices at both ends of the tunnel. First, you create the VPN components that you use to build the VPN, such as the security devices and the shared address objects. Next, you create the tunnel interfaces for each device and configure the VPN tunnel. You must also add the necessary static routes on each device to create the VPN tunnel. Finally, you create firewall rules in a security policy to control VPN traffic between the two sites.

1. Add the Tokyo and Paris security devices.
2. Configure the Tokyo device with the following interfaces:
 - Ethernet1 is the Trust IP (10.1.1.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (1.1.1.1/24) in the Untrust zone.
3. Configure the Paris device with the following interfaces:
 - Ethernet1 is the Trust IP (10.2.2.1/24) in the Trust zone.
 - Ethernet3 is the Untrust IP (2.2.2.2/24) in the Untrust zone.

Next, you create the address objects that you use in the VPN rule in the firewall rulebase (for details on creating VPN rules, see [“Adding VPN Rules” on page 661](#)).

4. Add the Tokyo Trust LAN (10.1.1.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
 - For Name, enter **Tokyo Trust LAN**.
 - For IP Address/Netmask, enter **10.1.1.0/24**.
 - For Color, select **magenta**.
 - For Comment, enter **Tokyo Trust Zone**.
5. Add the Paris Trust LAN (10.2.2.0/24) as a network address object. In Address Objects, click the Add icon and select **Network**. Configure the following, then click **OK**:
 - For Name, enter **Paris Trust LAN**.
 - For IP Address/Netmask, enter **10.2.2.0/24**.
 - For Color, select **magenta**.
 - For Comment, enter **Paris Trust Zone**.
 - Configure the Tokyo tunnel interface:
6. In the navigation tree, select **Device Manager > Security Devices**, then double-click the Tokyo device to open the device configuration.
7. In the device navigation tree, select **Network > Interface**. Click the Add icon and select Tunnel Interface. The General Properties screen for tunnel.1 appears.
8. Configure the following, then click **OK**:
 - For Zone, select **untrust**.
 - For Zone, select **untrust**.
 - For IP Options, select **Unnumbered**.
 - For Source Interface, select **ethernet3**.
 - Create the Tokyo VPN:
9. In the device navigation tree, select **VPN Settings > AutoKey IKE/Manual VPN**.

10. Select the Manual tab, then click the Add icon. The Properties screen appears. Configure the Properties tab as shown below:
 - For Name, enter **Tokyo_Paris**.
 - For Gateway, enter **2.2.2.2**.
 - For Local SP, enter **3020**.
 - For Remote SPI, enter **3030**.
 - For Outgoing Interface, select **ethernet3**.
 - For ESP/AH, select **ESP CBC**.
 - For Encryption Algorithm, select **3DES-CBC**.
 - Select **Generate Key by Password**, then enter the password **asdlk24234**.
 - For Authentication Algorithm, select **SHA-1**.
 - Select **Generate Key by Password**, then enter the password .
11. Select the Binding tab. Enable Tunnel Interface, then select tunnel1.
12. Click **OK** to save the new VPN.
13. Create Tokyo Routes:
14. In the device navigation tree, select **Network > Virtual Router** to display the list of virtual routers on the device. Double-click the trust-vr route to open the vr for editing.
15. In the virtual router dialog box, click **Routing Table**, then click the Add icon under destination-based Routing Table to add a new static route.



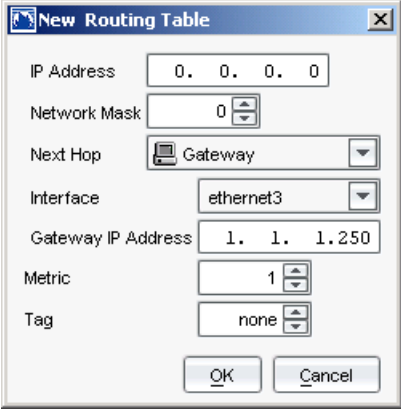
.....

NOTE: ScreenOS 5.0 devices display both destination-based and source-based routing tables. ScreenOS 5.1 and later devices display destination-based, source-based, and source interface-based routing tables.

.....

16. Configure a route from the untrust interface to the gateway, and then click **OK**.

Figure 105: Configure Tokyo Route for RB Site-to-Site VPN, MK



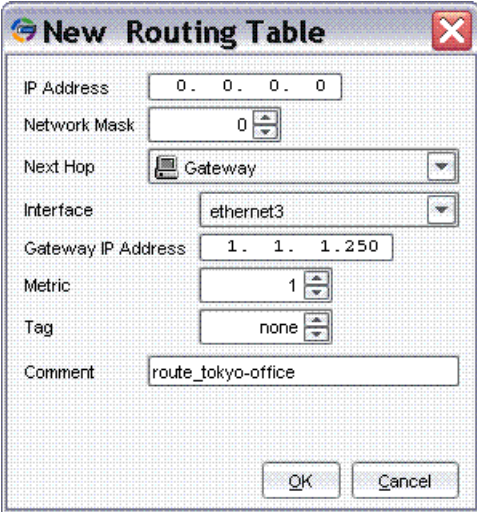
The dialog box titled "New Routing Table" contains the following fields and values:

- IP Address: 0. 0. 0. 0
- Network Mask: 0
- Next Hop: Gateway
- Interface: ethernet3
- Gateway IP Address: 1. 1. 1.250
- Metric: 1
- Tag: none

Buttons at the bottom: OK, Cancel.

17. Configure route from the trust zone to the tunnel interface, and then click **OK**.

Figure 106: Configure Tokyo Trust Route for RB Site-to-Site VPN, MK



The dialog box titled "New Routing Table" contains the following fields and values:

- IP Address: 0. 0. 0. 0
- Network Mask: 0
- Next Hop: Gateway
- Interface: ethernet3
- Gateway IP Address: 1. 1. 1.250
- Metric: 1
- Tag: none
- Comment: route_tokyo-office

Buttons at the bottom: OK, Cancel.

Your routing table should appear.

Figure 107: View Tokyo Routing Table for RB Site-to-Site VPN, MK

Routing Table				
Destination-based Routing Table				
<div> <div>+</div> <div>⊞</div> <div>⊞</div> </div>				
	IP Address /	Mask	Next Hop	Vsys
	0.0.0.0	0	Interface: ethernet3 Gateway IP Address: 1.1.1.250 Metric: 1 Tag:	...
	10.2.2.0	24	Interface: tunnel.1 Gateway IP Address: 0.0.0.0 Metric: 1 Tag:	...

- Click **OK** to save your changes to the virtual router, then click **OK** to save your changes to the Tokyo device.

Configure the Paris Tunnel Interface:

- In Device Manager, double-click the device icon for Paris to open the device configuration.
- In the device navigation tree, select **Network > Interface**. Click the Add icon and select **Tunnel Interface**. The General Properties screen appears.
- Configure the following, then click **OK**:
 - For Zone, select **untrust**.
 - For IP Options, select **Unnumbered**.
 - For Source Interface, select **ethernet3**.
- Create the Paris VPN:
 - In the device navigation tree, select **VPN Settings > AutoKey IKE/Manual VPN**.
 - Select the Manual tab, then click the Add icon. The Properties screen appears.
- Configure the following:
 - For Name, enter **Paris_Tokyo**.
 - For Gateway, enter **2.2.2.2**.
 - For Local SP, enter **3020**.
 - For Remote SPI, enter **3030**.
 - For Outgoing Interface, select **ethernet3**.
 - For ESP/AH, select **ESP CBC**.

- For Encryption Algorithm, select **3DES-CBC**, then select **Generate Key by Password** and enter the password **asdlk24234**.
 - For Authentication Algorithm, select **SHA-1**, then select **Generate Key by Password** and enter the password **PNas134a**.
6. Select the Binding tab. Enable Tunnel Interface, then select tunnel 1.
 7. Click **OK** to save the new VPN.

Create Paris Routes.

1. In the device navigation tree, select **Network > Virtual Router** to display the list of virtual routers on the device.
2. Double-click the trust-vr route to open the vr for editing.
3. In the virtual router dialog box, click **Routing Table**, then click the Add icon under destination-based Routing Table to add a new static route.



NOTE: ScreenOS 5.0.x devices display both destination-based and source-based routing tables; ScreenOS 5.1 and later devices display destination-based, source-based, and source interface-based routing tables.

4. Configure a route from the untrust interface to the gateway, then click **OK**:
5. Configure route from the trust zone to the tunnel interface, then click **OK**:
Your routing table should appear.
6. Click **OK** to save your changes to the virtual router, then click **OK** to save your changes to the Paris device.

Create the security policy:

1. In the main navigation tree, select **Policies**. Click the Add icon to display the New Security Policy dialog box.
2. Configure the following, then click **OK**:
3. For Security Policy Name, enter **Corporate Route-based VPNs**.

4. Optionally, add comments.
5. In the main navigation tree, select **Policies > Corporate Route-based VPNs**. The security policy appears in the main display area.

Figure 108: Configure Rules for RB Site-to-Site VPN, MK

No.	Match					Action	Install On
	From Zone	Source	To Zone	Destination	Service		
1	trust	Tokyo Trust LAN	untrust	Paris Trust LAN	ANY any	permit	Tokyo
2	untrust	Paris Trust LAN	trust	Tokyo Trust LAN	ANY any	permit	Tokyo
3	trust	Paris Trust LAN	untrust	Tokyo Trust LAN	ANY any	permit	Paris
4	untrust	Tokyo Trust LAN	trust	Paris Trust LAN	ANY any	permit	Paris

Example: Configuring a Policy-Based Site-to-Site VPN, Manual Key

In this example, a Manual Key tunnel provides a secure communication channel between offices in Tokyo and Paris, using ESP with 3DES encryption and SHA-1 authentication. The Trust zones at each site are in NAT mode. The Trust and Untrust security zones and the Untrust-Tun tunnel zones are in the trust-vr routing domain. The Untrust zone interface (ethernet3) serves as the outgoing interface for the VPN tunnel.

To set up the tunnel, you must configure the security devices at both ends of the tunnel. First, you create the VPN components that you use to build the VPN, such as the security devices and the shared address objects. Next, you configure the VPN tunnel and add the necessary static routes on each device. Finally, you create VPN rules in a security policy to create the VPN tunnel between the two sites.

Create VPN Components

1. Security Devices.
2. Address Objects.

Create the Tokyo VPN:

1. In the device navigation tree, select **VPN Settings > AutoKey IKE/Manual VPN**.
2. Select the Manual tab, then click the Add icon. The Properties screen appears. Configure the following:
 - For Name, enter **Tokyo_Paris**.
 - For Gateway, enter **2.2.2.2**.
 - For Local SP, enter **3020**.
 - For Remote SPI, enter **3030**.
 - For Outgoing Interface, select **ethernet3**.
 - For ESP/AH, select **ESP CBC**.

- For Encryption Algorithm, select **3DES-CBC**.
- Select **Generate Key by Password**, then enter the password **asdlk24234**.
- For Authentication Algorithm, select **SHA-1**.
- Select **Generate Key by Password**, then enter the password **PNas134a**.
- Select the Binding tab. Enable Tunnel Zone and select **untrust-tun**.
- Click **OK** to save the new VPN.

3. Create Tokyo Routes.

Create the Paris VPN

1. In the device navigation tree, select **VPN Settings > AutoKey IKE/Manual VPN**.
2. Select the Manual tab, then click the Add icon. The Properties screen appears.
3. Configure the following:
 - For Name, enter **Paris_Tokyo**.
 - For Gateway, enter **2.2.2.2**.
 - For Local SP, enter **3020**.
 - For Remote SPI, enter **3030**.
 - For Outgoing Interface, select **ethernet3**.
 - For ESP/AH, select **ESP CBC**.
 - For Encryption Algorithm, select **3DES-CBC**, then select **Generate Key by Password** and enter the password **asdlk24234**.
 - For Authentication Algorithm, select **SHA-1**, then select **Generate Key by Password** and enter the password **PNas134a**.
4. Select the Binding tab. Enable Tunnel Zone and select **untrust-tun**.
5. Click **OK** to save the new VPN.
6. Create Paris Routes.

Create the security policy

1. In the main navigation tree, select **Policies**. Click the Add icon to display the new Security Policy dialog box.
2. Configure the following, then click **OK**:
 - For Security Policy Name, enter **Corporate Policy-Based VPN**.
 - Optionally, enter comments.

3. In the main navigation tree, select **Policies > Corporate** Policy-Based VPN. The security policy appears in the main display area. Configure two VPN rules.
 - Rule 1 creates the VPN tunnel from the Tokyo device to the Paris device.
 - Rule 2 creates the VPN tunnel from the Paris device to the Tokyo device.
4. Save the security policy.

Example: Configuring a Policy-Based RAS VPN, L2TP

In this example, you create a RAS user group called Field Sales and configure an L2TP tunnel called Sales_Corp, using ethernet3 (Untrust zone) as the outgoing interface for the L2TP tunnel. The security device applies the default L2TP tunnel settings to the RAS user group.



NOTE: An L2TP-only configuration is insecure and is recommended only for debugging.

The remote L2TP clients are on Windows 2000 operating systems. For information on how to configure L2TP on the remote clients, refer to Windows 2000 documentation. Only the configuration for the security device end of the L2TP tunnel is provided below.

Configure the user and remote settings objects

1. Configure an L2TP user object for Adam, then click **OK**:
 - For Name, enter **Adam**.
 - Select **Enable**, then select **L2TP**.
 - Select **Password**, then enter and confirm the password: **AJbioJ15**.
2. Configure an L2TP user object for Betty, then click **OK**:
 - For Name, enter **Betty**.
 - Select **Enable**, then select **L2TP**.
 - Select **Password**, then enter and confirm the password: **BviPsoJ1**.
3. Configure an L2TP user object for Carol, then click **OK**:
 - For Name, enter **Carol**.
 - Select **Enable**, then select **L2TP**.
 - Select **Password**, then enter and confirm the password: **Cs10kdD3**.
4. Create a local user group called Field Sales that includes the Adam, Betty, and Carol local user objects.
5. Configure the following, then click **OK**:
 - For Name, enter **RM_L2TP**.
 - For Color, select **green**.

- For Dns1, enter **1.1.1.2**.
- For Dns2, enter **1.1.1.3**.
- For Wins1, enter **0.0.0.0**.
- For Wins2, enter **0.0.0.0**.
- Configure the IP Pool object. Configure the following, then click **OK**:
 - For IP Pool Name, enter **Global**.
 - For Color, select **magenta**.
 - For Start IP, enter **10.10.2.100**.
 - For End IP, enter **10.10.2.180**.

Configure the L2TP Tunnel

1. In Device Manager, double-click the device icon for the device on which you want to configure the L2TP tunnel.
2. In the device navigation tree, select **VPN Settings > L2TP**. In the main display area, click the Add icon. The null-L2TP tunnel dialog box appears.
3. Configure the following, then click **OK**:
 - For Name, enter **Sales_Corp**.
 - For Outgoing Interface, select **ethernet3**.
 - For Keep Alive, enter **60**.
 - For Peer IP, enter **0.0.0.0** (because the peer's ISP dynamically assigns it an IP address, enter 0.0.0.0 here).
 - Select **Use Custom Settings**, and leave the default authentication server as Local.
 - For User/Group, select **Dialup Group**, then select **Field Sales**.
4. Click **OK** to save your changes to the device.
5. Configure a rule in the Zone Rulebase of a security policy.

Auto-Connect Virtual Private Network

Hub-and-spoke configurations are deployed in large enterprises. Each branch site (spoke) is connected to a central site (hub). The communication between spoke sites must go through the hub, which does not scale as the number of spoke sites increases.

Using the auto-connect virtual private network (ACVPN) feature in devices running ScreenOS 6.0 and later, you can configure your hub-and-spoke network so that spokes dynamically create VPN tunnels between each other as needed. The dynamic tunnels time out when traffic ceases to flow through them, freeing network administrators from the time-consuming task of maintaining a complex network of static VPN tunnels.

With ACVPN, all spokes are connected to the hub by VPN tunnels. All VPN tunnels configured towards the hub must be route based. After you set up a static VPN tunnel between the hub and each of the spokes, you configure ACVPN, and then enable the Next Hop Resolution Protocol (NHRP).

Configuring ACVPN

You can configure ACVPN using VPN Manager.

To configure auto-connect VPN, perform the following steps:

1. Create a route-based auto-key IKE VPN.
 - In the main navigation tree, click **VPN Manager > VPNs**.
 - Click the Add icon and then select **AutoKey IKE VPN**. The New AutoKey IKE VPN dialog box appears.
 - Select the **Enable** check box to enable the VPN.
 - Select the default termination zone for the VPN tunnels from the Terminal Point drop-down list.
 - Select the type of VPN from the Type drop-down list, and then click **OK**.
2. Add the devices to the hub-and-spoke topology.
 - In the navigation tree, click **Security Devices**. The Security Devices dialog box appears.
 - Click the Add icon, and then select the devices to be included in the hub- and-spoke topology.
 - Click **OK**.
3. Configure the topology.
 - In the general configuration area of the VPN Manager, click the **Topology** link. The New Topology dialog box appears.
 - Select the device to be the hub for the topology from the Hub drop-down list.
 - Select the **Enable Auto-Connect VPN** check box.
 - In the Mains pane, select each device you want to be a main. Main devices can communicate with every other device in the topology.
 - Click **OK**, and then click the **Save** button to save the VPN configuration.
4. Configure the IP address for the tunnel interface on the hub and spokes.
 - In the configuration area of this VPN, click the **Device Tunnel Summary** link. A list of tunnels appears.
 - Right-click the tunnels and assign the IP address. The Tunnel Interface dialog box appears.
 - Enter the IP address and netmask, and then click **OK**.
 - Click the **Close** button in the AutoKey IKE VPN.

5. Specify the AutoConnect VPN parameters to complete the procedure.
 - In the configuration area of this VPN, click the **AutoConnect VPN Parameters** link. The AutoConnect VPN Parameters dialog box appears.
 - Click the **Import Gateway and AutoKey Parameters** button to import the existing hub-and-spoke configuration for the hub and spoke. You can configure the VPN and gateway by using ACVPN-Dynamic(Mains) or ACVPN-Profile(Hub) parameters in the navigation tree.
 - Click the **Save** button to save the VPN configuration.
6. Verify the NHRP settings for the hub and spokes, virtual router, and spoke virtual router.
7. For interface NHRP settings:
 - In the configuration area of this VPN, click the **Device Tunnel Summary** link.
 - Right-click the tunnels, and then click **Edit Interface**. The Tunnel Interface dialog box appears.
 - Click **Protocol > NHRP**. Ensure that the Enable NHRP check box is selected.
 - Click **OK**.
8. For the hub virtual router NHRP settings:
 - In the configuration area of this VPN, click the **Device Tunnel Summary** link.
 - Right-click the tunnels, and then click **Edit Virtual Router**. The Virtual Router dialog box appears.
 - Click **Dynamic Routing Protocol > NHRP > Parameters**.
 - Verify that the ACVPN-Profile setting has been populated.
 - Click **OK**.
9. For the spoke virtual router NHRP settings:
 - In the configuration area of this VPN, click the **Device Tunnel Summary** link.
 - Right-click the tunnels, and then click **Edit Virtual Router**. The Virtual Router dialog box appears.
 - Click **Dynamic Routing Protocol > NHRP > Parameters**.
 - Verify that the NHS IP Address field has been populated.
 - Click **OK**.

IVE VPN Monitoring

NSM real-time monitoring is available on Secure Access and Infranet Controller devices. For more information, see [“Realtime Monitoring” on page 705](#).

CHAPTER 13

Central Manager

Central Manager provides super administrators with the opportunity to manage up to ten concurrent regional servers from the Network and Security Manager (NSM) management system. With Central Manager, you can log into the system and perform operations such as enforcing global policies, adding regional servers, adding, modifying, and deleting pre/post rules and shared objects, and managing polymorphic objects.

This chapter contains the following sections:

- [Central Manager Overview on page 675](#)
- [Installing Global Policy to a Regional Server on page 678](#)

Central Manager Overview

Predefined shared objects are shared by Central Manager and regional servers. Any predefined shared objects that are used by Central Manager pre/post rules are available in regional servers, attack db, and so on. When you update pre/post rules, the Central Manager and regional server versions must match.



NOTE: You cannot create a new pre/post rule on the central manager and push it to a regional server.

Regional Server and Central Manager Self-Sufficiency

Both the Central Manager and the regional servers are self-sufficient and independent from the other being online.

Self-Sufficient Central Manager

Central Manager is independent of the regional servers and can run without any regional server online. The Central Manager administrator can add, modify, or delete pre/post rules and shared objects. Data is not lost when logging on and off of Central Manager. In addition, Central Manager does not use any of the shared objects that exist only in any of the individual regional servers.

Self-Sufficient Regional Server

Regional servers can enforce global policies even when Central Manager is not running. Regional servers maintain copies of pushed Central Manager pre/post rules and associated objects. In addition, regional servers do not use any Central Manager data directly.

Super Admin User

The Central Manager super administrator accounts have access to all features offered in Central Manager. The super user is created while installing Central Manager. It is also the only administrator account created during installation.

As a super administrator user, you can use the single sign-on feature to access regional servers directly from Central Manager without logging out of Central Manager inputting regional server login credentials.

Regional Server Management

Central Manager treats regional servers as objects similar to other objects. As with other objects, Central Manager can add, modify, and delete regional servers. Regional server objects are shared objects that contain essential connection information such as its own IP address, port, and so on. Central Manager administrators can use additional credential information in the regional server objects to sign onto each regional server.

Once logged into a Central Manager server, super administrators can select any of the regional servers managed by Central Manager and begin managing the servers using all assigned permissions. No extra log on/off steps are required for administrators to navigate from one regional server to another or from Central Manager server to a regional server. Any regional server accessible through Central Manager, is opened using a separate window. There is a maximum number of 25 concurrent regional servers Central Manager can open at any one time.

Management Modes for J Series and SRX Series Devices

With J Series and SRX Series devices, the NSM Central Manager can operate in either central management or device management mode.

In central management mode, a device references the central policy manager and central manager objects. In device management mode, a device does not reference the central policy manager or a central manager object.

The following sections briefly summarize these differences. For detailed information, refer to the configuration manuals.

Central Management Mode

In default central management mode, a device has a link to a central policy manager. All firewall, VPN, and IDP policy information and policy related configurations (shared configurations such as addresses and services) are hidden from device editor view.

Policies from the central policy manager are shared across ScreenOS-based firewall devices, standalone IDP devices, and J Series devices.



NOTE: When a J Series device is managed in central manager mode, if you select an IDP rulebase rule and specify an IP address for the source and destination instead of “any,” the rule policy is not pushed to the device.

Device Management Mode

In device management mode, you can use the NSM Device Editor to manage the complete device configuration. A device does not reference a central policy or central manager object. The device is not affected by any VPN manager configurations setup within NSM for policy-based VPNs.

You can update the latest configuration to the device. Later, you can reimport the latest configuration from the device to determine any configuration differences and so on.

You can use the NSM UI to switch from the central management to device management mode. This operation disconnects the device from the policy manager, central object manager, and the VPN manager.



NOTE: You cannot switch a J Series or SRX Series device from central management mode to device management mode if the device has an assigned policy.

Using Central Manager

This section provides procedures for the following tasks for Central Manager super users:

- “Adding a Regional Server Object” on page 677
- “Deleting a Regional Server Object” on page 678
- “Logging into a Regional Server” on page 678

Adding a Regional Server Object

For a Central Manager administrator to log onto a regional server, one of the regional server administrator credentials must be used and linked to the Central Manager administrator.

The following procedure assumes that a Central Manager administrator is logged onto a Central Manager client.

To add a regional server object:

1. In the main navigation tree, select **Object Manager > Regional Server**.
2. Click the **Add** icon in the toolbar.
3. Enter the following information for the regional server you want to add.
 - Name

- IP address
 - Backup IP address (optional)
4. Click the **Add** icon in the administrator table to add regional server admins.
 5. Enter the following information for the regional server you want to add.
 - Admin name
 - Admin password
 6. Click the **Add** icon in the administrator table to link the newly created regional server object to the Central Manager administrator.
 7. Select a Central Manager administrator from the drop down list box.
 8. Click **OK** in all the open boxes to save the options and close each window.

Deleting a Regional Server Object

The following procedure assumes that a Central Manager administrator is logged onto a Central Manager client and a regional server object has been created.

To delete a regional server object:

1. In the main navigation tree, select **Object Manager > Regional Server**.
2. Right-click the regional server you want to delete.
3. In the **Delete Regional Server** dialog box, click **Next** to delete the object.

Logging into a Regional Server

Central Manager administrators can log into regional servers directly from Central Manager.

The following procedure assumes that a Central Manager administrator is logged onto a Central Manager client and a regional server object has been created.

To switch from Central Manager to a regional server:

1. In the toolbar, click the **Login to Regional Server** drop down list.
2. Select a regional server to open to launch the selected regional server in a separate window.

Only a Central Manager administrator can log into any other regional servers. A regional server administrator cannot log into another regional server or a central manager server.

Installing Global Policy to a Regional Server

During the Global Policy Install on the Central Manager server, all pre/post rules as well as the global polymorphic and shared objects on the Central Manager server are updated to regional servers managed by Central Manager. The Central Manager administrator

can select which regional servers will receive the Central Manager rules and objects during the install.

Prerule and Postrule Updates during Global Policy Install

Pre/post rules exist on the Central Manager server under a separate policy object in the Policy Manager. When updating to the regional server during the Global Policy Install these rules are created in the regional server under a new global policy in the global domain. All global policies existing on the regional server prior to the Global Policy Install are removed, and replaced completely with the new global policies from Central Manager. Only those global rules which have the regional server object included in the Install On column, as well as rules with the Any entry in the Install On column, will be updated to that regional server.

Global Policy Install is a directive that is dispatched as a job with many tasks. Each task represents an install operation to a specific server. Job manager reports the status of each task. Each task is executed independently of other task's status. If one task fails, as well as Job Manager entries, this failure does not prevent another task from finishing successfully. Audit log entries are generated on both the Central Manager and regional servers.

Shared Objects Update During Global Policy Install

All shared objects (both polymorphic and regular) existing on the Central Manager server are updated to regional servers during the Global Policy Install if they are referenced in a pre/post rule. Note that objects are updated only if they are actually being used by the pre/post rules on the Central Manager server. All new shared objects are replicated/inserted into the global domain of the regional server. Objects that are not used are not updated.

Name Space Conflict Resolution for Shared Objects

When a regular shared object is replicated to a regional server during the Global Policy Install, the following name conflict scenarios could occur:

- Conflict with a regional server shared object of the same name and same type content—Existing shared objects will be kept and the incoming shared object will be discarded. The incoming global policy rules use the existing shared object.
- Conflict with a regional server shared object of the same name, but different content—An attempt is made to match the content of an incoming shared object with another shared object named "objname_n" where "n" is a sequentially increasing integer. The incoming global policy rules use the newly created shared object.
- Conflict with the previously replicated polymorphic object of the same type—The incoming shared object is renamed "objname_n" where "n" is a sequentially increasing integer and is inserted into the regional server's global domain.

Since polymorphic objects cannot be deleted by the regional server administrator, some of the polymorphic objects that exist in the global domain of the regional server are deleted as a first step in the Global Policy Install transaction.

All polymorphic objects are deleted if they are not used by any of the local policies in the regional server.

Name Space Conflict Resolution for Polymorphic Objects

When a polymorphic object is replicated to a regional server during the Global Policy Install, the following name conflict scenarios could occur:

- Name conflict with a previously replicated polymorphic object—To keep the customization information the regional server administrator added, existing polymorphic object are kept, and incoming global policy rules use existing polymorphic object. Incoming polymorphic object with the same name are discarded.
- Name conflict with a regional server regular shared object of the same type—The incoming polymorphic object is renamed “objname_n” where “n” is a sequentially increasing integer and inserted into the regional server’s global domain.

Only names are pushed for polymorphic objects.

CHAPTER 14

Topology Manager

- [Overview of the NSM Topology Manager on page 681](#)
- [About the NSM Topology Manager Toolbar on page 682](#)
- [Initiating a Topology Discovery on page 683](#)
- [Viewing a Network Topology on page 684](#)
- [About the NSM Topology Map Views on page 684](#)
- [About the NSM Topology Table Views on page 686](#)
- [About Topology Manager Preferences on page 688](#)
- [Adding Discovered Devices to NSM on page 688](#)

Overview of the NSM Topology Manager

- [About the NSM Topology Manager on page 681](#)
- [Requirements for a Topology Discovery on page 681](#)

About the NSM Topology Manager

The NSM Topology Manager is a tool provided in the NSM GUI that allows you to discover and manage the physical topology of a network of devices connected to a Juniper Networks EX Series Ethernet switch. The network can include J Series, M Series, MX Series, and EX Series devices, as well as ScreenOS and IDP devices, IP phones, desktops, printers, and servers. The Topology Manager also provides details about connections between a device and the EX Series switch.

Requirements for a Topology Discovery

To use the Topology Manager, first add one or more EX Series switches to the device manager in NSM. You can then use an added device as a seed device in initiating a topology discovery. Alternatively, if there are no devices added or managed in NSM, you can initiate a topology discovery by configuring preferred subnets; with this method, NSM discovers all the IP addresses in the included subnets range.

In addition to having either a seed device or configuring preferred subnets, you also need the following to initiate topology discovery:

1. The management IP address of the EX Series switch that acts as the seed IP address

2. SNMP credentials:
 - For SNMPv1 and SNMPv2c: Community string
 - For SNMPv3: Username, security level, authentication type, privacy type, privacy password, and authentication password
3. Enabled Layer 2 protocols such as LLDP, STP, RSTP in the switched network, because network discovery depends on these as well as the Address Forwarding Table information.

About the NSM Topology Manager Toolbar

You can use the Topology Manager toolbar to perform the following actions:

- **Zoom in and Zoom out:** Use these tools to view the network topology according to the detail required. These tools are only of use in the map view.
- **Save to file:** Use this tool to save the network topology map as an image file and the devices and links tables as text files from their respective views.
- **Print:** From different views, you can use this tool to print a network topology map as an image file and the devices and links tables as text files.
- **Manage Devices:** Use this tool to select one or more devices from a topology map and manage them in NSM. This tool is applicable only to map views and not the different table views. To add a device:
 - a. Select the **Manage Devices** icon. A dialog opens.
 - b. Enter the SSH user name and password.
 - c. Select **OK**.
- **Set Preferences:** Use this tool to set preferences according to which the discovery engine can perform a topology discovery. You can set preferences for default SNMP credentials, topology discovery intervals, and subnets to be included or excluded.
- **Start and Stop Topology Discovery:** Use these tools to initiate and stop a topology discovery based on the set of seed devices and credentials specified in the topology preferences.
- **Search:** You can search for a device, end point device, link, or port in any of the table views by providing a string in the search text box. NSM performs a sub string match against all attributes of the particular view and displays the results in the same table. If you navigate to another tab, your search results are lost. You can save the search output in a text file as comma-separated values.

The Topology Manager status bar at the bottom of the screen indicates the time stamp of the last completed topology discovery and whether a discovery is in progress.

Initiating a Topology Discovery

You can initiate network topology discovery in two ways.

- Select **Discover** (represented by a green arrow) from the Topology Manager toolbar. NSM displays the network topology in whichever view you choose.
- Select **Preferences** (represented by red check marks on a page) from the Topology Manager toolbar. The Topology Manager Preferences window opens.
 1. Select the **Refresh Intervals** sub-tab.
 2. Check **Run Topology discovery at regular intervals** and specify the interval or the time of the day to run the discovery engine.

The other settings for a successful discovery are the default SNMP credentials and optionally, the subnets to be included or excluded. You can adjust these settings in the **Default Credentials** and **Preferred SubNets** sub-tabs of the Topology Manager Preferences window.

To ensure a successful topology discovery, it is recommended that you:

- Use SNMP V1 or V2C versions.
- Ensure that UDP ports from 50001 to 50010 are free because Topology Discovery also uses these ports besides the standard SNMP UDP ports 161 and 162.
- Ensure that the EX Series switches in the network run Junos OS EX 9.3R2 or later.
- Ensure that LLDP or LLDP-MED is enabled on all switches and switch ports, as well as on all LLDP or LLDP-MED enabled devices such as IP Phones.
- Ensure that the included subnets specified in Topology Manager preferences are sufficient for all switches and routers and that they are SNMP enabled, in order that the maximum number of links are discovered.
- Check for NSM schema updates if some Juniper Networks devices are not discovered.
- Expand the range of the included subnets and ensure that all relevant routers are SNMP enabled if IP addresses for end-point devices connected to a switch are not discovered. Devices could go undiscovered if the router is not SNMP enabled or is not included in the subnets specified in the Topology Manager preferences.
- The broadest subnet mask recommended is 255.255.240.0 for any included subnet. The recommended limit of IP addresses in all included subnets put together is 15000.

The Topology Manager only discovers the product names of Juniper Networks devices. IC devices (version 3.0r1 onwards) and SA devices (6.4r1 onwards) are supported.

Vendor identification for non-Juniper Networks devices is based on the SNMP enterprise numbers registered with IANA. Therefore, it is possible that non-Juniper Networks discovered devices may not reflect most recent ownership changes caused by mergers, acquisitions or rebranding.



NOTE: Topology Discovery supports only IPv4 addresses. IPv6-based networks are not discovered.

Viewing a Network Topology

To view a network topology using the NSM Topology Manager:

1. Open the navigation tree within the **Configure** module of the NSM GUI.
2. Select (+) to expand the **Device Manager** branch.
3. Select **Topology**. You can choose between graphical and tabular views by selecting one of the sub-tabs in the display area.
4. Select the **Topology Map** sub-tab.
5. Select either a **SubNets** view or a **Groups** view of the topology from the navigation, to see the map in the main display area. You can use the mini map tool in the topology map to focus on particular areas in the topology map.
6. Select **Devices** to view a list of networking devices in the topology.
7. Select **End Point Devices** to view a list of end point devices such as IP phones, printers, wireless access points, PCs / desktops, etc. in the topology.
8. Select **Links** to view the links among network devices in the topology— both between network devices, as well as between network and end point devices.
9. Select **Free Ports** to view a list of EX Series switches and the available ports on these switches.

About the NSM Topology Map Views

The NSM Topology Manager provides both graphical and tabular views of your network topology. The map view is a graphic representation of all devices discovered in a network along with their linked elements.

In map view, each network element is represented by an icon indicating whether the element is a Juniper Networks product and whether it is managed by NSM. Each device type is represented by a unique icon on the map. Managed and unmanaged devices appear as different colored icons. The graphic view depicts physical connectivity between networking devices but does not show connectivity with endpoint devices. NSM offers two topology map views— SubNets and Groups.

- [SubNets View on page 685](#)
- [Groups View on page 685](#)
- [Menu Options in the Topology Map View on page 685](#)

SubNets View

The **SubNets** view offers you a complete topology view of the network by grouping together devices belonging to a subnetwork, representing them by a cloud icon. A default subnet contains devices that are not part of a specific subnet or with undiscovered subnet information. Double-click on a cloud icon to view devices within the particular subnet. Connections between devices in two different subnet clouds are depicted by a numbered link, that indicates the link count between devices in different subnet clouds.

Groups View

The **Groups** view shows groups of managed devices that have already been created; each group is represented by a cloud icon. You cannot assign devices and create groups from within the topology manager. Double-click on a cloud icon to view the devices within that particular group. Connections between devices in two different clouds are depicted by a numbered link, which indicates the link count between the devices in each cloud. Multiple discovered links that are part of a Link Aggregation Group (LAG) are displayed as a single distinctive link between the interfaces.

Menu Options in the Topology Map View

You can perform the following actions from the right-click menu in the topology map view.

- **Locate Devices:** Use this tool to locate a specific device within a particular topology view. To find a device within a topology cloud:
 1. Expand a topology cloud in the subnets or groups map view.
 2. Right-click in the map view and select the **Locate device** option.
 3. Provide the name, MAC address, or IP address of the device in the dialog.
 4. Select OK to locate the device with matching criteria in the map.
- **Locate Subnets:** Use this tool to locate a particular subnet cloud in the subnets topology map view.
 1. Open the subnets topology view.
 2. Right-click in the view and select the **Locate subnets** option.
 3. Provide the name or ID of the subnet in the dialog.
 4. Select OK to locate the subnet matching the given criteria.
- **Locate Groups:** Use this tool to locate a particular group cloud in the groups topology map view.

1. Open the Groups topology view.
 2. Select the **Locate Groups** option from the right-click menu.
 3. Provide the Group name in the popup dialog.
 4. Select OK to locate the group matching the given criteria.
- **Show Devices:** Use this menu option to view all the devices in a selected subnet cloud.
 - **Show Elements:** Use this option to view all the devices and groups within a selected group.
 - **Show Details:** Use this option to view the basic details about a device.
 - **Add:** Use this tool to add a device to NSM by selecting it from the topology map. This menu option is enabled for all unmanaged devices in NSM.
 - **Edit:** Use this tool to open the NSM configuration editor for a managed device. This option is enabled for all managed devices.
 - **Import:** Use this tool to import a device configuration to NSM. This option is only enabled for managed devices in NSM.
 - **Update:** Use this tool to update devices with new configuration. This option is only enabled for managed devices in NSM.



NOTE: The topology view is not automatically refreshed with every change in the actual network topology. While a discovery is in progress, the Topology Manager view is that of the last topology discovery. The Topology Manager view is refreshed only when the topology discovery is completed.

About the NSM Topology Table Views

The NSM Topology Manager provides both graphical and tabular views of your network topology. A tabular view of the topology lists all the network elements and devices connected to them. A tabular view does not display information related to the links and the types of links among various network elements in the topology.

In any of the table views, right-click on a device to view details corresponding to the table such as basic device details in the Devices table view, and detailed link details in the Links table view. NSM provides four different table views.

- [Devices View on page 687](#)
- [EndPoint Devices View on page 687](#)
- [Links View on page 687](#)
- [Free Ports View on page 687](#)

Devices View

The NSM Topology Manager provides a tabular view of all the discovered Juniper Networks devices in the network along with relevant details about each device. The **Devices** table lists details about the Juniper Network devices and other third party routers and switches. NSM indicates the managed status of a Juniper Networks device in the **Device Status** column. You can also track the status of your managed devices in the **Connection Status** and **Alarm Status** columns. The Devices view uses the same coloring scheme as the Device Manager.

EndPoint Devices View

The **End Point Devices** table provides detailed information about all the endpoint devices that exist in the current network topology. The table includes all endpoint devices connected with an EX Series switch, such as IP phones or desktops, but excludes all Juniper Networks devices, routers and switches. You can save the information in the table as comma-separated values in a file.

Links View

The **Links** table provides detailed information about all the links between devices discovered by the topology discovery engine. You can save the information in the table as comma-separated values in a file.

Multiple discovered Link Aggregation Group (LAG) links between two interfaces are displayed as a single link. When you launch the **Link Details** dialog box from the single displayed LAG link, you can view the aggregate ports on both devices.

Free Ports View

The **Free Ports** table lists all the free ports available on the devices discovered by the topology discovery engine. If the administrative status of a device port is down, it is considered a free port. The managed status of a Juniper Networks device is indicated in the **Device Status** column. You can save the information in the table as comma-separated values in a file.

You can right-click on a free port listed in the topology tabular view and launch a **Port Template** wizard. However, this feature is enabled only if the free port is discovered on an NSM-managed EX Series switch with its connection status up. The wizard allows you to view all the templates already applied on the device as well as those that can be applied to it. You can then select and apply a template to the device. For more details, see Chapter 13, Role-based Port Templates.



NOTE: The topology view is not automatically refreshed with every change in the actual network topology. While a discovery is in progress, the Topology Manager view is that of the last topology discovery. The Topology Manager view is refreshed only when the topology discovery is completed.

About Topology Manager Preferences

The topology discovery engine performs a topology discovery according to your preferences. Clicking **Preferences** in the tool bar opens the Topology Preferences Dialog. The topology preferences dialog has three tabs:

- [Default Credentials Tab on page 688](#)
- [Refresh Interval Tab on page 688](#)
- [Preferred Subnets Tab on page 688](#)

Default Credentials Tab

This tab contains a table where each row represents a set of credentials that you can add or delete. The topology discovery engine uses default SNMP settings in the absence of specified SNMP settings. You cannot delete credential SNMP version 1 or version v2c with the community string “public”, because NSM uses it as the default setting. When adding credentials, you must specify a community string if you choose versions SNMPv1 or SNMPv2c. If you choose version SNMPv3, then you also need to specify the following parameters:

- USM User Name
- Authentication Type and Password
- Privacy Type and Password

Refresh Interval Tab

This tab allows you to set the interval at which the topology discovery engine must discover the network topology. You can set a particular time of day or regular intervals. The time of your initial discovery serves as the basis of calculation for future discoveries.

Preferred Subnets Tab

This tab allows you to specify whether a particular subnet is to be explicitly included in the discovery. You can specify a list of allowed and denied subnets. You must specify included subnets because topology discovery happens only for those included subnets that you configure. Discovery does not take place if there are neither included subnets nor seed devices (managed devices). If you configure both included and excluded subnets, discovery happens only for included subnets and not for excluded subnets. When adding a subnet to either the allowed or denied list, specify the subnets base IP address and the subnets mask number.

Adding Discovered Devices to NSM

After a topology discovery, you can view a list of M Series, MX Series, J Series, EX Series, IDP, and ScreenOS devices that are not yet managed by NSM by selecting the **Manager Map** tab. You can add any of the devices on the list to NSM by following these steps.

1. Select and right-click on a device in the map. NSM launches a wizard to help you add devices to be managed.

2. Select one or more devices and specify SSH v2 credentials for them.
3. Verify the RSA key fingerprint for each of the devices.

The wizard detects each selected device and adds it to NSM. The wizard then imports the device configuration, hardware, software, and license inventory into NSM.

After adding a device to NSM, you can perform the following:

- View and edit device configuration: You can use the **Edit** menu to open the configuration editor to view and edit a device's configuration.
- Update device configuration: You can use the **Update** menu to update the changed configuration on the device.
- View device details in the topology map: You can view details of a managed device in the topology view.
- View link details between devices in the topology map: You can use the View details item on a selected link in the topology map to view link details between two managed devices, where one of the devices is the source and the other is the destination.

Role-based Port Templates

- [Using Role-Based Port Templates on page 691](#)

Using Role-Based Port Templates

Use port templates to apply a predefined configuration to interfaces in EX Series switches. Port templates apply recommended settings for security and class-of-service parameters. When applying a port template to an interface, you can specify values for parameters such as VLAN, Voice VLAN, and IP Address.

The port roles and the corresponding configuration are:

- Desktop plus Phone Port—Interface family is set to **ethernet-switching**, port mode is set to **access**, port security parameters (MAC limit =3; dynamic ARP Inspection, DHCP snooping enabled) are set, and recommended CoS parameters are specified for forwarding classes, schedulers, and classifiers.
- Desktop Port—Interface family is set to **ethernet-switching**, port mode is set to **access**, RSTP is enabled with the edge option, and port security parameters (MAC limit =1; dynamic ARP Inspection and DHCP snooping enabled) are set.
- Layer 2 Uplink Port—Interface family is set to **ethernet-switching**, port mode is set to **trunk**, port security is set to **dhcp-trusted**, and recommended CoS parameters are set for schedulers and classifiers.
- Routed Uplink Port—Port family is set to **inet**, and recommended CoS parameters are set for schedulers and classifiers.
- Wireless Access Point—Interface family is set to **ethernet-switching**, port mode is set to **access**, and RSTP is enabled with the edge option.

When you apply port templates on EX Series switches, NSM creates the required configuration in the following configuration groups and applies them at the top level configuration node:

- juniper-port-template-desktop
- juniper-port-template-desktop-phone
- juniper-port-template-layer2-uplink
- juniper-port-template-layer3-uplink

- [juniper-port-template-wireless-access-point](#)
- [juniper-port-template-cos-settings](#)

We strongly recommend that you do not change or delete these configuration groups. See [“Configuring Devices” on page 199](#) for more information.

1. [Managing Port Template Associations on page 692](#)
2. [Apply or Edit a Port Template on page 692](#)
3. [Detect and Resolve Configuration Conflicts on page 694](#)
4. [Clone a Port Template on page 694](#)
5. [Edit a Port Template on page 695](#)

Managing Port Template Associations

To manage port templates and the associated interfaces:

1. In the navigation tree, select **Device Manager > Port Templates**. The screen displays a list of the supported port templates.
2. Select a port template.
3. Click one:
 - **Apply/remove selected template on port**—Use this option to apply or remove the port template association with the interface. See [“Apply or Edit a Port Template” on page 692](#).
 - **Detect/resolve configuration conflicts**—Use this option to resolve conflicts between the port template configuration and the actual configuration on the associated device. See [“Detect and Resolve Configuration Conflicts” on page 694](#) for details.
 - **Customize port template CoS parameters**—Use this option to create a customized template by modifying CoS parameters. This option is activated only if the selected port template is Desktop plus Phone port, Layer 2 Uplink Port, or Routed Uplink Port. See [“Clone a Port Template” on page 694](#).
 - **Edit customized port template CoS parameters**—Use this option to edit the customized port template settings. See [“Edit a Port Template” on page 695](#).
 - **Delete customized port template**—Use this option to delete a customized port template.
 - **View port template configuration**—Use this option to view configuration details for a selected port template.

Apply or Edit a Port Template

The Manage Template Association screen displays the list of EX Series switches and their interfaces on which the selected port template is currently applied.

To apply a port template:

1. Click **+**. The **Add Ports** screen is displayed.
2. Select the device from the list. A list of available interfaces is displayed.
3. Select the interface or interfaces to which you want to apply the port template and click **>>**. The selected interfaces are displayed in the **Selected Ports** section.
4. If the selected template is **Desktop Port**, **Desktop and Phone Port**, or **Wireless Access Phone Port**, specify whether you want the port template to be applied to a VLAN. Select the VLAN from the list.

If the selected template is a **Layer 2 Uplink Port**, specify the list of VLANs and the native VLAN.

If the selected template is a **Routed Uplink Port**, specify the IP address.
5. Click **View Configuration** to view a summary of the resulting configuration.
6. Click **OK**.
7. The **Launch Update Device** screen is displayed with a list of modified devices. To update the device configuration immediately, select **Update Device now** and click **OK**.

Click **Finish** to save the changes and close the **Manage Template Port Association** screen.

To edit port template parameters:

1. Select the port template from the list in the **Manage Template Port Association** screen.
2. Click **Edit**. The Edit Port Template Association screen is displayed.
3. Modify the port template parameters. For example, if you select a Desktop Port template, modify the VLAN associated with the template.
4. Click **OK**.
5. In the Manage Template Association screen, click **Finish**.

To delete a port template from the selected ports:

1. Select the port template and click the **Apply/remove selected port template** button.
2. In the Manage Port Template Association screen, select the switch from the list.
3. Click **—** (the delete button).

In the Manage Template Association screen, you have the following options:

- **Save as Text**—Saves the details of port templates to port associations in a text file.
- **Save as HTML**—Saves the details of port templates to port associations in an HTML file.
- **Cancel**—Cancels all modifications and closes the **Manage Template Port Association** screen.

Detect and Resolve Configuration Conflicts

The Detect/Resolve Configuration Conflicts screen displays all the devices that are associated with the selected port template. To detect and resolve conflicts:

1. Select the devices from the list.
2. To resolve conflicts, select **Remove Conflicts in Configuration**.
3. Click **OK**. The port configuration conflicts are detected or resolved and the results are displayed.
4. Click **OK** to confirm.
5. The **Launch Update Device** screen is displayed with a list of modified devices. To update the device configuration immediately, select **Update Device now** and click **OK**.

Clone a Port Template

To customize a port template's CoS parameters:

1. Select the predefined port template from the list and click the **Customize Port Template** button and choose **Clone to new Port Template**. The Customize Port Template screen is displayed.



NOTE: As an administrator, you can create port templates using the Customize Port Template feature.

2. To modify the default template name, type a name in the **Template Name** field.
3. To modify the default description, type a description in the **Description** field.
4. To modify the default scheduler map name, type a name in the **Scheduler Map Name** field.

5. To edit scheduler settings, click **Edit Scheduler**. The Edit Scheduler screen is displayed. Specify the following:

- Scheduler name
- Transmit Rate—Select one: **Unconfigured** if you do not want to configure the parameter, **Percent** and enter a value, or **Remainder** to assign the remaining bandwidth available.
- Buffer Size—Transmit Rate—Select one: **Unconfigured** if you do not want to configure the parameter, **Percent** and enter a value, or **Remainder** to assign the remaining buffer available.
- Priority—Select a value from the list.

Click **OK** to save the settings or **Cancel** to cancel all modifications.

6. Click **Save** to create the customized port template.

See [“Detect and Resolve Configuration Conflicts” on page 694](#).

Edit a Port Template

To edit a customized port template's CoS parameters:

1. Select a customized port template from the list, click the Customize Port Template button and choose **Edit Port Template**. The Confirm Edit Port Template screen is displayed with a list of devices on which the template has been applied. Click **OK** if you want to proceed with the modification, else click **Cancel**.
2. To modify the default description, type a description in the **Description** field.
3. To modify the default scheduler map name, type a name in the **Scheduler Map** Name field.
4. To edit scheduler settings, click **Edit Scheduler**. The Edit Scheduler screen is displayed. Specify the following:
 - Scheduler name
 - Transmit Rate—Select one: **Unconfigured** if you do not want to configure the parameter, **Percent** and enter a value, or **Remainder** to assign the remaining bandwidth available.
 - Buffer Size—Transmit Rate—Select one: **Unconfigured** if you do not want to configure the parameter, **Percent** and enter a value, or **Remainder** to assign the remaining buffer available.
 - Priority—Select a value from the list.

Click **OK** to save the settings or **Cancel** to cancel all modifications.

5. Click **Save** to create the customized port template.

See [“Detect and Resolve Configuration Conflicts” on page 694](#).

CHAPTER 16

Unified Access Control Manager

- [Overview of the Unified Access Control \(UAC\) Manager Views on page 697](#)
- [Associating Enforcement Points with an Infranet Controller in the UAC Manager on page 698](#)
- [Disassociating Enforcement Points from an Infranet Controller in the UAC Manager on page 699](#)
- [Resolving Configuration Conflicts with the Infranet Controller in the UAC Manager on page 699](#)
- [Enabling 802.1X on Enforcement Point Ports in the UAC Manager on page 700](#)
- [Disabling 802.1X on Enforcement Point Ports in the UAC Manager on page 701](#)
- [Resolving Configuration Conflicts Between Devices and 802.1X Ports in the UAC Manager on page 701](#)

Overview of the Unified Access Control (UAC) Manager Views

Opening the Unified Access Control (UAC) Manager in the Configure module of the NSM UI allows you to view UAC policy attributes from the perspective of Infranet Controllers (IC) and Enforcement Points (EP).

- [The Infranet Controller View on page 697](#)
- [The Enforcement Point View on page 698](#)

The Infranet Controller View

The NSM main display area is horizontally divided into two tables. When you select the IC view, the upper table lists all the ICs managed by NSM's current domain. If the ICs are in cluster mode, the table also displays whether they are active/active or active/passive clusters. Selecting an IC causes NSM to list all the EPs and their location groups that are associated with the selected IC in the lower table. If the association is created with a load balancer option, then the load balancer is also displayed. In an active/active cluster mode, the IC cluster member name is displayed but not in the case of Standalone and active/passive cluster modes.

From the IC table, you can edit the configuration of a selected IC using the edit button provided above the IC table. The edit dialog is similar to the edit device action in the Device Manager.

The Enforcement Point View

When you select the Enforcement Points (EP) view, the NSM main display area is horizontally divided into the Enforcement Points table at the top and tab views of associated Infranet Controllers (ICs) and Port Details at the bottom. NSM displays only EX Series switches managed by a current domain in the EP table. Selecting an EP causes NSM to populate relevant information in the tab views. From the IC tab view, you can view the associated IC and its location group information. From the Port details tab, you can see the 802.1X enabled port names and their details.

Associating Enforcement Points with an Infranet Controller in the UAC Manager

To associate Enforcement Points (EP) with a selected Infranet Controller (IC):

1. Select the Add button (+) above the Enforcement Points table. NSM displays a list of EPs not managed by the selected IC. If the selected IC, is an IC cluster in Active-Active mode, then you must select the IC cluster member with which the EP association is to be created.
2. Select EPs to associate with the selected IC from the list. You can also search by strings for a particular EP.
3. Enter the shared secret between the IC and the EPs.
4. Select the Location Group the EPs must belong to in the selected IC. Each EP can be associated with only one Location Group available in the IC.
5. Enter the Infranet Controller port to which the EP should communicate. The default port is 1812.
6. Enter the IP address that should be used for RADIUS communication. If you do not specify an address, the EP's management IP address is used by default. You have the option to select the IP address of the RADIUS communication server only if you select a single EP because the IP address to communicate with an IC is unique.
7. Select **Use Load Balancer with IP Address** if the IC is load balancer administered. The IP address of the Load Balancer is then used as the RADIUS server in the EX Series switch configuration.

In an Active-Active cluster, with the load balancer selected, you can select one cluster member and perform an IC-EP association. You do not need to repeat the association for every cluster member.
8. Select the check box to run an Update Device task, which pushes configuration changes on both the IC and EPs. The configuration status of the EPs changes to Managed, InSync.

9. Select the check box to run a Summarize task that ensures the association between the IC and EP in the application database. The configuration status of these devices becomes Managed, NSM Changed.
10. Select OK. The selected EPs are listed under the associated IC.



NOTE: If you delete an IC or EP device, all related IC-EP associations are removed from the UAC Manager, and the configuration in the device is not modified. Therefore, first remove IC-EP and 802.1X port associations from the UAC Manager before you delete an IC or EP. This ensures that configuration changes take effect in the device.

Disassociating Enforcement Points from an Infranet Controller in the UAC Manager

To disassociate Enforcement Points (EP) from a selected Infranet Controller (IC):

1. Select from the EP table the EPs to disassociate from the IC.
2. Select the Delete button (-) above the Enforcement Points table.
3. Select the check box to run an Update Device task, which pushes configuration changes on both the IC and EPs. The configuration status of the EPs changes to Managed, InSync.
4. Select the check box to run a Summarize Delta Config task that ensures the association between the IC and EP in the application database. The configuration status of these devices becomes Managed, NSM Changed.
5. Select OK. The selected EPs are removed from the IC.

Resolving Configuration Conflicts with the Infranet Controller in the UAC Manager

Before you resolve configuration conflicts, perform an "Import Device" to identify the actual conflicts in the configuration. To ensure that configurations made in the Infranet Controller (IC) device are reflected in the UAC Manager:

1. Select an IC and right-click on it
2. Select **Resolve Config Mismatch** from the dropdown menu. The Resolve Config Mismatch window opens displaying differences between the RADIUS client configuration in the IC and the administered data in the UAC Manager.
3. Verify whether data currently administered in NSM by the UAC Manager is consistent with data in the RADIUS client device configuration.

4. Provide shared secret and IC port details. You cannot update information without this information.
5. Select the action to be performed:
 - Select **Update UAC Manager** to update the UAC Manager data with the data from the IC.
 - Select **Update the Enforcement Point Association in the Infranet Controller** to update the data in the IC with data from the UAC Manager.
6. Select the check box to overwrite the shared secret in the device.
7. Select the check box to run an Update Device task after you make an update.
8. Select OK.

When working in the Infranet Controller cluster mode, you face the following restrictions in conflict resolution:

- If the Infranet Controller in cluster mode is modified from active/active to active/passive or vice versa, you must delete the Infranet Controller from NSM and re-create it so that the UAC Associations are created in the context of the new mode. This action is required because in the active/passive mode, the Infranet Controller has a virtual IP address that the EX Series switches use to configure the RADIUS server tag. But in case of active/active, the RADIUS server tag contains the IP address of the individual cluster member.
- The "Resolve Configuration Conflicts" task is not supported if you modify the cluster mode.
- When an Infranet Controller is in cluster mode active/active, you can only use the "Resolve Configuration Conflicts" task to view conflicts but not resolve them.
- You can only view and resolve conflicts for new or modified entries. Entries requiring deletion do not appear because the Resolve Configuration Conflicts operation cannot identify these entries from the RADIUS client of the IC.

Enabling 802.1X on Enforcement Point Ports in the UAC Manager

To enable 802.1X on ports on Enforcement Points (EP):

1. Select an EP on whose ports you wish to enable 802.1X.
2. Select the Add button (+) below the Port Details tab.
3. Select one or more ports from the list. You can also search for a port name.
4. Select the optional Supplicant Mode attributes:

- **Single Secure**—Only one host is authenticated by the port.
 - **Multiple**—Multiple hosts are individually authenticated.
 - **Single** — Multiple hosts are authenticated using the first host's authentication.
5. Select the optional Authentication attributes—Whether reauthentication is allowed and the action to be taken if authentication fails.
 6. Select the check box to run an Update Device task, which pushes configuration changes on both the IC and EPs. The configuration status of the devices changes to Managed, InSync. The 802.1X enabled ports appear under the EP.
 7. Select the check box to run a Summarize Delta Config task that ensures the association between the EP and the ports in the application database. The configuration status of these devices become Managed, NSM Changed.
 8. Select OK.

Disabling 802.1X on Enforcement Point Ports in the UAC Manager

To disable 802.1X on ports on Enforcement Points (EP):

1. Select the Delete button (-) below the Port Details tab.
2. Select one or more ports on the EP on which to disable 802.1X.
3. Select the check box to run an Update Device task, which pushes configuration changes on the EP. The configuration status of the EPs changes to Managed, InSync. Ports on which 802.1X is disabled are removed from the EP.
4. Select the check box to run a Summarize Delta Config task that ensures the association between the EP and the ports in the application database. The configuration status of these devices become Managed, NSM Changed.
5. Select OK.

Resolving Configuration Conflicts Between Devices and 802.1X Ports in the UAC Manager

The Resolve Configuration Conflict option allows you to detect any inconsistency between the device configuration and 802.1X interfaces data administered in NSM by the UAC Manager. To resolve a conflict:

1. Right-click on a selected Infranet Controller.
2. Select Resolve Configuration Conflict. A dialog box displays details of port associations.

3. Select an action for NSM to execute:
 - Update the 802.1X port configuration in the Enforcement Point
 - Update NSM
4. Select the check box to run an Update Device task, which pushes configuration changes on the EP.
5. Select the check box to run a Summarize Delta Config task that ensures the association between the EP and the ports in the application database.
6. Select OK.

PART 4

Monitoring

- [Realtime Monitoring on page 705](#)
- [Analyzing Your Network on page 753](#)
- [Logging on page 783](#)
- [Reporting on page 855](#)

Realtime Monitoring

The Realtime Monitor module includes four views that you can use to monitor the status and traffic statistics for all the managed Juniper Networks devices in your network in real time.

To access, monitor, and configure the NSM management system, you use the Server Manager module. The NSM management system consists of a GUI Server and a Device Server.

This chapter contains the following sections:

- [About the Realtime Monitor on page 705](#)
- [Monitoring Managed Devices on page 706](#)
- [Monitoring IDP Sensors on page 732](#)
- [Monitoring VPNs on page 736](#)
- [Monitoring NSRP Statistics on page 739](#)
- [Monitoring IDP Clusters on page 742](#)
- [Using the Realtime Monitor on page 744](#)
- [Monitoring the Management System on page 744](#)

About the Realtime Monitor

The Realtime Monitor module in NSM enables you to monitor real-time status and statistics about all the managed devices, VPN tunnels, NSRP clusters, IDP sensors and IDP clusters in your network at a glance. You can use the Realtime Monitor to identify problems, track security events, and discover trends across multiple geographic regions and functional areas from a central management location.

The Realtime Monitor can also help you quickly identify potential device, network, and system-level problems, such as:

- Configuration status—At the device level, you can monitor the changing status of one or more devices in real time.
- Connection status—At the network level, you can monitor problems that could lead to failed devices.

- **Performance**—At the system level, you can monitor the activity between VPN members or NSRP cluster.

The Realtime Monitor tracks the integrity of your security perimeter by continually monitoring your security devices for security events (failed security devices, abnormal utilization, general errors). The Realtime Monitor does the work of a management expert by first gathering information about specific processes and network activity, then color-coding each event to organize problems.

Realtime Monitor Views

The Realtime Monitor includes the following views:

- **Device Monitor**—Displays status information on the managed devices in your network. This includes the name and type of the device managed in NSM, connection status, and current configuration status. From the Device Monitor, you can also access more detailed information and statistics on each security device, including ScreenOS 5.0 and later, mode, CPU utilization, memory, sessions, and network traffic. You can also use the Device Monitor to view status information on IDP sensors managed in your network.
- **VPN Monitor**—Displays status information on all VPN tunnel sessions that have been implemented within the domain you are working in. From the VPN Monitor, you can determine if a VPN tunnel is up, down, or not monitored.
- **NSRP Monitor**—Displays status information about NSRP (NetScreen Redundancy Protocol) clusters in your network. If you implement NSRP for the purpose of deploying clusters in your Juniper Networks security system, you can use the NSRP Monitor to view and troubleshoot the status of security devices in clusters within the domain you are working in.
- **IDP Cluster Monitor**—Displays status information about IDP clusters in your network. If you implement IDP clusters for the purpose of deploying clusters in your Juniper Networks security system, you can use the IDP Cluster Monitor to view and troubleshoot the status of IDP sensors in clusters within the domain you are working in.

Monitoring Managed Devices

Use the Device Monitor to get an at-a-glance view of the current status of all the managed devices and IDP sensors in your network.

If you have configured multiple subdomains, you can view all your managed devices from the global domain.

Viewing Device Status

[Table 51 on page 707](#) lists and describes device information that you can view through the Device Monitor:

Table 51: Device Status Information

Column	Description
Name	Unique name assigned to the device in NSM.
Domain	Domain in NSM in which the device is managed.
Platform	Model number of the device.
OS Version	Operating system firmware version running on the device.
Config Status	<p>Current configuration status of the device in NSM:</p> <ul style="list-style-type: none"> • None. No state has been set (does not show in Device Monitor). • Modeled. The device exists in NSM, but a connection to the device has not yet been established. • RMA. Equivalent to bringing the device into the Modeled state. RMA results from an administrator selection in the User Interface when a device goes down. • Waiting for 1st connect. NSM is waiting for the device to connect. You must enter a command on the device to make it connect to NSM. • Import Needed. You must import the configuration of the device into NSM. When you add a device for the first time, verify that your status indicates "Import Needed" before you attempt to import the device. During migration, this state indicates that import of the security device configuration is still required. • OS Version Adjustment Needed. The firmware version detected running on the device is different than what was previously detected in NSM. This could happen in the event that the automatic adjustment option was cleared during a change device firmware directive or an Update Device directive was issued to an IDP device with a firmware version mismatch. • Platform Mismatch. The device platform selected when adding the DMI device in NSM does not match the device itself. A device in this state cannot connect to NSM. • Device Firmware Mismatch. The OS version selected when adding a DMI device does not match the OS version running on the device itself. • Device Type Mismatch. The type of device specified when adding the device in NSM does not match the device itself. The device type might indicate whether the device is part of a vsys device, part of a cluster, or part of a virtual chassis. A device in this state cannot connect to NSM. • Detected duplicate serial number. The device has the same sequence number as another managed device. A device in this state cannot connect to NSM. • Update Needed. An update to this device is required. • Managed. The device is currently being managed by NSM. For devices running ScreenOS 5.0 and later, the Device Monitor can display the following additional configuration states: • Managed, In Sync. The physical device configuration is synced with the modeled configuration in NSM.

Table 51: Device Status Information (continued)

Column	Description
Config Status (continued)	<ul style="list-style-type: none"> Managed, Device Changed. The physical device configuration is out-of-sync with the modeled configuration in NSM. Changes were made to the physical device configuration (the configuration on the physical device is newer than the modeled configuration). For Junos devices with redundant Routing Engines, this status can indicate that a routing engine switchover has occurred. Managed, NSM Changed. The modeled device configuration is out-of-sync with the physical device configuration. Changes were made to the modeled configuration (the configuration on the NSM is newer than the physical device configuration). Managed, NSM and Device Changed. Both device configurations (physical and modeled) are out-of-sync each other. Changes were made to the physical device configuration and to the modeled configuration. Managed, Sync Pending. Completion of the Update Device directive is suspended and waiting for the device to reconnect. This state occurs only for ScreenOS devices that have the Update When Device Connects option selected during the device update.
Connection Status	<p>Connection status of the device in NSM:</p> <ul style="list-style-type: none"> Up. Device is currently connected to NSM. Down. Device is not currently connected to NSM, but has connected in the past. Never Connected. Device has never connected to NSM. <p>The Device Server checks the connection status of each device every 120 seconds by default. You can change this behavior by editing the value for the <code>devDaemon.deviceHeartbeatTimeout</code> parameter in the Device Server configuration file. Refer to the <i>Network and Security Manager Installation Guide</i> for more information on editing configuration files.</p> <p>NOTE: If the network connection goes down for a period longer than six to eight minutes, the device connection will permanently time out. If this occurs, and the device goes down for any reason, the device still appears as Up in the Device Monitor.</p>
Alarm	<p>Displays the current alarm status for each device in NSM:</p> <ul style="list-style-type: none"> If device has any alarms, the most severe alarm severity is displayed (either Major or Minor). None—The device has no alarms. Unknown—The device status is unknown. For example, the device might not be connected. N/A—The device's alarm is not pollable or discoverable, for example, this column shows "N/A" for ScreenOS and IDP devices. Alarm is colored: <ul style="list-style-type: none"> Red for Major. Orange for Minor. Green for Ignore, None, Unknown, or N/A.

Table 51: Device Status Information (continued)

Column	Description
H/W Inventory Status	Displays the inventory status for hardware on the device: <ul style="list-style-type: none"> • In Sync: The inventory information in the NSM database is synchronized with the information on the device. • Out Of Sync: The inventory information in the NSM database is not synchronized with the information on the device. • N/A: The connected device is a ScreenOS or IDP device, or the device is not connected and imported.
S/W Inventory Status	Displays the inventory status for software on the device: <ul style="list-style-type: none"> • In Sync: The inventory information in the NSM database is synchronized with the software on the device. • Out Of Sync: The inventory information in the NSM database is not synchronized with the software on the device. • N/A: The connected device is a ScreenOS or IDP device, or the device is not connected and imported.
License Inventory Status	Displays the inventory status for software on the device: <ul style="list-style-type: none"> • In Sync: The inventory information in the NSM database is synchronized with the licenses on the device. • Out Of Sync: The inventory information in the NSM database is not synchronized with the licenses on the device. • N/A: The connected device is a ScreenOS or IDP device, or the device is not connected and imported.
First Connect	The first time the device connected to the NSM Device Server.
Latest Connect	The last time the device connected to the NSM Device Server.
Latest Disconnect	The last time the device disconnected from the NSM Device Server.

Device Polling Intervals

NSM retrieves device statistics from the physical device. The device polling interval determines the number of seconds the Device Server waits before polling for new statistics.

To configure or view the device polling intervals, double-click the **Server Manager > Servers** node, then select the Device Server and click the **Edit** icon. The **Device Server** dialog box is displayed. Use the **Device Polling** tab to edit the intervals to meet your monitoring requirements:

Table 52: Device Polling Intervals

Statistic	Description	Poll Interval (in seconds)	Save Interval (in seconds)
Device	Details traffic, interface, zone, and system-related statistics on a specific device. Information appears in the Device Monitor.	300	300
VPN	Details VPN tunnels between your managed devices, including VPN tunnel status (Up, Down, Not Monitored), VPN name, VPN Type, VPN source, VPN destination, security parameter index (SPI), IP address, and protocol. Information appears in the VPN Monitor.	300	300
NSRP	Details high availability events and statistics, including VSD group ID, number of units in the cluster, state change counter, init counter, number of Master devices, number of Backup devices, and heartbeat information. Information appears in the NSRP Monitor.	300	300
Interface	Details the interface number, IP address, and zone to which the interface is mapped. Information appears in the Device Monitor, in the Device Summary.	300	300

Viewing Device Monitor Alarm Status

Alarms refresh automatically through periodic polling. To view the Alarm status and time:

1. From **Device Monitor**, right-click on the device row entry and select the **View Alarm** option.

The device **Alarm Status** dialog box displays the alarm list and polling time for the device.

2. To retrieve the current alarm status in the device, click the **Refresh** button.

The poll time is derived from the device server time.

Setting the Polling Interval For Device Alarm Status

The default polling interval is 900 seconds (15 minutes). To configure polling intervals for Alarm Status:

1. From **Device Manager > Devices**, double-click on the device to open it.

The Info tab dialog box is displayed.

2. Select the **Device Admin** page to set the polling interval for the device.

The minimum polling interval is 60 seconds. The maximum interval is 2,147,483,647 seconds. You cannot disable polling.

Viewing Additional Device Detail and Statistics

If a device is running, you can view additional status, using Device Detail Status and view traffic-related statistics and other information, using View Statistics.



NOTE: If a security device has never connected, the Device Detail Status and Statistics views for the device are not available.

Viewing Device Details

Double-click on any device to view additional details on the device related to resource usage. You can also right-click on the device and select View Details.



NOTE: The information in the Details window appears slightly different for firewall/VPN devices and IDP sensors. Details for IDP sensors include an additional tab showing process status.

Table 53 on page 711 lists and describes the information that you can view for a security device through the Device Detail Status:

Table 53: Device Detail Status Items

Item	Description
OS Version	ScreenOS firmware version running on the device.
Mode	Current operation mode of the device. Network Address Translation (NAT), Transparent, or Route.
Latest Reboot	Most recent date and time that the security device was powered off and on. You can use this information to determine how long the security device was down.
CPU Utilization	The percent of the CPU being used at the moment of the status snapshot.
One Min. Load	The percent CPU utilization average on the security device for the last minute.
5 Min. Load	The percent CPU utilization average on the security device for the last 5 minutes.
15 Min. Load	The percent CPU utilization average on the security device for the last 15 minutes.
Mem Allocated	The original amount of memory allocated to the security device.

Table 53: Device Detail Status Items (continued)

Item	Description
Mem Left	The amount of allocated memory that remains unused by the security device.
Mem Fragmented	The amount of fragmented memory.
Active Sessions	The number of active sessions on the security device.
Allocated Sessions	The number of sessions originally allocated to the security device.
Max Sessions	The maximum number of sessions on the security device.
Failed Sessions	The number of sessions that have failed on the security device.

Viewing Device Statistics

If a security device is running, you can access the Statistics view to see traffic, interface, zone, and other system-related information on the device. To view statistics on a particular security device, right-click the security device in either the **Device Monitor** or the **Device Manager** and select **View Statistics**. The Device Statistics Summary appears in a new window.



NOTE: The information in the Device Statistics window appears slightly different for firewall/VPN devices and IDP sensors.

Device Statistics Summary

The Device Statistics Summary displays the following details:

- Details describing the security device or virtual system, such as serial number and IP address, type, and firmware version.
- Interface information
- Device status
- Time-related statistics such as last connect or reboot

[Table 54 on page 712](#) details all the information you can view from the Device Statistics Summary.

Table 54: Device Statistics Summary

Item	Displays
Device	Device: Displays the name, serial number, and IP address of the security device. Vsys: Displays the serial number of the security device.

Table 54: Device Statistics Summary (continued)

Item	Displays
Vsys	The name of the virtual system (if applicable)
Version	The security device's build, model, and operation mode (this is not displayed in the Vsys view).
DC IP	The IP Address of the Device Server (this is not displayed in the Vsys view).
Interface Information	The employed interfaces. For example: Trust, Untrust, and Self.
Vsys Information	The virtual systems associated with this security device (this is not displayed in the Vsys view).
Last Known Connect Time	The last time the security device connected to the Data Collector (this is not displayed in the Vsys view).
Device Status	Whether the security device is currently up or down (this is not displayed in the Vsys view)
Last Reboot Time	The last time the system was restarted (this is not displayed in the Vsys view).
Last Known Uptime	If the security device is down, the entry lists the last time it was up. Used to determine how long a security device was down (this is not displayed in the Vsys view).
GMT Time Offset (Hours)	The hour the security device is set from Greenwich Mean Time (this is not displayed in the Vsys view).
DayLight-Saving	Whether you have enabled the security device to adjust time for daylight savings.

Additional Device Specific Views

From the Device Statistics Summary, you can access additional information enabling you to view and monitor key traffic, interface, zone, and other system-related information on a specific security device.

Table 55 on page 714 describes each device-specific view.

Table 55: Device-Specific Views

View Type	View	Purpose
Traffic	Policy Distribution	View traffic on the security device distributed by policy. Enables you to view a chart of the traffic distribution by policy.
	Protocol Distribution	View traffic on the security device distributed by protocol. Enables you to view a chart of the traffic distribution by protocol.
	VPN Distribution	View the up/down status and active statistics of VPNs on the security device (if applicable). Also enables you to view a chart of the VPN distribution by VPN tunnel.
Interface	Ethernet Statistics	View security device traffic over specific interfaces. Enables you to view a chart of the utilization distributed by interface.
	Flow Statistics	View security device traffic on flow counters over specific interfaces. Enables you to view a chart of flow statistics distributed by interface.
	Attack Statistics	View all of the attacks that have occurred on a security device over specific interfaces. Enables you to view a chart of attacks distributed by interface.
Zone	Ethernet Statistics	View security device traffic from specific zones. Enables you to view a chart of the traffic distributed by zone.
	Flow Statistics	View security device traffic on flow related statistics for specific interfaces. Enables you to view a chart of flow statistics distributed by zone.
	Attack Statistics	View all counters related to attacks that have occurred on a security device from specific zones. Enables you to view a chart of the attacks distributed by zone.
System	Resource Statistics	View CPU utilization and memory allocation statistics on the security device. Enables you to view CPU, Memory, and Session Utilization trends.
	Active Statistics	View administrator and user activities; active VPNs; and authenticated users on a security device. Also enables you to view a snapshot of the ongoing active sessions on the security device.
HA	NSRP Statistics	View NSRP statistics related to clusters created on the security device (if applicable).

Viewing Device Traffic Distribution

You can view statistics describing how the traffic on a specific security device is distributed, by policy, protocol, or VPNs (if applicable). You can use this information to help you

identify those policies, protocols, and VPN tunnels that are most and least frequently used on a security device.

Viewing Traffic Distribution by Security Policy

Click **Policy Distribution** to view security device traffic that matches the access policies configured for a security device. A bar graph appears (under the Chart tab) depicting the distribution of data by policy. The graph displays a percentage of the absolute number of bytes for the top 10 policies by default.

[Table 56 on page 715](#) describes all of the information that is available from the Policy Distribution view.

Table 56: Policy Distribution Items

Item	Description
Policy ID	The unique identifier of the policy.
Source IP	The IP address of the host generating the session.
Source IP Mask	The IP address mask for the host or network generating the session.
Destination IP	The IP address of the host receiving the session.
Destination IP Mask	The IP address mask for the host or network receiving the session.
Source Zone	Zone of the host generating the session.
Destination Zone	Zone of the host receiving the session.
VPN Name	Name of the virtual private network.
Service	The application or service associated with the policy. Examples include Mail, FTP, SNMP, AOL, Telnet, and LDAP.
Action	The activity to be performed, such as Permit, Deny, or Tunnel.
Total Connections	The total number of data connections.
Connection Rel%	The relative percentage of connections.
Delta Connection	The total numerical difference between the current connection value and the previous connection value.
Total Bytes	The total number of data bytes.
Bytes Rel%	The relative percentage of bytes.
Delta Bytes	The total numerical difference between the current bytes value and the previous bytes value.
Total Packets	The total number of data packets.

Table 56: Policy Distribution Items (continued)

Item	Description
Packets Rel%	The relative percentage of packets.
Delta Packets	The total numerical difference between the current bytes value and the previous packets value.
Policy Name	The name of the policy.

Adjusting Data Depicted Graphically

You can adjust all elements depicted in the graph, including the policies, data values (such as absolute or delta), and type of data (bytes in or out, packets in or out, utilization).

To adjust policies depicted graphically:

1. Right-click within the chart and select **Configure Policies**. A dialog box appears.
2. Clear the Default check box.
3. Click to select the policies that you wish to view on the graph from the list of Available Policies. Click **Add** to add the policies that you want to the list of Selected Policies. To remove a policy from the list of Selected Policies, click to select the policy and click **Remove**.
4. Click **OK**.

To adjust data and data types depicted graphically:

1. Right-click the Chart view.
2. From the **Data** option, select either **Delta** or **Absolute**.
3. From the **Data Type** option, select either **Connections**, **Bytes**, or **Packets**.
4. Click **OK**.

Viewing Traffic Distribution by Protocol

Click the **Protocol Distribution** node to view the distribution of traffic according to the protocols flowing through the device. Protocols are predefined services (such as HTTP, SNMP, or Telnet) that are enabled for each security device. You can view up to ten protocols. A bar graph displays a percentage of the absolute number of bytes for the top 10 protocols by default.

[Table 57 on page 716](#) describes all the information that is available from the Protocol Distribution view:

Table 57: Protocol Distribution Items

Item	Description
Protocol	The name of the predefined service (HTTP, SNMP, or Telnet) operating on the selected interface.

Table 57: Protocol Distribution Items (continued)

Item	Description
Interface	The type of interface through which the protocol is flowing.
Bytes In	The number of incoming bytes for the protocol through the security device.
Bytes In Rel%	Relative percentage of all incoming bytes.
Delta Bytes In	The total numerical difference between the current bytes in value and the previous bytes in value.
Bytes Out	The number of outgoing bytes for the protocol through the security device.
Bytes Out Rel%	Relative percentage of all outgoing bytes.
Delta Bytes Out	The total numerical difference between the current bytes out value and the previous bytes out value.
Packets In	The number of incoming packets handled by the protocol through the security device.
Packets In Rel%	Relative percentage of all incoming packets.
Delta Packets In	The total numerical difference between the current packets in value and the previous packets in value.
Packets Out	The number of outgoing packets handled by the protocol through the security device.
Packets Out Rel%	Relative percentage of all outgoing packets.
Delta Packets Out	The total numerical difference between the current packets out value and the previous packets out value.
Util. (Absolute)	The total number of the utilization of the current security device.
Util. (Delta)	The total numerical difference between the current utilization value and the previous utilization value.
Zone	The name of the zone associated with the protocol.

Adjusting Data Depicted Graphically

You can adjust the interfaces (such as Trust, Untrust, Management, NSRP, and Self) and data depicted graphically in the same way that you adjust the Policy Distribution graphs.

You can also adjust the data types in the Protocol Distribution graph by Bytes In, Bytes Out, Packets In, Packets Out, or Utilization, and by Interface.

Viewing Traffic Distribution by VPN (if applicable)

If you use your security devices to implement VPNs, you can view how traffic is distributed across each VPN tunnel on the security device. A bar graph (under the Chart tab) depicts the distribution of data traveling to and from each VPN tunnel. The graph uses a percentage of the absolute number of bytes traveling in to the top 10 VPN tunnels by default.

You can adjust all elements depicted in the graph, including the VPN tunnels, data values (absolute or delta), and type of data (bytes in or out, packets in or out, utilization).

Adjusting VPN Tunnels Depicted Graphically

1. Right-click the Chart view and select **Configure VPNs**.
2. Clear the Default check box.
3. Click to select the VPN tunnel that you wish to view on the graph from the list of Available VPN tunnels. Click **Add** to add the VPN tunnel to the list of Selected VPN tunnels. To remove a VPN tunnel from the list of Selected VPN tunnels, click to select the VPN tunnel and click **Remove**.
4. Click **OK**.

Adjusting Data Depicted Graphically

1. Right-click the Chart view and select **Data**, and either Delta or Absolute.
2. Right-click the Chart view and select **Data Type**, and either **Bytes In**, **Bytes Out**, **Packets In**, **Packets Out**, **Utilization**, **Last Session Duration**, **Avg Latency**, **Availability**.
3. Click **OK**.

Viewing VPN-Specific Information

Click the **VPN Monitor Table** tab to view specific information about your VPN. From the VPN Monitor Table, you can view the following details about a specific VPN:

- Key details describing the VPN (such as name, Policy ID, group and user associations, VPN type).
- Security Association (SA) information.
- Traffic over the tunnel (such as bytes in/out, packets in/out, utilization).

[Table 58 on page 718](#) describes all the information that is available from the VPN Monitor:

Table 58: VPN Monitor Table

Item	Description
Name	The name of the VPN.
VPN Type	Type of tunnel: Site-to-site or dial-up.
SA Id	The Security Association (SA) identification for the VPN at both ends of the tunnel.

Table 58: VPN Monitor Table (continued)

Item	Description
Policy Id--In/Out	A unique identifier specified when the policy was configured.
Status	The status of the VPN tunnel (up or down).
SA Status	Whether or not the current SA has been established.
Time-SA Status Change	Time that the SA status last changed.
Last SA Session Duration	Duration of last SA session.
Group	Group associated with the VPN.
User	User associated with the VPN.
DN Name	Distinguished Name (DN) of the VPN.
Avg. Latency	A rolling average of latency, presented in milliseconds.
Availability	Percentage of the time a tunnel is up over the last thirty samples.
Bytes In	The number of incoming bytes handled by the protocol through the security device.
Delta Bytes In	Total numerical difference between the current bytes in value and the previous bytes in value.
Bytes Out	The number of outgoing bytes handled by the protocol through the security device.
Delta Bytes Out	Total numerical difference between the current bytes out value and the previous bytes out value.
Packets In	The number of incoming packets handled by the protocol through the security device.
Delta Packets In	Total numerical difference between the current packets in value and the previous packets in value.
Packets Out	The number of outgoing packets handled by the protocol through the security device.
Delta Packets Out	Total numerical difference between the current packets out value and the previous packets out value.
Util. (Absolute)	Total number of the utilization of the current security device.
Util. (Delta)	Total numerical difference between the current utilization value and the previous utilization value.

Viewing Active VPN Information

Click the **Active VPN** tab to view specific information about active VPNs. From the **Active VPN** tab, you can view the following details:

- Key details describing the VPN (such as name, Policy IP, local and peer gateway IDs and IP addresses).
- Security established on the active VPN.
- Time-related statistics (such as lifetime, latency).

Table 59 on page 720 lists the information that is available from the active VPN.

Table 59: Active VPN Table

Item	Description
Name	Name of the active VPN.
VPN Type	Type of tunnel: Site-to-site or dial-up.
Policy Id--In/Out	A unique identifier specified when the policy was configured.
Status	Tunnel status for the active VPN is UP or Down.
Ave Latency	Rolling average latency (in milliseconds).
Last Latency	Latency for the last ping response.
Availability	Percentage of time a tunnel is available over the last 30 samples.
Local GW Id	Local gateway ID for the active VPN.
Peer GW Id	Peer gateway ID for the active VPN.
Local GW IP	Local gateway IP address for the active VPN.
Peer GW IP	Peer gateway IP address for the active VPN.
Local Address	Local IP address for the security device associated to the active VPN.
Peer Address	Peer IP address for the security device connected to the active VPN.
Monitor	Monitoring capability status for the VPN: ON or OFF.
IPSec	IPSec (IP security) protocol for the active VPN; AH (Authentication Header) or ESP (Encapsulating Security Payload).
SPI In	SPI (Security Parameter Index) key into the active VPN. A value that identifies a security association (SA).

Table 59: Active VPN Table (continued)

Item	Description
SPI Out	SPI (Security Parameter Index) key out of the active VPN. A value that identifies an SA.
Encryption	Algorithm used when a user encrypts communication between the security device and the server. Listed as either SDI or DES.
Authentication	Second algorithm used for user encrypted communication between the security device and server.
Key	Type of key associated with the VPN: Auto IKE (Internet Key Exchange) or manual key.
Lifetime P1	Time listed in seconds before re-keying.
Lifetime P2	Time reported in remaining bytes before re-keying. Independent from Lifetime P1.
Life Size	Predefined duration of the tunnel (in bytes).
P1 Status	P1 (phase 1) status for tunnel negotiation: enabled or disabled.
P2 Status	P2 (phase 2) status for tunnel negotiation: enabled or disabled.
P1 Auth	Associated with Auto IKE. This column displays the P1 (phase 1) authentication for the active VPN.

Viewing Interface Statistics

You can view traffic information as it is processed by a device on a specific interface:

- [Viewing Ethernet Statistics on page 721](#)
- [Viewing Flow Statistics on page 723](#)
- [Viewing Attack Statistics on page 724](#)

Viewing Ethernet Statistics

Click the **Ethernet Statistics** node to view traffic information as it is processed by a specific physical interface on a security device. Depending upon the specific security device, the following interfaces apply:

- Trust and Untrust interfaces available on all security devices.
- DMZ interface available on NetScreen-25, NetScreen-50 and NetScreen-500 devices; the NetScreen-5XP device has no DMZ interface.
- HA interface and management interface available on NetScreen-500 devices.

Ethernet Statistics apply only to security devices, and not to virtual systems.

A graph displays security device percentage utilization traffic on the interface. Right-click within the chart to select a desired interface (such as Ethernet or HA). The active interface is listed below the graph. The graph also provides the total errors. You can view up to 12 samples in the chart. [Table 60 on page 722](#) describes the information available from the Ethernet Statistics view:

Table 60: Ethernet Statistics View Data

Item	Description
Interface	The data for each interface.
Bytes In	The number of bytes of incoming traffic processed through the security device over the selected interface.
Delta Bytes In	The total numerical difference between the current bytes in value and the previous bytes in value.
Bytes Out	The number of outgoing bytes handled by the interface through the security device.
Delta Bytes Out	The total numerical difference between the current bytes out value and the previous bytes out value.
Packets In	The number of incoming packets handled by the interface through the security device.
Delta Packets In	The total numerical difference between the current packets in value and the previous packets in value.
Packets Out	The number of outgoing packets handled by the interface through the security device.
Delta Packets Out	The total numerical difference between the current packets out value and the previous packets out value.
Broadcast	The number of broadcast-type packets processed through the security device over the selected interface.
CRC Errors	The number of packets generating a cyclic redundancy code error processed through the security device over the selected interface.
Alignment Errors	The number of Frame Checksum (FCS) errors.
ShortFrame	The number of frames that are not of the correct length.
RXCollision	The number of times that two packets collide, resulting in damage to both. This indicates that the network is overloaded.
Speed (Mbps)	This is useful in calculating the speed of the interface.
Status	Whether the security device is currently Up or Down.

Table 60: Ethernet Statistics View Data (continued)

Item	Description
Direction	Whether the security device is in half or full duplex mode.
Zone	The name of the zone associated with the interface.

Viewing Flow Statistics

Click the **Flow Statistics** node to view data for various flow counters on a specific security device or virtual interface. For each security device, the data and statistics are separated by all available interfaces.

To change the interface setting, right-click in the chart and select the interface that you want. [Table 61 on page 723](#) describes all the information that is available from the Flow Statistics view:

Table 61: Flow Statistics View Data

Item	Description
Interface	The name of the virtual interface.
Bytes In	The number of bytes of incoming traffic processed through the security device over the selected interface.
Bytes Out	The number of bytes of outgoing traffic processed through the security device over the selected interface.
Packets In	The number of incoming packets processed through the security device over the selected interface.
Packets Out	The number of outgoing packets processed through the security device over the selected interface.
VLAN In	The number of VLAN packets received through the security device; applies to virtual systems.
VLAN Out	The number of VLAN packets sent through the security device; applies to virtual systems.
Connections	The number of connections that occurred for a given interface.
Packets Dropped	The number of incoming packets dropped by a given interface.
Packets Denied	The number of incoming packets denied on the virtual interface by the policy.
Authentication Failed	The number of packets dropped because of an authentication failure.

Table 61: Flow Statistics View Data (continued)

Item	Description
URL Blocking Dropped	The number of packets dropped because of URL blocking.
IPSec Dropped	The number of IPSec packets dropped.
Zone	The name of the zone associated with the interface.

Viewing Attack Statistics

Click the **Attack Statistics** node to view distribution of the attacks that have occurred on a specific security device. The report separates the data and statistics for all available interfaces. [Table 62 on page 724](#) describes each of the attack counters available from the Attack Statistics view:

Table 62: Attack Counters

Item	Description
Interface	Name of the interface.
SYN Attack	SYN packets overwhelm a network by initiating so many connection attempts or information requests that the network can no longer process legitimate connection requests, resulting in a Denial of Service.
Tear Drop	When the first and second parts of a fragmented packet overlap, the server attempting to reassemble the packet can crash. If the security device sees this discrepancy in a fragmented packet, it drops the packet.
Source Route	This option applies in an IP header and allows an attacker to enter a network with a false IP address and have data sent back to the attacker's real address.
Ping of Death	Intentionally oversized or irregular ICMP packets can trigger a Denial of Service condition, freezing, or other adverse system reactions. You can configure a security device to detect and reject oversized or irregular packet sizes.
Address Spoofing	You can enable a security device to guard against spoofing attacks by checking its own route table. If the IP address is not in the route table, traffic through the security device is not allowed.
Land Attack	Combining a SYN attack with IP spoofing, a Land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address. This creates an empty connection. Flooding a system with such empty connections can overwhelm the system, causing a Denial of Service. Security devices automatically block any attempt of this nature and records such attempts as a Land attack.
ICMP Flood	ICMP pings can overload a system with so many echo requests that the system expends all its resources responding until it can no longer process valid network traffic. If you set a threshold to invoke ICMP flood attack protection when exceeded, ICMP flood attacks are recorded as statistics.

Table 62: Attack Counters (continued)

Item	Description
UDP Flood	Similar to the ICMP flood, UDP flooding occurs when UDP packets are sent with the purpose of slowing down the system to the point that it can no longer handle valid connections. After enabling the UDP flood protection feature, you can set a threshold that once exceeded invokes the UDP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, the security device ignores further UDP packets for the remainder of that second.
WinNuke	WinNuke can cause any computer on the Internet running Windows to crash. WinNuke introduces a NetBIOS anomaly that forces Windows to restart. Security devices can scan any incoming Microsoft NetBIOS Session Service packets, modify them, and record the event as a WinNuke attack.
Port Scan	Port scan attacks occur when packets are sent with different port numbers with the purpose of scanning the available services in hopes that one port will respond. The security device internally logs the number of different ports scanned from one remote source. If a remote host scans 10 ports in 0.3 seconds, the device flags this as a port scan attack, and rejects further packets from the remote source.
IP Sweep	Also called an address sweep attack, an IP sweep is similar to a port scan attack. It occurs when an attacker sends ICMP echo requests (or pings) to different destination addresses hoping that one will reply, thus uncovering an address to a target. If a remote host pings 10 addresses in 0.3 seconds, the security device flags this as an address sweep attack and drops the connection.
Block Java/ActX	Malicious Java or ActiveX components can be hidden in Web pages. When downloaded, these applets install a Trojan horse on your computer. Similarly, Trojan horses can be hidden in compressed files such as .zip, .gzip, .tar, and executable (.exe) files.
SYN Frag	A SYN fragment attack floods the target host with SYN packet fragments. The host catches the fragments, waiting for the remaining packets to arrive so it can reassemble them. When a server or host is flooded with connections that cannot be completed, the host's memory buffer eventually fills. No further connections are possible, and damage to the host's operating system can occur. The security device drops ICMP packets when the protocol field indicates ICMP packets, and the fragment flag is set to 1 or an offset is indicated.
TCP no Flag	TCP packet that does not have any bits set in the flags.
Unknown Prot	The security device drops packets where the protocol field is set to 101 or greater. These protocol types are reserved and undefined at this time.
Bad IP Opt	Triggered when the list of IP options in the IP datagram header is incomplete or malformed.
IP Rec Route	The security device blocks packets where the IP option is 7 (Record Route). This option is used to record the route of a packet. A recorded route is composed of a series of Internet addresses, which an outsider can analyze to learn details about your network's addressing scheme and topology.

Table 62: Attack Counters (continued)

Item	Description
IP Timestamp	The security device blocks packets where the IP option list includes option 4 (Internet Timestamp).
IP Security	This option provides a way for hosts to send security, compartmentation, TCC (closed user group) parameters, and Handling Restriction Codes compatible with DOD requirements.
IP Loose Src	The security device blocks packets where the IP option is 3 (Loose Source Routing). This option provides a means for the source of a packet to supply routing information to be used by the gateways in forwarding the packet to the destination. This option is a loose source route because the gateway or host IP is allowed to use any route of any number of other intermediate gateways to reach the next address in the route.
IP Strict Src	The security device blocks packets where the IP option is 9 (Strict Source Routing). This option provides a means for the source of a packet to supply routing information to be used by the gateways in forwarding the packet to the destination. This option is a strict source route because the gateway or host IP must send the datagram directly to the next address in the source route, and only through the directly connected network indicated in the next address to reach the next gateway or host specified in the route.
IP Stream	The security device blocks packets where the IP option is 8 (Stream ID). This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support the stream concept.
ICMP Frag	When the protocol field indicates ICMP packets, and the fragment flag is set to 1 or an offset is indicated.
Large ICMP	An ICMP packet with a length greater than 1024.
SYN n FIN	Both the SYN and FIN flags are not normally set in the same packet. However, an attacker can send a packet with both flags set to see what kind of system reply is returned and thereby determine what kind of system is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks. Enable this option to have the security device drop packets that have both the SYN and FIN bits set in the flags field.
FIN no ACK	TCP packet with a FIN set but no ACK set in the flags field.
Mal URL	When you enable Malicious URL Detection, the security device monitors each HTTP packet and detects any URL that matches any of several user-defined patterns. The security device automatically drops any such packet.
Limit Session	Security devices can limit the number of sessions that can be established by a single IP address. For example, session resources on a Web server can be exhausted if there are many requests from the same client. This option defines the maximum number of sessions the security device can establish per second for a single IP address. (The default threshold is 128 sessions per second per IP address.)

Table 62: Attack Counters (continued)

Item	Description
Block Frag	As packets traverse different networks, it is sometimes necessary to break a packet into smaller pieces (fragments) based upon the network's maximum transmission unit (MTU). IP fragments may carry an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the target system receives these packets, the results range from not processing the packets correctly to crashing the entire system. When you enable the security device to deny IP fragments on a security zone, the security device blocks all IP packet fragments that it receives at interfaces bound to that zone.
Zone	The name of the zone associated with the attack.

Viewing Zone Statistics

You can view traffic information as it is processed by a security device over specific zones. You can view Ethernet statistics, flow statistics, and attack statistics.

Viewing System Statistics

You can also view system-related information for a security device.

Viewing Resource Statistics

Click the **Resource Statistics** node to view the resources for a security device.

[Table 63 on page 727](#) describes all the information that is available from the Resource Statistics view:

Table 63: Resource Statistics Items

Item	Description
Avg. CPU Utilization	The average CPU usage of the security device.
Memory Allocated	The current memory allocation to the security device.
Memory Left	The remaining usable memory.
No. of Fragment Blocks	The percentage of blocks that are fragmented.
Active Sessions	The number of currently active sessions.
Allocated Sessions	The number of allocated sessions.
Max. Sessions Allowed	The maximum sessions allowed.
Failed Sessions	The number of sessions that failed to allocate (after maximum reached).

Viewing Active Statistics

Click the **Active Statistics** node to view administrator and user activities for a security device. The **Administrators** tab displays information about the administrators, including when, where, and how they logged in to the system. [Table 64 on page 728](#) describes all the information that is available from the Administrators view:

Table 64: Administrators View

Item	Description
Administrator ID	The administrator's login ID.
IP Address	The administrator's IP address.
Service Used	The type of service, for example, Console, Web, or Telnet.
Time	The time that the administrator logged in.

[Table 65 on page 728](#) describes all information that is available from the Authenticated Users view:

Table 65: Authenticated Users View

Item	Description
User ID	User login ID.
Source IP Address	Source IP address.
Time	Time that the user logged in.

You can also access VPN information from the Active VPN view, and Active Session information from the **Active Sessions** view.

Viewing Active Sessions

You can view a snapshot of ongoing active sessions on the security device. You can view active sessions from the **Active Statistics** view.

When you click the **Active Sessions** tab, a short view of the active sessions displays basic information (such as source IP, destination IP, translated IP (if applicable), source port, destination port, translated port (if applicable), policy ID, time the session starts, and protocol type) about the active sessions on the security device by default. You can also view extended information about the session, such as session ID, ICMP type (if applicable), total incoming bytes, total outgoing bytes, total packets count, and how long the session has been active.

[Table 66 on page 729](#) describes all of the information that is available from the Active Sessions view:

Table 66: Active Sessions Items

Item	Description
Session ID	A unique identifier specified with the active session.
Source IP	IP address of the sending node of the connection.
Source Port	Port number of the sending node of the connection.
Destination IP	IP address of the receiving node of the connection.
Destination Port	Port number of the receiving node of the connection.
Translated IP	Translated IP address.
Translated Port	Translated port number.
Duration (sec)	Length in seconds of the connection session.
Policy ID	A unique identifier specified when the policy was configured. <i>None</i> means no name was specified during policy configuration.
Protocol ID	A unique identifier specified when the protocol was configured.
ICMP Type	The type of ICMP protocol.
Bytes In	The total number of bytes sent in.
Bytes Out	The total number of bytes sent out.
Total Packets	The total number of packets sent.
Duration	The length in seconds of the connection session.
Start Time	The time that the session started.

Using the Session Filter

You can control the information that is provided in the Active Sessions view by configuring a session filter. You can use the session filter to fetch specific sessions on a security device that match specific criteria that you set. The session filter defines the overall data set that you can view from the Active Sessions view. After you configure and apply the session filter, you can configure additional session display filters to view more specific session information.

Configuring the Session Filter

To configure the session filter:

1. Use the **Options** menu, and select **Session Filter**. The **Session Filter** dialog box appears.
2. Click the Long Form check box to display additional information about the Active Session.
3. Click the Maximum number of sessions to retrieve check box and enter the total number of sessions you want the Session Filter to retrieve.
4. Specify criteria for the sessions that you would like to view. You can specify an active session according to the following:
 - Source, Destination, and Translated IP (IP Address, Net Mask, and Port Range)
 - Session Duration
 - Session Start Date and Time
 - Policy ID
 - Session Type
 - Protocol ID
 - Policy with Logging

Click **More** to view additional criteria.

Click **Reset to Default** to reset all criteria back to their default settings.

Click **OK** when you are done.

Click **Refresh** to apply the criteria to the active session table view.

Configuring a Session Display Filter

You can apply a session display filter to view only specific active sessions.

1. From the **Options** menu, select **Session Display Filter**. The **Session Filter** dialog box appears.
2. From the **Source** tab, you can specify the sessions that you want to view according to the Source IP Address and Port number, or Port Range.
3. Click in the **Destination** tab to specify the sessions that you want to view according to Destination IP Address and Port number, or Port Range.
4. Click in the **Translated** tab to specify the sessions that you want to view according to Translated IP Address and Port number, or Port Range.
5. Click in the **Protocol** tab to specify the sessions that you want to view according to protocol.
6. Click in the **Other** tab to specify the sessions that you want to view according to Session Duration, Session Start Time, or Policy ID.
7. Click **OK** when you are done.
8. Click **Refresh** to apply the Session Display criteria to the active session table view.

Troubleshooting

From the **Device Manager**, you can right-click on any security device and select the **Troubleshoot** option to issue common CLI commands (such as **get**, **exec**, **debug**) to a security device using Telnet or a Secure Command Shell to troubleshoot problems. You can also add, delete, edit or search for custom CLI commands using the **Add/Delete Custom Troubleshoot Commands** option under the **Tools** menu in the **Device Manager** window.

To troubleshoot a device:

1. From the **Device Manager**, right-click on any device and select **Troubleshoot**. The **Troubleshoot Device** window appears.
2. Select the appropriate command from the list of **Predefined Commands** in the **Shortcuts** window. The command appears in the **Command** field.
3. Click on the **Execute Command** button. The status of the command appears in the field below. All commands that you execute appear in the **History** window.

You can use the **Add to Shortcut** or **Remove from Shortcut** buttons in the **Troubleshoot Device** window to create a list of CLI command shortcuts.



NOTE: Commands from NSM originate from the UI client to the security device. If you intend to issue **get** commands from NSM, you must plan and implement security policy rules in your network accordingly.

Viewing High Availability (HA) Statistics (if applicable)

If you have configured security devices to be highly available, you can view NSRP-related statistics on the device by accessing the HA Statistics view. [Table 67 on page 731](#) describes all of the information that is available from the HA Statistics view:

Table 67: HA Statistics View

Item	Description
VSD Group ID	The group ID that is associated with the VSD (or RTO).
Number of Units	The number of units associated with the VSD (or RTO).
State Change Counter	The number of times a security device changes operational states.
Init Counter	The transient state of a VSD (or RTO) group member while it was in the process of joining the VSD (or RTO) group.
Master	The number of master security devices.
Primary BackUp	The number of primary backup security devices.
BackUp	The total number of backup security devices.

Table 67: HA Statistics View (continued)

Item	Description
Ineligible	An administrator purposefully assigned a security device so that it cannot participate in selecting a new master security device.
InOperable	A VSD (or RTO) group security device has an internal problem.
Master Conflict	The number of conflicts that occurred on the master security device.
Primary Backup Conflict	The number of conflicts that occurred on the primary backup security device.
Tx Heartbeat	The number of transmitted heartbeats on the security devices.
Rx Heartbeat	The number of received heartbeats on the security devices.

Monitoring IDP Sensors

Use the Device Monitor to get an at-a-glance view of the current status of all the IDP sensors in your network.

Viewing IDP Device Status

[Table 68 on page 732](#) lists and describes information about IDP sensors that you can view through the Device Monitor:

Table 68: Device Status Information

Column	Description
Name	Unique name assigned to the sensor in NSM.
Domain	Domain in NSM in which the sensor is managed. NOTE: If you have configured multiple subdomains, you can view all your managed devices from the global domain.
Platform	Model number of the sensor.
OS Version	IDP firmware version running on the sensor.

Table 68: Device Status Information (continued)

Column	Description
Config Status	<p>Current configuration status of the sensor in NSM:</p> <ul style="list-style-type: none"> • None. No state has been set (does not show in Device Monitor). • Modeled. The sensor exists in NSM, but a connection to the sensor has not yet been established. • RMA. Equivalent to bringing the sensor into the Modeled state. RMA results from an administrator selection in the User Interface when a sensor goes down. • Waiting for 1st connect. NSM is waiting for the sensor to connect. You must enter a command on the sensor to make it connect to NSM. • Import Needed. You must import the configuration of the sensor into NSM. When you add a sensor for the first time, verify that your status indicates "Import Needed" before you attempt to import the sensor. During migration, this state indicates that import of the sensor configuration is still required. • Await Migration. After you have migrated your data in IDP Manager, the Config Status on each sensor displays that it is awaiting migration. It remains in this state until you have migrated the sensor. • Update Needed. An update to this sensor is required. • Managed. The sensor is currently being managed by NSM.
Connection Status	<p>Connection status of the sensor in NSM:</p> <ul style="list-style-type: none"> • Up. Sensor is currently connected to NSM. • Down. Sensor is not currently connected to NSM, but has connected in the past. • Never Connected. Sensor has never connected to NSM. <p>The Device Server checks the connection status of each sensor every 120 seconds by default. You can change this behavior by editing the value for the <code>devDaemon.deviceHeartbeatTimeout</code> parameter in the Device Server configuration file. Refer to the <i>Network and Security Manager Installation Guide</i> for more information on editing configuration files.</p> <p>Note: If the network connection goes down for a period longer than six to eight minutes, the sensor connection will permanently time out. If this occurs, and the device goes down for any reason, the sensor still appears as Up in the Device Monitor.</p>
First Connect	The first time the sensor connected to the NSM Device Server.
Latest Connect	The last time the sensor connected to the NSM Device Server.
Latest Disconnect	The last time the sensor disconnected from the NSM Device Server.

Viewing IDP Device Detail and Statistics

If a sensor is running, you can view additional status using Device Details and view traffic-related statistics and other information using Device Statistics.



NOTE: If a sensor has never connected, the Device Detail Status and Device Statistics views for the sensor are not available.

Viewing IDP Device Details

Double-click on any IDP sensor to view additional details on the sensor related to resource usage. You can also right-click the sensor and select **View Details**.



NOTE: The information in the Details window appears slightly different for firewall/VPN devices and IDP sensors. Details for IDP sensors include an additional tab showing process status.

Table 69 on page 734 lists and describes the information that you can view for an IDP sensor through the Device Detail Status:

Table 69: IDP Device Detail Status Items

Item	Description
OS Version	IDP firmware version running on the sensor.
Mode	Current operation mode of the device.
CPU Idle	Percentage of the time the CPU was idle.
CPU User	Percentage of CPU utilization that occurred while executing at the user level.
CPU Kernel	Percentage of CPU utilization that occurred while executing at the system level.
CPU Usage	Percentage of CPU utilization.
1 Min. Load	One minute load average.
5 Min. Load	Five minute load average.
15 Min. Load	Fifteen minute load average.
Total Mem	Total amount (in megabytes) of memory.
Used Mem	Amount (in megabytes) of used memory.
Mem Usage	Percentage of used memory.
Total Swap	Total amount (in megabytes) of swap space.
Used Swap	Amount (in megabytes) of used swap space.

Table 69: IDP Device Detail Status Items (continued)

Item	Description
Swap Usage	Percentage of used swap space.

Viewing IDP Process Status

For IDP sensors, use the **Process Status** tab to view information on various processes running on the sensor.

[Table 70 on page 735](#) lists and describes the information that you can view for an IDP sensor through the Process Status:

Table 70: IDP Sensor Process Status Items

Item	Description
Process Name	Name of the process running on the sensor.
Total Mem Usage	Amount (in megabytes) of memory used.
Phys Mem Usage	Amount of memory (in kilobytes) a process currently has in physical memory (not in swap).
CPU Usage	Percentage of CPU used.

Viewing IDP Device Statistics

If a sensor is running, you can also access the Statistics view to access traffic and other system-related information on the device. To view statistics on a particular sensor, right-click the sensor in either the **Device Monitor** or the **Device Manager** and select **View Statistics**. The **Device Statistics Summary** appears in a new window.



NOTE: The information in the Device Statistics window appears slightly different for firewall/VPN devices and IDP sensors.

IDP Device Statistics Summary

The **Device Statistics Summary** displays the following details:

- Details describing the sensor, for example, firmware version and mode.
- Packet and flow information

[Table 71 on page 736](#) details additional information you can view from the Device Statistics Summary for IDP sensors.

Table 71: Device Statistics Summary (for IDP Sensors)

Item	Description
OS Version	IDP firmware version running on the sensor.
Mode	Current operation mode of the device.
ICMP Packets	Total number of ICMP packets.
TCP Packets	Total number of TCP packets.
UDP Packets	Total number of UDP packets.
Other Packets	Total number of other packets.
ICMP Flows	Total number of ICMP flows.
TCP Flows	Total number of TCP flows.
UDP Flows	Total number of UDP flows.
Other Flows	Total number of other flows.

Monitoring VPNs

Use VPN Monitor to get an at-a-glance status of the up/down status of VPN tunnels as well as other statistics relevant to your VPN.



NOTE: You must enable the **VPN Monitor** option on the tunnel when configuring the tunnel for the device.

Viewing the VPN Status Summary

The VPN Monitor lists a summary of all the VPN tunnels that have been implemented in your system. It includes visual indicators that depict whether an existing VPN tunnel is Up, Down, or Not Monitored. The Summary also includes information describing the VPN name, VPN type, source, destination, security parameter index, IP address, and protocol.

Table 72: VPN Tunnel Summary

Column	Description
VPN	Name of the active VPN.
VPN Type	Type of tunnel: Dialup or Site-to-Site.

Table 72: VPN Tunnel Summary (continued)

Column	Description
From Hostname (IP)(Vsys)	Source security devices used in the VPN. For example, a root security device named NS5000 with an IP address of 1.1.1.1 appears as NS5000(1.1.1.1). For a Vsys 1, "NS5000(1.1.1.1)(1)" appears.
From Domain	Domain in NSM in which the source security device used in the VPN is managed.
To Hostname(IP)(Vsys)	Destination security devices used in the VPN. For example, a root security device named NS5000 with an IP address of 1.1.1.1 appears as NS5000(1.1.1.1). For a Vsys 1, "NS5000(1.1.1.1)(1)" appears.
To Domain	Domain in NSM in which the destination security device used in the VPN is managed.
Status	VPN Status: Up or Down
SPI (in/out)	Security Parameter Index (SPI) key into and out of the active VPN. This is the encryption method.
IP (Local-Peer)	Peer gateway IP address for the active VPN.
Protocol	Protocol used for the active VPN
Peer GW ID	Peer gateway ID for the active VPN.

Configuring a VPN Filter

You can configure a VPN filter to control the information that is provided in the VPN Monitor. You can view VPN information related to the type, status, or the specific security device or virtual system associated with the VPN tunnel that you want to view.

To create a VPN filter:

1. Select **Report Manager > VPN Monitor**
2. From the **Option** menu, select **VPN Filter**.
The **VPN Filter** dialog box is displayed.
3. Click the add icon.
The **New VPN Filter** dialog box is displayed.
4. In the Type area, select **RAS**, **Site to Site**, or **Default** to view both types of VPN.
5. In the Device/Vsys area, select the type of device or virtual system associated with the VPN tunnel. Select **Include all selected devices**.



TIP: In the Selected Devices/Vsys area, by default, all devices or virtual systems are included in the filter. To improve system performance, you can remove devices or virtual systems by selecting them and clicking **Remove**. Next, select **Exclude all selected devices**.

7. Click **OK** to save the filter.

The new VPN filter is updated in the Filter Table.

8. Click **Apply** to activate the selected filter. The filter results are shown in the VPN Monitor view.

Modifying a VPN Filter

To modify a VPN filter:

1. Select **Report Manager > VPN Monitor**
2. From the **Option** menu, select **VPN Filter**.

The **VPN Filter** dialog box is displayed.

3. Select the edit icon.

The **View VPN Filter** dialog box is displayed.

4. Change the existing field values and click **OK**.

The modified field values are updated in the Filter Table.

Deleting a VPN Filter

To delete a VPN filter:

1. Select **Report Manager > VPN Monitor**
2. From the **Option** menu, select **VPN Filter**.

The **VPN Filter** dialog box is displayed.

3. From the Filter Table, select the row for the filter that you want to delete.

4. Select the delete icon.

The selected filter is deleted.

Configuring a VPN Display Filter

You can control the information that is provided in the VPN Monitor by configuring a VPN display filter. From the **Options** menu, select **Display Filter** to configure a VPN display filter. You can view VPN information related to the type, status, or the specific security device or virtual system associated with the VPN tunnel that you want to view. Click the **Refresh** button to apply the Session Display criteria to the active session table view.

Viewing Active VPN Details

To view the details on the active VPN, click to select the VPN. From the **View** menu, select **Active VPN Details** (alternatively, you can also right-click the VPN tunnel and select **Active VPN Details**).

Refer to “[Viewing Active VPN Information](#)” on page 720 for more information on the Active VPN Details table.

Viewing Device-Specific VPN Information

To view security device-specific information about your VPN, right-click the VPN tunnel and select **Monitor Data** and then select the security device. A Monitor info window appears where you can access the VPN Monitor table, Active VPN table, and a chart, which enable you to view the distribution of VPN tunnels on the security device.

Monitoring NSRP Statistics

If you have implemented NetScreen Redundancy Protocol (NSRP) for the purpose of deploying clusters for redundancy, you can use the NSRP Monitor to get an at-a-glance status of your Juniper Networks systems that are in clusters. These systems include both the NetScreen-500 and the NetScreen-1000. To launch the NSRP Monitor, click **NSRP Monitor**.

Viewing NSRP Summary Information

Double-click an NSRP device to view a summary of the top-level information on the selected cluster. From the NSRP Summary, you can view the following details about a specific cluster:

- Key details describing the cluster (such as name, number of VSDs, and number of RTOs)
- Security details
- The total number and type of events

[Table 73 on page 739](#) describes all of the information that is available from the NSRP summary:

Table 73: NSRP Device Summary

Item	Description
Cluster	Name of this cluster.

Table 73: NSRP Device Summary (continued)

Item	Description
Domain	Domain in NSM in which the NSRP device is managed.
No of VSD's	The total number of virtual security devices (VSDs) that are attached to this cluster.
No of RTO's	The total number of runtime objects (RTOs) that are attached to this cluster.
Encryption	Whether encryption has been enabled or disabled.
Authentication	Whether authentication has been enabled or disabled.
No. of Gratuitous arps	The number of gratuitous ARPs.
Critical Events	The total number of Critical events that occurred.
Major Events	The total number of Major events that occurred.
Minor Events	The total number of Minor events that occurred.
Warning Events	The total number of Warning events that occurred.
Intermediate Events	The total number of Intermediate events that occurred.
Clear Events	The total number of Clear events that occurred.

Viewing VSD/RTO Information

Double-click the cluster security device icon or click the **+** icon that corresponds to the cluster security device icon to view the virtual security devices (VSD) and runtime objects (RTO) that have been attached to this cluster.

Click the **VSD** or **RTO** icon to see summary information describing the object. [Table 74 on page 740](#) describes the information available from the VSD/RTO summary:

Table 74: VSD/RTO Summary

Item	Description
Cluster	The name of the cluster associated with this VSD.
VSD(RTO)	The name of this VSD (or RTO).
No of Devices	The total number of security devices that are associated with this VSD.
Init Hold Time (sec)	The initial hold time state (in seconds) of the VSD.

Table 74: VSD/RTO Summary (continued)

Item	Description
Heartbeat Interval (ms)	The time interval (in milliseconds) between each heartbeat.
Heartbeat Lost Threshold (ms)	Threshold level required to change over to the backup security device.
Master	The master system.
Primary Backup	The primary backup system.

Viewing VSD Counter Details

Click the **Counters** tab to view specific information about your VSD counters. [Table 75 on page 741](#) describes the information that is available from the VSD counters view:

Table 75: VSD Counter Details

Item	Description
Device	The devices that are associated with the VSD or RTO.
Number of Units	The number of units associated with the VSD or RTO.
State Change Counter	The number of times a device changes operational states.
Init Counter	The transient state of a VSD or RTO group member while it was in the process of joining the VSD or RTO group.
Master	The number of master devices.
Primary BackUp	The number of primary backup devices.
BackUp	The total number of backup devices.
Ineligible	An administrator purposefully assigned a device so that it cannot participate in selecting a new master device.
InOperable	A VSD or RTO group device has an internal problem.
Master Conflict	The number of conflicts that occurred on the master device.
Primary Backup Conflict	The number of conflicts that occurred on the primary backup device.
Tx Heartbeat	The number of transmitted heartbeats on the devices.
Rx Heartbeat	The number of received heartbeats on the devices.

Viewing RTO Counter Details

Click the **Counters** tab to view specific information about your RTO counters. [Table 76 on page 742](#) describes the information that is available from the RTO counters view:

Table 76: RTO Counters Details

Item	Description
Device	The devices that are associated with the RTO.
Member ID	The member identification associated with this RTO.
Status	The current status of the RTO: Active or Down.
Direction	The direction of the RTO: In or Out.
Lost Heartbeat	The number of heartbeats not received from the RTOs peers.
Counter to Active	The number of times that the RTO was placed to Active.
Counter to Set	The number of times that the RTO was placed to Set.
Counter to Lost Peer	The number of times that the RTO was placed to Lost Peer.
Counter to Group Detach	The number of times that the RTO was placed to Group Detach.

Monitoring IDP Clusters

If you have implemented IDP clusters for the purpose of redundancy, you can use the IDP Cluster Monitor to get an at-a-glance status of your IDP sensors that are in clusters.

[Table 77 on page 742](#) describes all of the information available from the IDP Cluster Monitor:

Table 77: IDP Cluster Monitor

Item	Description
Name	Name of the cluster.
Status	Status of the cluster (OK, Warning, or Fail).
Domain	Domain in NSM in which the source IDP cluster is managed.

Viewing IDP Cluster Summary Information

Click **IDP Cluster Monitor** to view a summary of the top-level information on all IDP clusters. From the IDP Cluster Summary, you can view the following details about a specific cluster:

- Key details describing the cluster (such as name, Status, Cluster ID)
- HA mode

[Table 78 on page 743](#) describes all of the information that is available from the IDP Cluster summary:

Table 78: IDP Cluster Summary

Item	Description
Domain	Domain in NSM in which the source IDP cluster is managed.
Name	Name of the cluster.
Cluster ID	Number uniquely identifying a cluster on a given Ethernet segment (retrieved from all nodes)
HA Mode	Whether the Cluster is in Hot-standby or Load-Sharing mode.
Total Members	Total number of IDP sensors that are associated with the cluster.
Heartbeat Interval	Time interval (in ms) between each heartbeat
Heartbeat Lost Threshold	Threshold level (number of heartbeat intervals) required to declare that a device in the cluster has gone down.
No. of OK Members	Number of cluster members that are in OK state.
No. of Failed Members	Number of cluster members that are in FAIL state.
No. of Initializing Member	Number of cluster members that are in INIT state.
Master	Name of the master node.
Backup Availability	Whether a backup is available in the event that the master node goes down.
No. of Backup Members	Number of active backup devices.

Monitoring IDP Cluster Members

Click any IDP cluster to view details of each member in the cluster.

[Table 79 on page 744](#) describes all of the information that is available from the IDP Cluster Member Monitor.

Table 79: IDP Cluster Member Monitor

Item	Description
Name	Name of the device.
Status	State of the device (INIT, OK, or FAIL).
Domain	Domain in NSM in which the IDP cluster member is managed.

Using the Realtime Monitor

The following example describes a typical use case for monitoring your security devices, VPNs, and NSRP clusters in NSM.

In this example, you are a network administrator responsible for monitoring the day-to-day operation of all the security devices managed in your network. You are using NSM to manage your network, and **Realtime Monitor** to monitor the up/down connection status of all your security devices.

One day, you notice that the Connection Status on a mission-critical security device indicates that the security device is DOWN. You wait several minutes to see whether the connection status resolves itself because intermittent network problems might cause a security device to temporarily indicate as DOWN. The Device Monitor still indicates that the security device is DOWN.

You next try to ping the security device. If you are successful in reaching the device, you can send a get status command to check the status of the security device.

If you cannot ping the security device, you need to investigate further. You scan the Log Viewer for the log entry indicating that the security device has disconnected. You can filter the log entries in the Log Viewer to display only the log entries generated for the security device during the time just before it went down. Viewing these log entries will also provide you with context for events leading to the security device disconnection. This will help you to determine the cause of the problem.

You notice several very suspicious log entries that indicate that this security device may have been the target of an attack. You flag the log entries using the predefined flag types in the Log Viewer, and assign them to your security experts for further investigation.

Monitoring the Management System

Use the Server Manager to access, configure, and monitor the NSM management system. The management system includes a GUI Server and Device Server. Refer to the *Network and Security Manager Installation Guide* for more information about the GUI Server and Device Server.

The Server Manager contains the following:

- Servers

- Server Monitor (Machine-wide info)
- Schema Information

Configuring Servers

Use Servers to add, configure and view key information about the GUI Server and Device Server:

[Table 80 on page 745](#) lists and describes Device Server and GUI Server information that you can view from Servers .

Table 80: Server Information

Item	Description
Name	Name of the GUI Server or Device Server.
Server Type	Whether the current server is Device Server, Device Server Cluster, GUI Server, or GUI Server Cluster.
IP Address	IP address of the server.
Device Server Manager Port	The port open on the Device Server for security devices running ScreenOS 5.0 and later. (Read Only)
IDP Device Server Manager Port	The port open on the IDP Device Server for security devices.
DMI Device Server Manager Port	The port open on the DMI Device Server for security devices.
IP Address of Secondary Server	IP address of the secondary server.

Configuring Device Servers

You can also add and configure a Device Server. You might need to configure a Device Server when installing the GUI Server and Device Server on separate servers, or when installing the management system with High Availability (HA) enabled.

You can configure the following parameters on a Device Server:

- Name—Name of the Device Server.
- IP Address—IP address of the Device Server.
- Server Type—Either Device Server or Device Server Cluster. If you are installing the management system with HA enabled, you need to configure the Device Server as part of an HA Cluster. After you specify that a Device Server will act as a Device Server Cluster, you can access additional tabs allowing you to further configure cluster details including the IP Address and port number of the secondary server, and e-mail notification.

- **Device Server Manager Port**—The port is set to 7800 by default. This field is read only.
- **IDP Device Server Manager Port**—The port is set to 7803 by default. This field is read only.
- **DMI Device Server Manager Port**—The port is set to 7804 by default. This field is read only.
- **Device Server ID**—A unique ID assigned by NSM to each Device Server.
- **Mapped IP Address (MIP)**—If applicable, you can define multiple mapped IP addresses.
- **Device Polling**—The Device Server polls security devices it manages for Device, VPN, NSRP, or Interface statistics every 300 seconds by default. If you wish to change this behavior, you can edit the interval, using the Device Polling tab.
- **High Availability (HA)**—To configure a secondary Device Server, you need to specify the IP Address and port, and Mapped IP Address (if applicable).
- **E-mail Notification**—You can configure an SMTP server to send you an e-mail notifying you of various events on the Device Server.
- **Disk and Log Management**—You can enable an alert and action when available disk space reaches a specified minimum limit.

Refer to the *Network and Security Manager Installation Guide* for more information on adding and configuring the Device Server.

Configuring the GUI Server

Table 81 on page 746 lists and describes device information that you can view through the Server:

Table 81: GUI Server Table

Item	Description
Name	Name of the GUI Server.
Server Type	Whether the current server is a GUI Server or GUI Server Cluster.
IP Address	IP address of the GUI server.
IP Address of secondary server.	IP address of the secondary server.

You can configure the following parameters for the GUI Server:

- **Server Type**—Select GUI Server or GUI Server Cluster. If you are installing the management system with HA enabled, you need to configure the GUI Server as part of an HA Cluster.
- **IP Address**

- HA parameters
- E-mail Notification—You can configure an SMTP server to send an e-mail to notify you of events on the GUI Server. (See “Sending E-mail Notification of Downed Device” on page 840 for details.)



NOTE: You must restart the GUI Server to apply any changes that you have made.

Using Server Monitor

You can use the Server Monitor to view the status of the running GUI Server and Device Server. The Server Monitor lists all GUI Servers and Device Servers in your management system. For example, if you have installed a primary and secondary GUI Server in a high availability configuration, you can use the Server Monitor to monitor which GUI Server is currently active.

The Server Monitor provides two categories of information:

- Server status—Displays information about the GUI Server or Device Server's status, CPU, and memory. You can also choose to view the status of each server in the Server Monitor, or view additional server status details in a separate dialog box.
- Process status—Displays information about the individual processes on a GUI Server or Device Server.

Viewing Server Status

To view the status of any server in the management system, select **Server Manager** in the navigation tree, and then select **Server Monitor (Machine-wide Info)**.

Figure 109 on page 747 shows the Server Monitor Window.

Figure 109: Server Monitor (Machine-wide Info)

Server Monitor (Machine-wide Info) - Juniper Networks - NSM - global : current										
File View Devices Tools Help										
Server Monitor (Machine-wide Info)										
Device Server										
Name /	Server Ty...	Status	CPU	Mem	Disk	CPU Usage	Disk Usage	Peer Devi...	Active Se...	
server_1	Device Server	OK	OK	OK	OK		1%	N/A	Primary	

GUI Server										
Name /	Server Ty...	Status	CPU	Mem	Disk	CPU Usage	Disk Usage	Peer GUI ...	Active Se...	Last GUI ...
server_0	GUI Server	OK	OK	OK	OK		1%	N/A	Primary	N/A

Table 82 on page 748 lists and describes the columns that appear in the Server Monitor:

Table 82: Server Monitor (Machine-wide Info) Data

Indicator	Description
Name	Name of the GUI Server or Device Server.
Server Type	Whether the current server is a GUI Server, GUI Server Cluster, Device Server, or Device Server Cluster.
Status	<p>Status of the server based on CPU or memory utilization:</p> <ul style="list-style-type: none"> • OK • Warning • Critical • Down <p>Note: By default, the Status field for each server appears Green (OK) if the usage on either the CPU, memory, or disk is less than 90%. It appears Yellow (Warning) if the usage on either the CPU, memory, or disk is greater than 90%. You can edit the settings that apply for the Status indicators using Tools>Preferences>Alert Settings.</p>
CPU	<p>Status based on CPU utilization:</p> <ul style="list-style-type: none"> • OK (CPU usage < 90%) • Warning (CPU usage = 90-95%) • Critical (CPU usage > 95%) • N/A (when the server is down)
Mem	<p>Status based on memory utilization:</p> <ul style="list-style-type: none"> • OK (memory usage < 99%) • Warning (memory usage > 99%) • Critical (memory usage = 100%) • N/A (when the server is down)
CPU Usage	Percentage of CPU used.
Peer Device Server State	State of the server's peer server (only applicable if you have added a secondary server and configured it in an HA Cluster).
Active Server	Whether the currently active server is the primary or secondary server in an HA Cluster.
Last GUI Server Replication Time	Time of day that the GUI Server database was last replicated to its peer server.

You can sort data in the Server Monitor according to any column header by clicking that column.

Viewing Additional Server Status Details

If you are interested in monitoring additional details about your server's status, you can view the **Server Detail Status** window by double-clicking any of the servers that appear

in the Server Monitor. You can also right-click anywhere on the **Server Monitor** and select **View Details**.

Table 83 on page 749 describes information available in the Server Detail Status:

Table 83: Server Detail Status

Name	Description
OS	Operating system running on server machine.
Type	The server's machine processor type.
CPU Idle	Percentage of the time the CPU was idle.
CPU User	Percentage of CPU utilization that occurred while executing at the user level.
CPU Kernel	Percentage of CPU utilization that occurred while executing at the system level.
CPU Usage	Percentage of CPU utilization.
1 Min Load	One minute load average.
5 Min Load	Five minute load average.
15 Min Load	Fifteen minute load average.
Total Mem	Total amount (in megabytes or gigabytes) of memory.
Used Mem	Amount (in megabytes or gigabytes) of used memory.
Mem Usage	Percentage of used memory.
Total Swap	Total amount (in megabytes or gigabytes) of swap space.
Used Swap	Amount (in megabytes or gigabytes) of used swap space.
Swap Usage	Percentage of used swap space.

Viewing Process Status

From the Server Monitor, you can also view the status of all running server processes on the GUI Server or Device Server. This view is useful for troubleshooting. If you are having problems with the server, you can quickly identify if a specific process on the server is the source of that problem.

To view process status, select **Server Manager** in the navigation tree, and then select **Server Monitor**. Double-click the **Server Monitor** or click the node to expand the navigation tree. You can also right-click the **Server Monitor** to open it in a new window. Click to select a server to view the status of the processes running on it.

Figure 110 on page 750 shows process status for the Device Server.

Figure 110: Process Status for the Device Server

Name	Status	Total Mem Used	Phys. Mem Used	CPU Usage	Version
devSvrDataCollector	Up	670 MB	19 MB	0%	Version NSM 2004 FP3
devSvrDirectiveHandler	Up	691 MB	8 MB	0%	Version NSM 2004 FP3
devSvrLogWalker	Up	5 MB	3 MB	0%	v1.3.1 (build LGB3z1bm)
devSvrManager	Up	16 MB	4 MB	0%	v1.3.1 (build LGB3z1bm)
devSvrStatusMonitor	Up	5 MB	3 MB	2%	v1.3.1 (build LGB3z1bm)

Figure 111 on page 750 shows process status for the GUI Server.

Figure 111: Process Status for the GUI Server

Name	Status	Total Mem Used	Phys. Mem Used	CPU Usage	Version
guiSvrDirectiveHandler	Up	650 MB	15 MB	0%	Version NSM 2004 FP3
guiSvrManager	Up	58 MB	46 MB	0%	v1.3.1 (build LGB3z1bm)
guiSvrMasterController	Up	649 MB	7 MB	0%	Version NSM 2004 FP3
guiSvrStatusMonitor	Up	5 MB	3 MB	2%	v1.3.1 (build LGB3z1bm)

Table 84 on page 750 lists and describes the information that appears in the Process Status:

Table 84: Process Status

Name	Description
Name	Name of the GUI Server or Device Server process.
Status	Displays if the process is Up or Down.
Total Mem Used	Total amount (in megabytes) of memory utilized.
Phys Mem Used	Total amount (in megabytes) of physical memory utilized.
CPU Usage	Percentage of CPU utilized.
Version	Process version.

You can sort server monitor data according to any column header by clicking that column.

Using Management System Utilities

Table 85 on page 751 describes management system utilities.

Table 85: Management System Utilities

Name	Description
logcount.sh	<p>Provides information on peak/average logging rate, total log database size, and average log size.</p> <p>This utility is located on the Device Server at /usr/netscreen/DevSvr/utls</p>
setSrsDbParams.sh	<p>Configures the GUI Server for connection to the Statistical Report Server.</p> <p>This utility is located on the GUI Server at /usr/netscreen/GuiSvr/utls</p>
tech-support.sh	<p>Collects and compresses technical support data.</p> <p>This utility is located in the utls directory on both the Device Server and GUI Server.</p>
xdbAuditLogConverter.sh	<p>Exports Audit Log data to a csv file or Syslog server.</p> <p>CSV Command Usage:</p> <p>For csv file, issue the csv command: ./xdbAuditLogConverter.sh <xdb root> csv [csv full file path]</p> <p>For example, ./xdbAuditLogConverter.sh /usr/netscreen/GuiSvr/var/xdb csv /tmp/audit.csv</p> <p>This creates a file called audit.csv in the /tmp directory with audit logs in csv format. If the csv file path is not specified, audit logs in csv format are written to a default file called auditlog.csv in the current directory.</p> <p>Syslog Command Usage:</p> <p>To export data to a Syslog server, issue the syslog command: ./xdbAuditLogConverter.sh <xdb root> syslog [remote IP Address]</p> <p>For example, ./xdbAuditLogConverter.sh /usr/netscreen/GuiSvr/var/xdb syslog 172.23.9.94</p> <p>If you want to syslog to the host machine on which you are running this command, do not specify an IP address. For example, ./xdbAuditLogConverter.sh /usr/netscreen/GuiSvr/var/xdb syslog</p>

In NSM, enhancements to the audit log exporter tool allow you to:

- Invoke detailed help messages from the audit log exporter tool with `./xdbAuditLogConverter - help`
- Use `--showdiff` to view audit log details before and after a modification. This feature is only supported on objects such as addresses, NAT, VLAN objects, access profiles, VSYS profiles and so on. It is not supported on policies.
- View details about modifications to objects (add, delete, modify actions) in the audit log output. For example,
`ACTION=delete,ADMIN=super,DEVICE=,TARGET=addressObj/address/global,TIME= Tue Dec 30 11:02:49 2008`
- Filter exported audit log data according to Date/Time and Date/Time Range, Device, Action, Admin, Domain, Working Domain. Multiple filters are also allowed. For example,
`./xdbAuditLogConverter <xdb path> <csv/syslog> [CSV full filepath | Remote IP Address] - --domain=<domain id> - --device =<device name>`

Using Schema Information

From NSM, you can select **Schema Information** to view current and running schema and update schema for devices whose schema are defined using XML.

For information about downloading and activating XML-based device schema without the need to upgrade NSM, see [“Managing Device Schemas Through the Juniper Update Mechanism” on page 338](#)



NOTE: This feature does not apply for ScreenOS or IDP devices.

Viewing Device Schema

To view current and running schema:

1. In the User Interface, click **Administer**.
2. In the navigation tree, select **Server Manager > Schema Information**

The main display area displays the current staged and running schema details. The staged schema is the most current schema available for download. The running schema is the schema currently applied in NSM.

CHAPTER 18

Analyzing Your Network

You can use the Security Monitor module to learn about your internal network. The Security Monitor includes the Dashboard, Profiler, and Security Explorer network analysis tools.

This chapter contains the following sections:

- [About the Dashboard on page 753](#)
- [About the Profiler on page 753](#)
- [Setting Up the Profiler on page 755](#)
- [About Profiler Views on page 760](#)
- [Recommended Profiler Options on page 771](#)
- [Accessing Data in the Profiler Database on page 775](#)
- [About Security Explorer on page 775](#)
- [Using Security Explorer on page 779](#)

About the Dashboard

Known targets and sources of attacks or suspected targets and sources of attacks can be added to source or destination watch lists. The Dashboard is a near-real-time monitor of these watch lists and the top 10 attacks within the previous hour. The interval at which these lists are updated ranges from 2 minutes (default rate) to 30 minutes. The lists are updated automatically after the refresh rate is configured.

Use the Dashboard to create and configure both a destination and a source watch list as well as to display device status. This feature is available by default to the system administrator. Use the watch list to add or modify known or suspected targets and sources of attacks.

About the Profiler

The Profiler is a network-analysis tool that helps you learn about your internal network so you can create effective security policies and minimize unnecessary log records. After you configure the Profiler, it automatically learns about your internal network and the elements that constitute it, including hosts, peers (which host is talking to which other

host), ports (non-IP protocols, TCP/UDP ports, RPC programs), and Layer-7 data that uniquely identifies hosts, applications, commands, users, and filenames.

The Profiler is supported in all IDP modes and in HA configurations, and it queries and correlates information from multiple devices.

To use the Profiler, you must first configure the networks and hosts on your internal network that you want to monitor. The device monitors traffic at the network and application levels. You can use this data to investigate and analyze potential problems in the network and to resolve security incidents.

During profiling, the device records network activity at Layer-3, Layer-4, and Layer-7 and stores this information in a searchable database called the Profiler DB. The device uses session creation, session teardown, and protocol contexts to generate this database, which defines all unique activities occurring on your network. Unique activities include attempts, probes, and successful connections. The device logs normal events only once, and it logs all unique events as often as they occur. A normal event is an event that reoccurs frequently and does not change. A unique event is an event that is new, unexpected, or does not match the normal traffic patterns of your network.

Example of Unique Events

For example, you allow users to use a laptop to connect to the corporate network while working in a conference room.

- **Normal Event.** Wendy holds a meeting every Tuesday at 4:00 PM in conference room A. Every meeting, she connects her laptop to the network and accesses documents on the primary fileserver. Because the same event occurs multiple times, the device logs the event once and includes a timestamp that indicates the first and last times Wendy accessed the network from conference room A.
- **Unique Event.** The device logs changes from normal activity as a unique event in the Profiler.
 - During one of Wendy's Tuesday meetings, she discovers she needs a document that resides on the Engineering server. She connects to that server and downloads the needed files. Because this connection differs from her usual activity, the device logs it as a unique event and records the IP and MAC addresses for both Wendy's laptop and the Engineering server.
 - The device also logs other unique qualifiers, such as user name and e-mail address for each individual that participated in the connection. If Wendy is out sick and another person logs into her laptop to run the meeting, the device records the connection as a unique event because the user name has changed.

To see all normal and unique events on your network, you configure and start the Profiler on multiple devices. This enables the Profiler to aggregate and display a complete view of your internal network.



NOTE: Profiler DBs remain on individual devices even if the devices restart.

After your devices have started profiling, you can begin to use the profiled data to perform the following tasks:

- Set a network baseline— A baseline can help you track the servers and hosts on the network, as well as the protocols and services those components use to communicate. By immediately locating new components on your network, you can ensure that those components are protected (with a security policy) and that you can track their status (with the Profiler). For details, see [“Configuring a Network Baseline” on page 771](#).
- Update vulnerable systems—The Profiler uses passive fingerprinting to provide you with an inventory of operating-system and software applications, their versions, and what components use them. As new versions or security updates are announced, you must first determine if your network is affected, locate the affected components, and patch as appropriate. For details, see [“Keeping Your Network Current” on page 772](#).
- Immediately locate the source of an internal worm or trojan—The Profiler can show you exactly when the worm or trojan entered your network, how it was introduced, and which network components are infected. By filtering the profile data, you can quickly identify the source and contain the attack to minimize impact, then investigate and recover from any damage. For details, see [“Stopping Worms and Trojans” on page 773](#).
- Detect violations of your corporate security policy— The Profiler can help you confirm suspected violations such as rogue servers running on the network. Most of the time, however, you do not know exactly what you are looking for on the network. In these cases, it is easier to specify exactly what should be on the network, then detect any traffic that violates that specification. To detect violations, you can use a special type of object, called a permitted object, to define what you should see on the network.

The following sections detail how to set up, configure, and use your profiled data as described previously.

Setting Up the Profiler

Using the Profiler involves the following steps:

- Configure the Profiler to collect specific information about your internal network.
- Update Profiler Settings on the device after you configure the Profiler.
- Start the Profiler to enable your device to begin collecting data.
- Customize Profiler preferences.

You configure your device to collect specific information and compile it into the Profiler DB.



NOTE: Because devices collect data from network components on your internal network, it is helpful to create network objects to represent those components before you begin configuring the Profiler. Alternatively, you can create new network objects directly from the Profiler.

Configuring the Profiler

To configure the Profiler, use the Profiler settings that are available on the device settings in the Device Manager. From Device Manager, double-click on a device and select **IDP Profiler Settings**.

The **IDP Profiler Settings** dialog box appears with the **General** tab selected. After selecting the device you want to use for profiling, you can then configure how that device collects data from your internal network.

[Table 86 on page 756](#) describes the Profiler settings that you can configure from the **General** tab:

Table 86: General IDP Profiler Settings

Setting	Description
Enable IDP Profiling	Enables the IDP Profiler.
Enable Application Profiling	Enables the Profiler to collect and track application data.
Include Probe and Attempt	Enables the Profiler to collect and track specific probes and attempts.
Enable AVT	Enables the Profiler to perform Application Volume Tracking.
Include Non-tracked IP Profiles	Enables the Profiler to collect and track data from external hosts.
db limit (in MB)	Maximum database size for the Profiler on each device. By default, the maximum database size is 3 GB.
Enable OS Fingerprinting	Enables the Profiler to perform passive OS fingerprinting to determine the operating system of an end host. Not applicable for ISG family devices with IDP.
Refresh Interval (in secs)	Time interval (in seconds) that the Profiler refreshes OS fingerprinting. By default, the Profiler refreshes OS fingerprinting data every 3600 seconds (60 minutes). Not applicable for ISG family devices with IDP.

If you enable AVT in the Profiler settings, you can view AVT reports on ISG devices running ScreenOS 6.2 and later. From **My Reports** under the **Investigate** module, you can create a new AVT report format with your defined settings.

The AVT feature is limited by its dependency on the NSM agent's report delivery which might be unreliable, affecting the accuracy of information. Also, the AVT feature displays the cumulative count of all the traffic on a port, which could be over many sessions.

Enabling OS Fingerprinting

OS fingerprinting passively detects the operating system of an end-host by analyzing TCP handshake packets. To ensure that this works, you need to verify that OS fingerprinting is first enabled on the profiled device. After you have configured the Profiler with the tracked hosts, contexts, you must update the device.

OS fingerprinting works only for packets that contain a full-fledged TCP connection, one that has a SYN, a SYN/ACK, and a FIN connection. OS fingerprinting only works for operating systems that are supported on the device. A list of the supported operating systems is available on the device in a file called **fingerprints.set** in the **/usr/idp/device/cfg/** directory.

Configuring Network Objects

The first part of configuring the Profiler is to tell the device which network objects you want the device to profile. When you start the Profiler, the device begins collecting data from the selected hosts.

In the **Tracked Hosts** tab, select the network objects that represent your internal hosts. The device collects detailed information about traffic that passes between internal hosts, and groups traffic that does not match an internal host in a special IP: **73.78.69.84**. Communication between an internal host and an external host is recorded only once. For example, the device records internal host A communicating to www.yahoo.com and www.cnn.com as one entry in the Profiler DB.

You can select unlimited internal network objects.

You can also use the **Exclude List** tab to select the network objects that represent internal hosts you do not want to include in IDP profiling. You might want to exclude a host from the Profiler if you selected a group of network objects in the **Tracked Host** tab but want to exclude specific members of that group.

Configuring Context Profiles

Next, determine which contexts you want the device to record. In the **Contexts to Profile** tab, the context list includes only the contexts that can clearly identify a host, a user, or an application. Select contexts that the device profiles. When you start the Profiler, the device begins collecting data on traffic that matches the selected contexts.

Example: Selecting Contexts

To track FTP logins, usernames, and commands, select the FTP contexts in the **Contexts to Profile** tab. After the Profiler is started, the device begins collecting information about FTP logins, usernames, and commands, enabling you to quickly identify who is using FTP on your network and what they are doing over that protocol.

When you first configure the Profiler, select all contexts. This enables the device to collect data about every context on your network, giving you a complete view of your network traffic. Later, when you have analyzed your traffic, you can eliminate contexts that you know will not be used on your network.

Select **Profile Context** to include context information. If you clear Profile Context, IDP profile data only includes higher-level traffic data such as source, destination, and service. If you want Profiler information to include context values and network probes (for example, port scans), also configure the Profiler to “Include Probes and Attempts” in the **General** tab.

Configuring Alerts

Use the **Alert** tab to configure the Profiler to indicate the appearance of a new host, protocol, or port on your internal network. When you enable New Host Detected, New Protocol Detected, or New Port Detected, the device generates a specific log record, such as PROFILER_NEW_HOST, in the Profiler Logs section of the Log Viewer, when the device discovers a new host, protocol, or port.

If you are configuring the Profiler for the first time, do not enable the new host, protocol, or port alerts. As the Profiler runs, the device views all network components as new, which can generate unnecessary log records. After the Profiler has learned about your network and has established a baseline of network activity, you should reconfigure the device to record new hosts, protocols, or ports discovered on your internal network. For details, see [“Configuring a Network Baseline” on page 771](#).

Enable the **Database Limit Exceeded** alert to indicate when you have reached the maximum limit of the database size. You can configure the maximum limit of the Profiler DB using the dbLimit parameter in the **General** tab of the **Profiler Settings** dialog box. The default limit is the value that has been set for Profiler preferences (see [“Customizing Profiler Preferences” on page 760](#)). After a device reaches this limit, it begins purging the database.

Example of Using Alerts

For example, a network host performs the normal connections required for Internet connectivity (SMTP, POP3, HTTP, and so on). The host becomes infected by a worm and begins making outbound connections on an arbitrary port. The device logs the unique event and generates PROFILER_NEW_PROTO and PROFILER_NEW_PORT log records. The system immediately e-mails these log records to the Security Administrator, who can investigate the worm and take action to contain it.

Repeat the configuration process for each device in your network. When you have configured all devices on your network, you are ready to start the Profiler.

Updating Profiler Settings

After you configure settings on the Profiler, you must update those settings on the device.

To update the settings on the device:

1. From the **Device Manager**, right-click on the device and select **Update Device**.

The **Device Update Options** window prompts you to **Restart IDP Profiler After Device Update**.

2. Click **OK**.

The **Job Information** window shows the status of the update. After the operation finishes, the device begins collecting data for the Profiler DB.

Starting the Profiler

To manually start the Profiler, use the **Devices** menu, and select **IDP Profiler > Start Profiler**. In the **Start Profiler** dialog box, select the devices you want to use for profiling, then click **OK**, or optionally, right-click on any device from the Device Manager and select **IDP Profiler > Start Profiler**.



NOTE: After you start the Profiler for a specific device, the Enable Protocol Profiler setting in the device is automatically enabled.



NOTE: The Profiler is actually a service, located in `/usr/idp/device/bin/profiler.sh`.

As your devices begin profiling your internal network, they gather information about your network hosts, their peers, ports, and Layer 7 data.

Stopping the Profiler

To manually stop the Profiler, use the **Devices** menu, and select **IDP Profiler > Stop Profiler**. In the **Stop Profiler** dialog box, select the appropriate devices, then click **OK**, or optionally, right-click on any device from the Device Manager, and select **IDP Profiler > Stop Profiler**.



NOTE: After you stop the Profiler for a specific device, the Enable Protocol Profiler setting in the device is automatically disabled.

Starting Profiler Operations on ISG Devices Without IDP Rules

When a Start Profiler operation is performed on a regular or virtual system (vsys) device that has no IDP rules, profiling for the device is disabled, and NSM returns the following message:

```
IDP Profiling for <device-name> will be disabled because there are no IDP rules
currently associated with this device.
```

If profiling for the device was enabled before the Start Profiler operation, profiling is disabled for the vsys device; however, the profiler process itself might not be disabled because other virtual systems associated with the device might be using the profiler process.

Customizing Profiler Preferences

To configure the following Profiler preferences, use the **Tools** menu, and select **Preferences > Profiler Settings**:

- Profiler:
 - Purge Profiler Database If Size Exceeds —NSM purges the profiler database size if it exceeds 1 GB (1000 MB) by default.
 - Max Profiler Database Size After Purging—If the database size exceeds its maximum limit, NSM purges the profiler database size until the size reaches 750 MB by default.
 - Profiler Query Timeout (120 seconds or 2 minutes by default).
 - Hour Of Day To Perform Database Optimization (midnight GMT by default).
- AVT View Settings:
 - Number Of Sessions To Display Per Application—The range is 5–10,000 sessions; the default is 10 sessions.
 - Hours Of Session Data To Display from Present Time—Configure from 1–24 hours; the default is one hour.

About Profiler Views

The Profiler includes four main views that you can use to analyze data about your profiled network:

- Protocol Profiler—Displays a snapshot of Layer-7 traffic on your internal network including source, destination, service context, and value. Use this view to analyze specific applications that are running on your network, their versions, and the values for each supported context.
- Network Profiler—Displays a high-level snapshot of static information (Layer-3, Layer-4, and RPC protocols, ports, and program numbers) on your internal network along with the Source/Destination IP, and Source/Destination MAC and Organizationally Unique Identifier (OUI). Use this view to quickly see which hosts are communicating with other hosts, and what services are passing between them.
- Violation Viewer—Similar to the Network Profiler, the Violation Viewer displays a high-level snapshot of network traffic— Layer-3, Layer-4, and RPC protocols, ports, and program numbers along with the corresponding Source/Destination IP, and Source/Destination MAC and OUI. The Violation Viewer, however, enables you to more effectively view content that does not match or is in violation of certain patterns that you can set in a shared object called a permitted object. You must configure permitted objects before data appears in this view.
- Application Profiler—Displays the network traffic information at the application and application-group level. Applications that are running in a network are grouped by the common functionality they provide to the network user. These applications form a hierarchal structure called an application hierarchy. For example, Yahoo messenger,

MSN, and AIM are chat applications; Kazaa, Bittorrent, and Gnutella are file sharing applications. In the application hierarchy, you view both chat and file-sharing applications are grouped under peer-to-peer applications.

About the Protocol Profiler

The Protocol Profiler view is a table of information that, like the Log Viewer, enables you to view and analyze dynamic application (Layer-7) traffic within a specific context. By default, this view contains only the data collected during the configured time interval; additionally, if you select a specific device, the Profiler displays only the information gathered by that device.

Table 87 on page 761 lists and describes the information that you can view using the Protocol Profiler:

Table 87: Protocol Profiler Data

Column	Description
Src IP	Source IP address of the traffic profiled.
Dst IP	Destination IP address of the traffic profiled.
User	The user associated with the traffic profiled.
Role	The role group to which the user that is associated with the traffic profiled belongs.
Context	All contexts of traffic that the devices selected in the Device table recorded.
Value	When you select a context, the values that your devices recorded for a selected context.
Src MAC	Source MAC addresses of traffic profiled.
Dst MAC	Destination MAC addresses of traffic profiled.
Src OUI	Source OUIs of traffic profiled. NOTE: OUI stands for Organizationally Unique Identifier. This value is a mapping of the first three bytes of the MAC address and the organization that owns the block of MACs. You can obtain a list of OUIs at http://standards.ieee.org/regauth/oui/oui.txt .
Dst OUI	Destination OUIs of traffic profiled.
Src OS Name	Operating-system version running on the source IP of the traffic profiled.
Dst OS Name	Operating-system version running on the destination IP of the traffic profiled.
Hits	Number of occurrences that match the traffic profiled.

Table 87: Protocol Profiler Data (continued)

Column	Description
First Time	Timestamp for the first time the device logged the event (within the specified time interval).
Last Time	Timestamp for the last time the device logged the event (within the specified time interval).
Domain	Domain in which the device is managed in NSM.
Device	Device that profiled the data displayed.

About the Network Profiler

The Network Profiler view is a table of information that, like the Log Viewer, enables you to view and analyze data related to static traffic (Layer-3, Layer-4, and RPC protocols, ports, and program numbers) within the context of data corresponding to peer, host, and operating system.

[Table 88 on page 762](#) lists and describes the information that you can view from the Network Profiler:

Table 88: Network Profiler Data

Column	Description
Src IP	Source IP address of the traffic profiled.
Dst IP	Destination IP address of the traffic profiled.
User	The user associated with the traffic profiled.
Role	The role group to which the user that is associated with the traffic profiled belongs.
Service	All services of traffic profiled.
Access Type	Type of the traffic profiled: Access indicates a successful connection, during which the device recorded valid requests and responses from the server to a client. Attempt indicates a request that did not receive a reply. The device recorded a packet from a client to a server, but never saw a reply. Probe indicates a request that does not expect a reply. For non-TCP sessions, the device recorded an ICMP error; for TCP sessions, the device recorded a SYN packet from the client followed by a RST from the server.
Src MAC	Source MAC addresses of traffic profiled.
Dst MAC	Destination MAC addresses of traffic profiled.

Table 88: Network Profiler Data (continued)

Column	Description
Src OUI	Source OUIs of traffic profiled. NOTE: OUI stands for Organizationally Unique Identifier. This value is a mapping of the first three bytes of the MAC address and the organization that owns the block of MACs. You can obtain a list of OUIs at http://standards.ieee.org/regauth/oui/oui.txt .
Dst OUI	Destination OUIs of traffic profiled.
Src OS Name	Operating-system version running on the source IP of the traffic profiled.
Dst OS Name	Operating-system version running on the destination IP of the traffic profiled.
Hits	Number of occurrences that match the traffic profiled.
First Time	Timestamp for the first time the device logged the event (within the specified time interval).
Last Time	Timestamp for the last time the device logged the event (within the specified time interval).
Domain	Domain in which the device is managed in NSM.
Device	Device that profiled the data displayed.

About the Violation Viewer

The Violation Viewer is similar to the Network Profiler view. The Violation Viewer displays the same data that you can view in the Network Profiler view, but only for those object entries that do not match specific address and service criteria. By creating specific permitted objects, you can configure the Violation Viewer to display only those items that violate the criteria that you set.

Configuring Permitted Objects

Permitted objects are shared objects specific to the Profiler. They enable you to configure objects in the Profiler containing simple rules, consisting only of source IP, destination IP, and service. The implied action is “permit”. You can then use the object to define what you should see on the network—as opposed to an attack object, which defines what you do not want to see as a set of rules. After you have created your permitted objects, the Violation Viewer displays all traffic that does not match the criteria that you have configured in these objects.



NOTE: In previous versions of IDP Manager, permitted objects were called *violation objects*.

Example: Using the Violation Viewer to Detect Traffic That Uses Nonstandard Ports

The Profiler can help you confirm suspected violations such as SQL servers running on the network. Most of the time, however, you do not know exactly what you are looking for on the network. In these cases, it is easier to specify exactly what should be on the network, then detect any traffic that violates that specification.

For example, you want to detect internal traffic that uses a nonstandard port for its service.

1. From the Violation Viewer, click on the **+** icon that appears on the top of the right-hand window. You can also right-click anywhere in the right-hand window and select **Add**. A New Permitted Object window appears.
2. Name the object "Non-Standard-Ports".
3. Right-click on the Service column and select **Add Service**.
4. Select all predefined services.
5. Click **OK** to save the permitted object. After you have created and saved the permitted object, the object automatically becomes available in the Profiler.
6. From the Violation Viewer, select the new permitted object "Non-Standard-Ports". The Profiler uses the object to filter the data collected from the devices. Traffic that matches the object (uses a standard service port) is filtered out, leaving only the traffic that does not match (uses a nonstandard service port).

You can now review the data in the Violation Viewer to see all traffic on your network that uses non standard service ports.

Now that you can see the traffic you do not want on your network, take the appropriate security measures, for example, remove the unauthorized network components, incorporate the components/services into your existing corporate security policy, or create rules in your security policy to restrict the traffic to specific network components.

Example: Detecting Traffic That Is Not Using Primary Services

For example, you want to detect internal traffic that is not using the primary services for Web access, e-mail, ping, and DNS services. In the Object editor, create a permitted object to permit all traffic that uses a standard service port. For Service, select the following predefined services:

- dns
- http
- https
- ping
- pop3
- smtp
- ssh

After you have created and saved the permitted object, the object automatically becomes available in the Profiler.

1. From the Violation Viewer, select the permitted object **Internal Services**. The Profiler uses the permitted object to filter the data collected from the devices. Traffic that matches the object (uses a service specified in the object) is filtered out, leaving only the traffic that does not match (does not use a service specified in the object).
2. Take appropriate measures to secure the network, for example:
 - Investigate the source IP of the traffic and contact the user. If the traffic is legitimate, you might need to add the service to your corporate security policy to allow it on your network. You should also edit the permitted object to include the service so you no longer see the service in the Violation view.
 - Create a rule in your security policy that drops connections between your internal network objects if the traffic uses a service that you do not allow on your network.

About the Application Profiler

The Application Profiler view is a table of information that allows you to view network traffic information at the application and application group level. By default, this view contains only the data collected during the configured time interval. You can also specify filters to display only the data you want to view.

As the applications that are running on the network form a hierarchy, the volume of traffic that is generated by individual applications or application groups can be added to provide the aggregate traffic volume information from the parent application group. As you move up the root of the application hierarchy, you can view the total network traffic volume.

The Application Profiler view is divided into two sections: on the left panel, the hierarchical application view displays a tree of the application categories with volume information (in bytes and packets); on the right panel the application session view displays the application name and the aggregated bytes and packets for the application. You can expand the application name to see latest session data for that application.

[Table 89 on page 765](#) lists and describes the information that you can view using the Protocol Profiler:

Table 89: Application Profiler Data

Column	Description
Src IP	Source IP address of the traffic profiled.
Dst IP	Destination IP address of the traffic profiled.
VLAN ID	VLAN ID associated with the traffic profiled.
Application	Application ID associated with the traffic profiled.
Byte Count	Byte count for the traffic profiled.

Table 89: Application Profiler Data (continued)

Column	Description
Packet Count	Packet count for the traffic profiled
User	The user login name.
Role	The role group to which the user that is associated with the traffic profiled belongs.
First Time	Timestamp for start of session.
Last Time	Timestamp for end of session or the last update time.
Domain	Domain in which the device is managed in NSM.
Device	Device that profiled the data displayed.

Using Profiler Views

You can perform all of the following functions within the Profiler views, as described in the following sections:

- [Filtering and Sorting from the Protocol Profiler, Network Profiler, and Violation Viewer on page 766](#)
- [Filtering and Sorting from the Application Profiler on page 767](#)
- [Refreshing Profiler Data on page 768](#)
- [Viewing Database Information on page 769](#)
- [Viewing Detailed Network Information on page 769](#)
- [Purging the Database on page 770](#)

Filtering and Sorting from the Protocol Profiler, Network Profiler, and Violation Viewer

To help you view and analyze data in the Profiler, you can use criteria to filter information for each cell or column that appears in the Profiler view.

Double-click in any of the columns that appear in the Filter Criteria. A dialog box lets you add entries that match the column you selected as a criterion to filter the Profiler view. The Profiler view automatically updates, displaying the data that matches the criterion that you have set in the filter. You can also right-click on any entry in the Profiler view and select **Add to Filter** to add that entry as a filter criterion. Similarly, you can use the **Remove From Filter** option to remove that entry as a filter criterion.

Click on the **Negate** option to hide entries that match the criteria that you have set as a filter. You can also right-click on any entry in the Profiler view and select **Toggle Filter Negation** to hide entries that match that criterion.

Right-click on any filter criteria on any entry in the Profiler view and select **Clear Column Filter** to disable filtering on entries that match the criteria set.

Right-click on any filter criteria or on any entry in the Profiler view and select **Clear All Column Filters** to disable all filtering.

Other options that you can set in the Profiler views include:

- **Sorting**—Click the column header to sort columns in ascending or descending order.
- **Rearranging Columns**—Drag a column header to a new position in the table to rearrange the order of the columns in the viewer.

All filter criteria are saved each time you log out of the UI.

Example: Setting a Time Interval

Profiler queries that return a very large amount of data may not complete in an acceptable amount of time. By default, queries that take longer than 120 seconds (2 minutes) will time out.

You can reconfigure the timeout using the Profiler settings under the **Tools** menu.

If the Profiler times out before your query is finished, refine your filter criteria. For example, if you want to analyze a specific event in your internal network, set a time interval for the data.

The Profiler records the first-seen and last-seen timestamps for each entry in the database. You can set a time interval based on these timestamps. You can set an exact From and To time, or enter the Last day, hour, minute, or seconds:

- Use the **First Seen** setting to define a start timestamp threshold. If the device logged an event for the first time, and the event timestamp is after the start timestamp, the event appears in the Profiler view. For example, to see all new events in the last 2 days, configure the First Seen timestamp as the last 2 days.
- Use the **Last Seen** setting to define a last timestamp threshold. If the device logged an event, and the event timestamp is before the last timestamp, the event appears in the Profiler view. For example, to see what network components have been idle over the last 10 days, configure the Last Seen timestamp as the last 10 days.

After you have configured a time interval, the selected Profiler view automatically applies the time interval as a filter criterion.

Filtering and Sorting from the Application Profiler

To help you view and analyze data in the Application Profiler, you can use criteria to filter information.

You can double-click in the Src IP and Dst IP address column fields that appear in the **Filter Criteria** to add entries that match the column you selected as a criterion to filter the Application Profiler view. You can also right-click on any entry in the Profiler view and select **Add to Filter** to add that entry as a filter criterion. Similarly, you can use the **Remove From Filter** option to remove that entry as a filter criterion.



NOTE: You cannot filter on the Application column of the Application Profiler.

Click on the **Negate** option to hide entries that match the criteria that you have set as a filter. You can also right-click on any entry in the Profiler view and select **Toggle Filter Negation** to hide entries that match that criterion.

Right-click on any filter criteria on any entry in the Profiler view and select **Clear Column Filter** to disable filtering on entries that match the criteria set.

Right-click on any filter criteria or on any entry in the Profiler view and select **Clear All Column Filters** to disable all filtering.



NOTE: The Application Profiler view is not automatically updated when you add, delete, or modify filters. To update the Application Profiler view with the latest application session data, select the application you want to refresh and then click the refresh icon, in the top right corner of the Profiler view.

You can perform any of the following actions on any Application header:

- Select All—Selects all the applications in the session view.
- Clear All—Clears all selections in the session view
- Expand All— Expands all applications and loads new data
- Collapse All— Collapses all expanded applications
- Collapse Selected—Expands selected applications and loads new data

Other options that you can set in the Application Profiler view include:

- Sorting— Sort on any column except the Application column. The Application column does not support sorting because application values are similar for each application group. When you perform a sort on any other column, only the rows inside the loaded tables are rearranged.
- Rearranging Columns—Drag a column header to a new position in the table to rearrange the order of the columns in the viewer.

Refreshing Profiler Data

NSM fetches data from each of the devices that you are profiling automatically. By default, data is fetched from three devices in sequence. When the Device Server finishes fetching data from one device, it begins the operation on the next available device.



NOTE: You can change the default behavior that the Profiler uses to fetch data by editing parameters in the Device Server configuration file. Refer to the *Network and Security Manager Installation Guide* for more information.

Click on the **Refresh** icon periodically to refresh the Profiler view with the latest data available.

Viewing Database Information

Click on the **Show DB Information** icon to view specific details about the Profiler database, including the database size.

Viewing Detailed Network Information

Click on the **Show Detail Viewer** to access the Detailed Network Information view. This view displays details for selected IP addresses, enabling you to further investigate the details of a connection, such as a username for an account on a host, open ports, or RPC services. You can also right-click on any entry in the Profiler view and select **Show Detail Viewer** to access the Detailed Network Information view for that entry.

Use the IP and MAC areas to select the IP and MAC addresses of each interface on the selected host. For hosts with multiple interfaces, select the MAC address for the interface you want to investigate.

Click on any of the tabs in the right-hand window to view different types of information that your devices have recorded for the selected IP address and MAC address.

Table 90 on page 769 lists and describes the information that you can view using the Detailed Network Information View:

Table 90: Detailed Network Information Data

Column	Description
Host	<p>Details about the selected host IP, including:</p> <ul style="list-style-type: none"> • IP Address • MAC Address • OUI (Organizationally unique identifier), a mapping of the first three bytes of the MAC address and the organization that owns the block of MACs (to obtain a list of OUIs, see http://standards.ieee.org/regauth/oui/oui.txt). • VLAN tag (if applicable) • First Seen date and time that the devices first recorded traffic on that interface • Last Seen date and time that the devices last recorded traffic on that interface <p>The OUI value enables you to see immediately the vendor of the network interface card (NIC) that is generating the packets.</p>
TCP Ports	Details about the outbound TCP ports on the selected host IP.
UDP Ports	Details about the outbound UDP ports on the selected host IP.
Layer-3 Protocols	Details about the outbound Layer-3 protocols (IP) on the selected host IP.
RPC Services	Details about the outbound RPC services on the selected host IP.

Table 90: Detailed Network Information Data (continued)

Column	Description
Profiles	<p>Details about the contexts and values on the selected host IP. Use the context and value fields to identify:</p> <ul style="list-style-type: none"> • Software version of the application • Username and password of an account on that host • Computer name
Peers	<p>Details about how the selected host IP communicates with other hosts, including:</p> <ul style="list-style-type: none"> • Source IP address of peer traffic • Destination IP of peer traffic • Hits (number of times the source IP communicated with the destination IP) • First Access date and time that IDP first recorded the peer traffic • Last Access date and time that IDP last recorded the peer traffic <p>If the source IP is the selected interface, you can use the destination IP address to identify all the hosts that the interface sent traffic to, both internal and external, as well as the number of times traffic was sent.</p>
Summary	Lists all the details of the other View box selections in HTML format.

Purging the Database

When the Profiler DB reaches a maximum size (4 GB by default), it begins purging records (oldest first) automatically. The Profiler DB stops purging records when it reaches a certain set minimum size (3 GB by default). You can use the Profiler settings in the Tools preferences menu to change these parameters.

To manually purge the Profiler DB of all records, click **Clear All DB**. This operation can take up to one minute. During this time, a message appears on all other connected UIs indicating that a Clear All DB operation is in progress. After the data clears, you can click on the **Refresh** icon to get new data.



NOTE: The **Clear All DB** option is not supported for Application Profiler views. Application Profiler data is purged every 1–24 hours, depending on the value configured for the AVT View Settings from **Preferences > Profiler Settings** in the **Tools** menu.



NOTE: NSM supports purging all records for the entire Profiler DB (all subdomains) but does not support purging records for specific subdomains. Consequently, an IDP administrator cannot delete Profiler DB records for a subdomain. Only the NSM super administrator can delete records for the Profiler DB, and all records for all the subdomains are purged.

Recommended Profiler Options

The following are recommended for using the Profiler:

- [Configuring a Network Baseline on page 771](#)
- [Keeping Your Network Current on page 772](#)
- [Proactively Updating Your Network on page 772](#)
- [Reacting to Vulnerability Announcements on page 772](#)
- [Stopping Worms and Trojans on page 773](#)

Configuring a Network Baseline

A baseline is a static view of your network traffic patterns. This view, which is compiled from multiple views of traffic over time, represents the normal, known activity that occurs on your network. By setting a baseline for your network, you can quickly detect any traffic that deviates from or violates that baseline and take appropriate measures.

Identifying a Baseline

Your devices begin learning your network baseline as soon as the Profiler starts. As your devices profile the network for the first time, each component appears as new. To avoid unnecessary log records generated by Profiler alerts, you should ensure that alerts are not enabled in the **Alerts** tab of the **Profiler Settings** dialog box.

During the learning phase, your devices profile the network hosts, servers, and software applications that they protect; the Profiler synchronizes profile information from the devices and creates an initial view of your network. Each time you synchronize the devices, the Profiler incorporates any new data into this view, creating a more complete, up-to-date image of your network. You should continue to synchronize data daily until you feel the Profiler is accurately depicting your normal traffic patterns.

Because all networks are different, the learning phase can range from a few hours to a few weeks.

Setting a Baseline

When you are satisfied that the Profiler has detected each host, protocol, and port that you want to profile, you have successfully created a network baseline. By itself, this baseline view can help you implement software and hardware upgrades, take inventory for new support contracts, plan for a network ROI investigation, and so on.

However, the true power of your network baseline is to enable your devices to identify network deviations. The Profiler uses the baseline to identify new or unknown hosts or software that might represent a network vulnerability. Network deviations can be a simple application update, or a serious security breach.

When enabled, if the device discovers a new host, protocol, or port, the device generates a log record, such as PROFILER_NEW_HOST, in the subcategory column of the Log Viewer.

Keeping Your Network Current

Typical networks include multiple servers and hosts, each running different operating systems and software applications that are important to users on the network. While this variety helps users accomplish their tasks, it can make it difficult to keep your network systems current. As new versions or security updates are announced, you must first determine if your network is affected, locate the affected components, then patch as appropriate.

To help you maintain control of your network software versions, the Profiler uses passive application fingerprinting to identify the application version for each service used in your network. Additionally, your devices collect and store the user name and other important information for each application. These profiling activities provide you with an inventory of operating-system and software applications, versions, and the components that use them.

You can use this information to proactively update your network or respond quickly to vulnerability announcements.

Proactively Updating Your Network

To eliminate security holes, you should update your software applications regularly.

Some guidelines:

- **Research known vulnerabilities.** Compare the information in the Profiler with a software vulnerability database, such as Security Focus at <http://www.securityfocus.com/bid> or Common Vulnerabilities and Exposures (CVE) at <http://www.cve.mitre.org/cve/>.
- **Plan to patch.** After you identify your vulnerable systems, schedule a regular maintenance time to keep downtime and disruption to a minimum.

Even if your network components do not require security patches or updates, they might use default configurations. Many network device vendors use a common phrase, the vendor name, or other simple word as the default password for accessing the administration interface of their device. Because these passwords can be guessed easily, the vendor recommends that users change the default password immediately. However, for convenience, some users leave the default configuration password, unknowingly opening a security hole in the network. The Profiler captures user information that you can use to see who is logging in to network devices so you can verify that they are from trusted IP addresses.

Reacting to Vulnerability Announcements

New network attacks and exploits are discovered every day. When new security patches are issued, use the Profiler to quickly identify which systems are running the affected software version, then patch them appropriately.

For large networks, it is difficult to patch everything immediately. Plan your patching process by prioritizing based on the importance of the resources. Critical, high-risk, and heavily used resources should be patched first, while less important, minimally used resources might be able to wait.

Example: Identifying Vulnerable Components

For example, Microsoft announces a vulnerability in version 6.0 of the Microsoft Internet Informations Services (IIS). To quickly identify all network components running the vulnerable version:

1. Select the **Protocol Profiler** to see the applications running on the network.
2. In the Context Filter data table, select **HTTP Header Servers**. The value data table lists all Web servers currently running. The network uses the following Web servers:
 - Apache (two versions)
 - Microsoft IIS, version 6.0
3. Select the **Microsoft IIS 6.0** value to find out which IP addresses are running the IIS server. The Protocol Profiler displays the destination IP address of the service, which is the IIS server.
4. Patch the vulnerable IIS server by using the information supplied with the Microsoft Security Bulletin.

Stopping Worms and Trojans

Worms and trojans often bypass firewalls and other traditional security measures to enter a network. Because worms and trojans operate inside a network, external firewalls might not be able to detect them.

Use the Profiler to determine when a worm or trojan entered your network, how it was introduced, and which network components were infected. By filtering the profile data you can identify the source and contain the attack to minimize impact, before investigating and recovering from any damage.

Example: SQL Worm

For example, your corporate security policy does not permit SQL servers on the internal network. However, during a regular Microsoft update, SQL applications are installed on a network server, without your knowledge. Because you are not aware that an SQL server is running on your network, you do not attempt to block SQL attacks at your firewall or IDP system. Suddenly, the SQL Slammer worm attacks and infects your network.

Using the Profiler:

1. Create a custom TCP service object to represent Microsoft SQL (default port: TCP/1433).
2. Restart the Profiler.
3. Select the **Network Profiler** to quickly see the source, destination, and service of traffic on your network.
4. In the Service data table, select the SQL service you just created. The Network data viewer lists all network components current running SQL servers.
5. Take appropriate measures to secure the network, such as:

- Apply patches.
- Remove the components from your network.
- Remove SQL from all components.
- Create a rule in your security policy that drops all SQL connections between your internal network objects.

Example: Blaster Worm

For example, the Blaster worm uses a special ICMP (ping) packet to exploit a vulnerability in Remote Procedure Call (RPC), a Microsoft networking tool that enables desktops to share files over a remote network. Your corporate firewall denies RPC filesharing traffic to protect sensitive corporate files from Internet users, but enables RPC filesharing on a local network for convenience.

A laptop user uses a wireless network to access the Internet. Because the laptop is configured to allow RPC, it contracts the Blaster worm from an infected user on that network. When the user returns to the office and connects the laptop to the corporate network, the worm immediately begins scanning the internal network and infecting all components that have RPC enabled.

Because the Profiler records all unique activity on the network, it identifies the ICMP packet scans as a new event. Because you have configured the Profiler to send alerts for new hosts, you also receive a log record on your pager indicating that a new host has joined the network. A quick check of the Profiler's Network view tells you that the new event is a user laptop suddenly scanning the entire network using ICMP, a possible sign of the Blaster worm.

From the Profiler:

1. Restart the Profiler.
2. Select the **Network Profiler** to quickly see the source, destination, and service of traffic on your network.
3. In the Service data table, select the **ICMP service**. The Network data viewer displays all network components using ICMP.
4. In the Access data table, select **probe**. The Network data viewer displays all network components that used ICMP to probe the network.
5. Set a Last Seen time interval of two hours.

The Network Profiler displays all network components that used ICMP to probe the network in the last two hours. You can now see that one IP address, **192.168.4.66**, is currently probing your network using ICMP. However, because you use DHCP to dynamically assign IP addresses, you need to identify which user laptop is currently using that IP address.

6. From Network Profiler, select the source address you want to investigate. The MAC/View area displays the host detail for the IP address.
7. In the **View** menu, select **Profiles**. The MAC/View area displays the context/value information about the IP/Mac address.

The IP/MAC address has the unique asset tag "darkness". After checking your IT inventory, you determine who the laptop user is and patch the infected system.

Accessing Data in the Profiler Database

The Profiler database is located on the NSM Device Server.

To query the actual records in the database:

1. Log in to the Device Server.
2. Execute `/usr/bin/psql -U nsm -d profilerDb`; where **nsm** stands for the username and **profilerDb** is the database name. By default, the PostgreSQL user is set to **nsm**. You can define the PostgreSQL user when installing NSM.

The following is a sample query:

```
[root@ bin]# which psql
usr/bin/psql
root@]# /usr/bin/psql -U nsm -d profilerDb
Welcome to psql 8.1.9, the PostgreSQL interactive terminal.

Type: \copyright for distribution terms
      \h for help with SQL commands
      \? for help with psql commands
      \g or terminate with semicolon to execute query
      \q to quit

profilerDb=# select * from host;
 id | device | os | ip | mac | oui | hits | vlan | first | last
----+-----+---+---+---+---+-----+-----+-----+-----
(0 rows)
profilerDb=#
```

About Security Explorer

The Security Explorer is a powerful, graphical tool that enables you to visualize and correlate network behavior based on data collected in the Profiler, Log Viewer, and Report Manager. You can use the Security Explorer to perform the following tasks:

- Get a dynamic, interactive view of your network.
- Drill down on a particular host or server and view all the different attacks, open ports, destination or peer IP addresses, and so on.

- Move between hosts and peers and trace a connection or attack.
- Toggle between different views or slices of the network, as well as explore the contextual information (logs, reports, IDP attacks, IP addresses, and so on) within the Security Explorer panel.

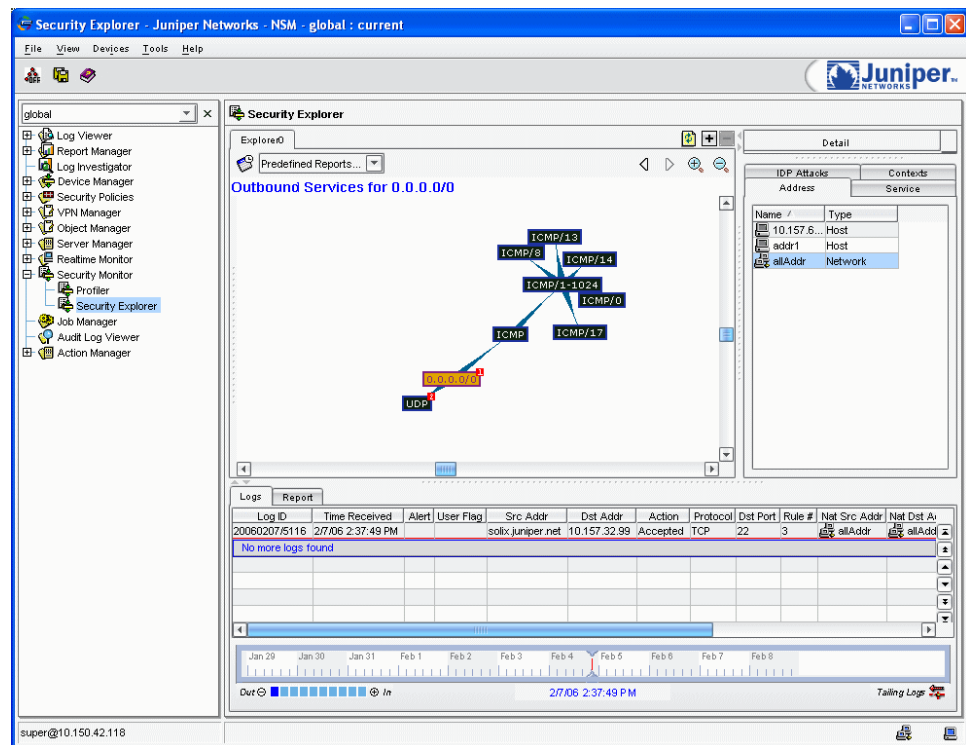
The main component is a graph that represents the relationships between data objects, such as hosts, services, attacks, etc.

There are five main views in the Security Explorer:

- “Security Explorer Main Graph” on page 776
- “Connections Detail Pane” on page 778
- “Reference Point Pane” on page 778
- “Log Viewer” on page 778
- “Reports Viewer” on page 778

Figure 112 on page 776 shows the Security Explorer.

Figure 112: Security Explorer



Security Explorer Main Graph

The main component of Security Explorer is a graph that displays the following nodes:

- Host—Displayed as an IP address
- Network—Displayed using CIDR notation (ip/class: 8/16/24)
- Protocol—These include TCP, ICMP, and so on
- Attack—Specific attack object name
- Service—Displayed in protocol/port notation
- Service range—Displayed in protocol/port range notation, for example, TCP/1-1024
- Context—Plain text describing one protocol attribute, for example, 'SSL server version'
- Value—Value specific for a context, for example, for 'SSL server version' the value '3.1'

Graph Types

Depending on the starting reference point, the following graphs appear in Security Explorer depicting relationships between objects (for example, peers of a host or services for a host):

- Peer IP—Selected host and all peers of this host. Hosts are grouped in networks /8, /16, /24. Every network appears as a graph node and is connected to its subnetworks and hosts.
- Outbound Services—A host or network and its outbound services. The services are grouped by protocol (TCP, ICMP, and so on) and by service port range. Every group is displayed as a separate node, for example, for a host 10.150.151.3 with outbound service FTP, the graph structure is: [Host:10.150.151.3]- [Protocol:TCP]- [Service range:TCP/0-1024]-[Service:TCP/21]
- Inbound Services—A host or network and its inbound services. The services are grouped by protocol (TCP, ICMP, and so on.) and by service port range. Every group is displayed as a separate node, for example, for a host 10.150.151.3 with outbound service FTP, the graph structure is: [Host:10.150.151.3]- [Protocol:TCP]- [Service range:TCP/0-1024]-[Service:TCP/21]
- Server Profiles—One host or network and the context for server-related traffic. Every context is connected to its host/network related value, for example, on a host is an SSL server running version 3.1. The graph displays the host and its relationship to the host connected 'SSL Server Version' context and to the context connected value '3.1'.
- Client Profiles—One host or network and context for the client related traffic. Every context is connected to its host/network related value, for example, on a host is an SSL server running version 3.1. The graph displays the host and its relationship to the host connected 'SSL Server Version' context and to the context connected value '3.1'.
- Outbound IP—All outbound hosts/networks for a selected service, context, or value.
- Inbound IP—All inbound hosts/networks for a selected service, context, or value.
- Value—All values for a selected context.

Connections Detail Pane

Use the Connection Details pane to view a list of all objects connected to the currently selected object in the graph. For example, if you are viewing an Outbound Services graph, and a host is selected, the Connections Detail pane contains all services for this host. If a Peer IP graph appears, the Connections Detail pane contains all peers for the selected object.

Double-clicking on one of the objects in the **Details** pane displays the relationship graph for it.

Reference Point Pane

Use the **Reference Point** pane to view a list of reference objects. Next to the graph, a list of the possible initial objects appears:

- Address objects—For host/network selection. The selected host appears in a Peer Graph, with selected network in a Outbound Service graph.
- Service objects—Displayed in an Outbound IP graph.
- IDP Attacks—Displayed in an Attack-Network graph.
- Context objects—Displayed in a Value graph.

Double-click on an object to set it as a point of reference in the main graph.

Log Viewer

Use the **Logs** tab in the viewer that appears below the main graph to view all logs related to the point of reference that you have selected.

Reports Viewer

Use the **Reports** tab to generate and view one of the following reports in Security Explorer:

- Top Alarms
- Top Traffic Alarms
- Top Traffic Logs
- Top IDP/DI Attacks
- Top Screen Attacks
- Top Destination IPs
- Attacks over time
- Attacks by severity
- Attacks by flag
- Critical severity attacks
- High severity attacks

Use the **Source/Destination Address** pull-down menu to view these reports based on either the Source IP or Destination IP.

Use the **time duration** pull-down menu to view data in these reports during a specific time frame. You can select to view data from the last 24, 12, 8, 4, 2, 1 hours.

Using Security Explorer

You can launch the Security Explorer in any of the following ways:

- From the Security Monitor tree node, select Security Explorer.
- From the Log Viewer, for any anomaly, signature or custom log, right-click on a **Source/Destination Address**, **Protocol/Destination Port**, or **NAT Source/Destination Address** and select **Launch Security Explorer**.



NOTE: For all other log categories, the **Launch Security Explorer** option is not available.

- From the Profiler, right-click on a **Source/Destination Address**, **Context**, **Service**, **Subcategory**, or **Time Received** in any entry and select **Launch Security Explorer**.

If you launch Security Explorer from the Security Monitor without any reference point, the main graph appears empty.

If you launch Security Explorer using a starting point of reference from the Log Viewer or the Profiler, that reference point is depicted in the main graph.

Permissions

To use the Security Explorer, you must have the proper administrative privileges including the View Security Explorer activity. By default, all IDP administrators have this privilege. Depending on other activities you may want to use with Security Explorer, you also may need proper administrative privileges to:

- View Profiler
- View Device Logs
- View Historical Log Reports
- View Devices
- View Shared Objects

If you do not have proper permission to perform all activities, you may only have access to a reduced set of Security Explorer features. For example, if you do not have View Profiler privileges, you cannot view graphs related to the Profiler, for example, Peer graph.

If an administrator is logged into the global domain, all devices in the address point of reference list appear. This includes all global and subdomain objects. If an administrator is logged into a subdomain, only those devices in the current subdomain devices appear.

Analyzing Relationships

After you have established your main point of reference, a list of navigation options appears. Every option represents a transition from one graph to another.

Viewing Data

The following view options are available, making it easier for you to view and analyze each node in the main graph:

- **Center Node**—Click on any node that appears in the main graph to center it. You can also right-click on the node and select **Center Node**
- **Expand Node**—To view all objects related to a specific node, right-click on any node and select **Expand Node**
- **Collapse Node**—If you are not interested in viewing all objects related to a specific node, you can reduce extraneous data by right-clicking on any node and selecting **Collapse Node**
- **Hide Node**—If you are not interested in viewing a specific node, and you want to remove it from viewing in the graph, right-click on any node and select **Hide Node**
- Use the zoom in and zoom out icons to increase or decrease the size of the nodes on the graph.
- Use the back arrow to view the previous graph. Use the forward arrow to view the next graph.

Transitioning to Other Relational Graphs

Use the icons that appear in the main graph to quickly access additional information related to your point of reference. Depending upon the type of icon that you select, you can transition to another graph. [Table 91 on page 780](#) describes the graphs that you can transition to:

Table 91: Transitional Graphs

Node	Transitions to:
Host	Peer IP, Outbound Services, Inbound Service, Server Profiles, Client Profiles, Attacks
Network	Outbound Services, Inbound Service, Server Profiles, Client Profiles, Attacks
Service	Outbound IP, Inbound IP
Service range	Outbound IP, Inbound IP
Context	Outbound IP (if selected context is related to a client), Inbound IP (if selected context is related to a server, for example, SSH Server Version), Values
Value	Outbound IP, Inbound IP

Table 91: Transitional Graphs (continued)

Node	Transitions to:
Attack	Source IP, Destination IP, Protocol Ports

Setting a Time Duration

Click on the Time Period icon to set a specific time period during which you want to view data.

Viewing Predefined Reports

Use the **Predefined Reports** pull-down menu to view a predefined report of that data. You can access three predefined reports:

- Top Attacks—Most common attacks on the network
- Top Attackers—Most common origins of attacks on the network
- Top Targets—Top destination targets of attacks on the network.

Refreshing Data

Click on the **Refresh** icon to update the Security Explorer with the latest data available.

Adding and Removing Panels

You can also view additional data and graphs by adding and removing additional panels to Security Explorer.

- Use the **+** icon to add a Security Explorer panel. The new panel appears as a new tab in the main graph area. Click on the tab to access the new panel.
- Use the **–** icon to remove the current Security Explorer panel.

Exporting to HTML

You can export any data depicted in the Security Explorer to an HTML file by using the **Export to HTML** option.

CHAPTER 19

Logging

Network and Security Manager integrates log information from multiple devices to help you access and distill data about the traffic on your network.

The Log Viewer presents log data as a log entry in a table; a log entry contains the details of the traffic that triggered the log, such as IP address, port number, and source and destination zones. This log data is also automatically used to generate predefined reports, helping you to interpret event information in a specific context. To perform your own investigation, use the Log Investigator to view cross-tabulations between sources, destinations, subcategories, and destination ports.

The Audit Log Viewer presents log entries triggered by administrative changes (changes made to the NSM system). An audit log entry includes details about the administrative event, such as the administrator name, timestamp of the change, and job details.

You can configure each managed device to generate and export specific log records to multiple formats and locations, such as syslog, xml, or e-mail servers. You can also forward logs that meet specified criteria to predefined formats and locations.

This chapter contains the following sections:

- [About Logging on page 783](#)
- [Viewing Logs on page 786](#)
- [Configuring the Device for Logging on page 787](#)
- [Using the Log Viewer on page 800](#)
- [Using the Log Investigator on page 821](#)
- [Using the Audit Log Viewer on page 832](#)
- [Managing Log Volume on page 837](#)
- [Forwarding Logs on page 840](#)

About Logging

Logging is the act of recording information about an event. In NSM, each event that occurs on your network or in your management system can be recorded and stored as a log entry. To view log entries from the NSM UI, you can use one or more of the logging-related UI components, such as the Log Viewer or the Log Investigator.

It may be helpful to visualize log entries being sent or *pushed* from the device to the NSM Device Server, which then pushes the log entries to the logging database. A UI module (the Log Viewer or **Report Manager**) requests or pulls the log entries in the logging database and displays the entries in the UI.

About Log Entries

A managed device generates a log entry when an event matches the configured logging conditions. The log entry, which contains details of the event, is sent to the NSM Device Server and stored in the logging database. You can view log entries in the NSM UI.

In a single log entry, you can view detailed information about where traffic comes from (the source address), where traffic goes (the destination address), and a description of the event that triggered the log entry. You can also view summarized information about events and alarms for multiple log entries. This data can help you analyze log entries and determine the effectiveness of your current security policies and device configurations.

About Log Events

Managed devices generate log entries based on events. Typically, devices generate log entries when:

- An event matches a rule in which logging is enabled. When you configure a rule for logging, the device creates a log entry for each event that matches that rule.
- An event matches a predefined set of conditions configured on a managed device or the management system.

Some events generate log entries that appear in the Log Viewer, while others appear in the Realtime Monitor. [Table 92 on page 784](#) details event-generated log entries.

Table 92: Event-Generated Log Entries

Events	Description	Destination
Attack, Alarm, Other	Generates log entries for events related to network activity on the device that violates a set threshold.	Log Viewer
VPN Events	Generates log entries for events related to VPN tunnels. These log entries are used to produce statistical information for monitoring.	Realtime Monitor >VPN Monitor
Configuration, Information, Self, Policy, Traffic	Generates log entries for events related to device configuration, NSM configuration, security policy rules, and traffic activity on the managed device.	Log Viewer
Flow, Ethernet, Attack, Policy	Generates log entries for events related to packet flow, Ethernet objects, network attacks, and security policy rules. These log entries are used to produce statistical information for monitoring.	Realtime Monitor >Device Monitor
Protocol Distribution	Generates log entries for events related to protocols used in network activity. These log entries are used to produce statistical information for monitoring.	Realtime Monitor >Device Monitor

About Log Severity

The log severity level defines the urgency of the information contained in the log entry. The severity level of a log entry depends on the log category, such as information, traffic, or configuration log entries.

You can configure a managed device to generate log entries only for those events that meet a specific severity level. Additionally, you can configure the device to forward log entries that contain a specific severity to a specific destination, such as a console location or syslog server. You can forward multiple log entries with different severity levels to the same log destination.

Juniper Networks assigns a predefined severity level in the firmware of each Juniper Networks device. However, these severity levels are not the same as the severity levels that appear in the log entries viewed in an NSM UI module.

[Table 93 on page 785](#) details how NSM handles severity levels as defined by DMI devices:

Table 93: Log Entry Severity Levels for DMI Devices

NSM Severity	Description
Alert	Log entries triggered when system requires immediate action.
Critical	Log entries triggered when system encounters critical conditions.
Emergency	Log entries triggered when system becomes unusable.
Error	Log entries triggered when system encounters errors, such as errors in device function.
Notice	Log entries triggered when system encounters normal but significant conditions.
Warning	Log entries triggered when system encounters warning conditions.
Informational	Log entries triggered by general system operations such as when a device connects or disconnects.
Debug	Log entries triggered when system receives debug-level messages.

[Table 94 on page 785](#) details how NSM handles severity levels as defined by ScreenOS and IDP devices:

Table 94: Log Entry Severity Levels for ScreenOS and IDP Devices

NSM Severity	Severity	Description
Device_critical_log	Emergency	Log entries triggered when traffic matches a critical severity attack object. Also includes log entries triggered by the SCREEN-level attacks, SYN attacks, Tear Drop attacks, and Ping of Death attacks.
	Alert	Log entries triggered by the general firewall SCREEN-level attacks or other conditions that require immediate attention, such as the expiration of license keys.

Table 94: Log Entry Severity Levels for ScreenOS and IDP Devices (continued)

NSM Severity	Severity	Description
Major	Critical	Log entries triggered when traffic matches a major severity attack object. Also includes log entries triggered by changes in the device function, such as high availability (HA) status changes.
Minor	Error	Log entries triggered when traffic matches a minor severity attack object. Also includes log entries triggered by errors in device function, such as a failure in antivirus scanning or in communicating with SSH servers.
Device_warning_log	Warning	Log entries triggered when traffic matches a warning severity attack object. Also includes log entries triggered by questionable device activity, such as a failure to connect to e-mail servers and authentication failures, timeouts, and successes.
Info	Notification	Log entries triggered when traffic matches an informational severity attack object. Also includes log entries triggered by normal events, such as device configuration changes.
	Information	Log entries triggered by general system operations such as when a device connects or disconnects.
Not Set	Other	No severity is set.



NOTE: From NSM release 2008.1 onwards, critical and warning logs from ScreenOS and IDP devices are displayed as `Device_critical_log` and `Device_warning_log`. If upgrading from an earlier release, you may need to modify your action manager criteria to match the new conventions.

Viewing Logs

NSM logging tools provide a high-level view of the activity on your network, enabling you to view summaries as well as detailed information. You can choose to view log entries for an event that occurs across domains (you must have the requisite permissions), as well as for specific device groups, clusters, firewalls, and so on.

Because you collect log entries from multiple devices, log analyzing, log volume, and log management are important concerns. To control the amount of log data displayed on screen, use tools such as filters, flags, and custom views to help identify patterns, and even isolate log entries from devices that appear to be the source of problems. For further investigation, use the Log Investigator tools to cross-tabulate source, destination, and attacks. Based on your analysis, you can then edit the rules in your security policies to modify how NSM handles your log entries.

NSM includes three primary logging modules:

- **Log Viewer**—Presents complete, summarized, or detailed log-entry information in a table format. You can view an individual log entry to analyze the raw log data, or use

filters to view a subset of log entries. You can also use column settings and flags to control how the UI presents log information. The Log Viewer displays each log entry as it enters the database in realtime, displaying its fields in the Log Viewer. For details, see [“Using the Log Viewer” on page 800](#).

- **Log Investigator**—Enables you to correlate log data. The Log Investigator is an exploratory data analysis tool that cross-tabulates on two dimensions. Log entries are linked to the Log Viewer, to help you perform an interactive analysis. For details, see [“Using the Log Investigator” on page 821](#).
- **Audit Log Viewer**—Tracks administrative changes made to a managed device by an NSM administrator. Log-entry details include the administrator that performed the change, when the change occurred, and the job results. For details, see [“Using the Audit Log Viewer” on page 832](#).

Device Limitations for Viewing Logs

For J Series devices running Junos 9.0 software and later versions, only partial structured syslogs are generated for logs. For these devices, all the non-structured log information is captured in the Details column and specifies “Self” for the Category and SubCategory columns in Log Viewer. Consequently, the following limitations apply for these devices:

- The **Jump to Policy** feature is not supported.
- Log Investigator analysis can only be applied to those partially structured syslogs that provide the source address and destination address in related columns.
- Log Viewer provides only limited support.

Configuring the Device for Logging

Before your managed device can generate log entries or log data, you must configure your devices and the NSM system for logging. You can configure an individual device to generate attack, alarm, configuration, information, and self log entries for specific destinations.

To view log entries and log data in the NSM UI, you must configure the individual device to generate log information for NSM, and enable one or more severity settings for NSM. However, you are not required to configure the settings for other destinations if you do not use those destinations for log management.

At the device level, you can configure how and where the device sends its log entries. For each destination, you can define:

- The category of log entries you want the device to generate and forward to a specific destination
- The severity of log entries you want the device to forward to a specific destination

The severity setting applies to all log types for that destination. For example, if traffic log entries are enabled for , but the severity setting specifies critical and major severities, receives only critical and major traffic logs; all other severity traffic log entries are

generated, but never sent to the management system. Unsent traffic log entries are stored on the device and discarded when the device log storage capacity is exceeded.

Configuring Severity Settings

Use the **General** settings to select the severity levels of the log entries you want to forward to a specific location. Juniper Networks assigns a predefined severity level for each event that generates a log entry on a managed device; using NSM, you can configure a device to send log entries with specific severity levels to specific destinations.

For each destination (except **Firewall Options**), you can specify one or more severity levels; for details on severity levels, see [“About Log Severity” on page 785](#).

Not all destinations support all log entry severities. [Table 95 on page 788](#) details the log entry severities accepted by each destination (except Firewall Options):

Table 95: Destinations of Log Entry Severities

Destination	Description	Severities
Console	The PC you use to view log entries in NSM.	All severities
E-mail	An e-mail server to which you want log information forwarded.	Emergency Alert Critical Notification
SNMP	A Simple Network Management Protocol destination.	Emergency Alert Critical
Syslog	The syslog server that you specify from the NSM UI or the log2action utility.	All severities
WebTrends	A WebTrends server to which you want log entries forwarded.	All severities
Network and Security Manager	The NSM server.	All severities
PCMCIA	A PCMCIA device to which you want log entries forwarded.	All severities
Internal	A destination within the current device to which you want log entries forwarded.	All severities

To select log entry severities for a destination, open a device configuration and select **Report Settings > General**, then select the destination.

Forwarding Self Log Entries (Firewall Options)

Self log entries typically display information on traffic that was dropped by the managed device or terminates on the device. Any packet that terminates at the device generates a self log entry; Telnet, Ping, BGP, and OSPF connections all terminate at the device, and trigger a self log entry.

A self log includes the date and time a packet was dropped, the source address of the packet, the destination address of the packet, the duration for which the packet was active, and the service associated with the packet. You can disable or enable logging of dropped packets for specific traffic types, including ICMP, IKE, SNMP, and multicast packets.

To configure self log entries, open a device configuration and select **Report Settings > General**. Click the **Firewall Options** tab and configure the following settings. See [Table 96 on page 789](#).

Table 96: Self Log Entry Settings

Setting	Description
Log ICMP Packets to Self	Creates a log entry for an ICMP (ping) packet that was dropped or terminated at the device.
Log IKE Packets to Self	Creates a log entry for an IKE packet that was dropped or terminated at the device. When negotiating an IKE key, the VPN client communicates with the security device.
Log SNMP Packets to Self	Creates a log entry for an SNMP packet that was dropped or terminated at the device.
Log Multicast Packets to Self	Creates a log entry for a multicast packet that was dropped or terminated at the device.

Configuring e-mail Server Settings

Use the **Email** option to configure a managed device to send messages using e-mail whenever a system event of Emergency, Alert, Critical, or Notification severity level occurs. You can configure the e-mail and SMTP settings at the device level, or skip this section and configure the GUI server to handle e-mail; see [“Exporting to E-mail” on page 842](#).

To configure e-mail server settings and enable the device to send e-mail messages, open a device configuration and select **Report Settings > Email**. Configure the following settings. See [Table 97 on page 789](#).

Table 97: Email Server Settings for Log Entries

Setting	Description
Enable Notification for Alarms	When alarm is enabled for a rule in the installed security policy and traffic matches the rule, the device sends an e-mail notification to the specified SMTP server.

Table 97: Email Server Settings for Log Entries (continued)

Setting	Description
Include Traffic Log	When logging is enabled for a rule in the installed security policy and traffic matches the rule, the device sends the traffic log entry to the specified SMTP server.
SMTP Server Name	The name of the Simple Mail Transfer Protocol server that receives e-mail notification messages. You must specify the SMTP server name and at least one e-mail address to receive e-mail notification.
Email Address 1	The primary e-mail address that receives e-mail notification messages from the device.
Email Address 2	The secondary e-mail address that also receives e-mail notification from the device.

Configuring Events Reporting Settings

Use the **Events** reporting settings to configure the managed device to report specific events to NSM.

Select the appropriate NSM Device Server, then select the events that are logged on the device and reported to NSM. The following sections detail each event.



NOTE: For security devices running ScreenOS 5.0 and later, you must also select **Enable Logging**.

Screen Alarm Log Entries

The device generates screen alarm log entries when a device detects network traffic that matches the screen settings enabled on the device.

To receive screen alarm log entries, you must:

- Enable the device to generate screen alarm log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with the desired severity settings to NSM in **Report Settings > General > NSM**.

Screen alarm log entries appear in the Log Viewer and display the following columns of information in the Log Viewer:

- Source Address
- Destination Address
- Service
- Action

- Category (Screen)
- Subcategory (for details on Screen subcategories, see Appendix E.)
- Severity

Event Alarm Log Entries

The device generates event alarms for any security event that has a predefined severity level of emergency, critical, or alert. Event alarms generate log entries that appear in the Alarm category.

To receive event alarm log entries, you must:

- Enable the device to generate event alarm log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with emergency, alert, and critical severity settings to NSM.

Event alarms appear in the Log Viewer under the Alarm category. For details on Attack subcategories, see [“Alarm Log Entries” on page 921](#).

Traffic Alarm Log Entries

The device generates traffic alarm log entries when your device detects network traffic that exceeds the specified alarm threshold in a security policy rule. The traffic alarm log entry, which displays in the Log Viewer, describes the security event that triggered the alarm. Traffic alarms generate log entries that appear in the Alarm category.

To receive traffic alarm log entries, you must:

- Enable the device to generate traffic alarm log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with the desired severity settings to NSM.
- Enable counting and alarms in the security policy installed on the device. For details on configuring traffic alarm logging in your security policy rules, see [“Configuring Counting and Alarms” on page 501](#).

Traffic alarms appear in the Log Viewer under the Alarm category. For details on alarm subcategories, see [“Alarm Log Entries” on page 921](#).

Alarm log entries contain information in the following Log Viewer columns:

- To Zone
- From Zone
- Source IP
- Destination IP
- Threshold (displayed in the Misc. column of the Log Viewer)

Deep Inspection Alarm Log Entries

The device generates Deep Inspection alarm log entries when a device with Deep Inspection (DI) detects network traffic that matches an attack object specified in a security policy rule. When matched in a rule, protocol anomaly attack objects, signature attack objects, and custom attack objects all generate Deep Inspection alarm log entries that appear in the Log Viewer.

To receive Deep Inspection alarm log entries, you must:

- Enable the device to generate Deep Inspection alarm log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with the desired severity settings to NSM in **Report Settings > General > NSM**.
- Enable Deep Inspection detection in the security policy installed on the device. For details on configuring Deep Inspection logging in your security policy rules, see [“Configuring a DI Profile/Enable IDP for Firewall Rules” on page 507](#).

Deep Inspection alarm log entries appear in the Log Viewer and display the following columns of information in the Log Viewer:

- Source Address
- Destination Address
- Service
- Action
- Category (Predefined or Custom)
- Subcategory (for details on Deep Inspection alarm subcategories, see [“Deep Inspection Alarm Log Entries” on page 922](#))
- Severity

Configuration Log Entries

The device generates configuration log entries for events that change the configuration on the device. Specifically, any command issued that the ScreenOS **get config** command statement captures and displays in ScreenOS generates a configuration log. For each configuration change, the device generates a configuration log entry that contains information about the change in the Log Viewer Detail column.

To receive configuration log entries, you must:

- Enable the device to generate configuration log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with a notification severity setting to NSM in **Report Settings > General > NSM**.

Configuration log entries appear in the Log Viewer under the category Configuration. For details on configuration subcategories, see [“Configuration Log Entries” on page 997](#).

Information Log Entries

The device generates information logs when it detects that an administrator has made a change to the basic settings of the device, such as logging in or out, setting a new password for the device, issuing a key value for the device, or entering an MD5 authentication password to enter a device. For each administrative change, the device generates an information log entry that contains information about the change in the Log Viewer Detail column.

To receive information log entries, you must:

- Enable the device to generate information log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with the info, warning, and error severity settings to NSM in **Report Settings > General > NSM**.

Information log entries appear in the Log Viewer under the category Information. For details on information subcategories, see [“Information Log Entries” on page 999](#).

Self Log Entries

The device generates self log entries for any packet that terminates at the device. Self log entries display information on traffic that was dropped by the device or that terminates on the device.

To receive self log entries, you must:

- Enable the device to generate self log entries for NSM in **Report Settings > NSM**.
- Enable the device to send specific self log entries to NSM in **Report Settings > General > Firewall Options**. For details, see [“Forwarding Self Log Entries \(Firewall Options\)” on page 789](#).

Self log entries appear in the Log Viewer under the category Self, which contains information in the following Log Viewer columns:

- Source
- Destination
- Services

Self log entries have the category “Self” and the subcategory “Self Log”.



NOTE: In some cases, Junos devices return log messages in syslog format, which map to “Self” for Category and Subcategory columns. Consequently, self log entries are not necessarily the result of packets that terminate at the device or packets that were dropped by a security device.

Traffic Log Entries

The device generates traffic log entries when your device detects network traffic that matches the source, destination, and service specified in a security policy rule.

To receive traffic log entries, you must:

- Enable the device to generate traffic log entries for NSM in **Report Settings > NSM**.
- Enable the device to send log entries with the desired severity settings to NSM.
- Enable logging in the security policy installed on the device. For details on configuring traffic logging in your security policy rules, see [“Configuring Logging and Alerts” on page 501](#).

Traffic log entries appear in the Log Viewer under the Traffic category. For details on traffic subcategories, see [“Traffic Log Entries” on page 1001](#).

Policy Statistics

The device forwards statistics on the policy distribution of the traffic that entered the device. Policy distribution statistics do not generate log entries; the statistics are used by the **Realtime Monitor** module. For details on how policy distribution is displayed in **Realtime Monitor**, see [“Viewing Traffic Distribution by Security Policy” on page 715](#).

Attack Statistics

The device forwards statistics for attacks detected in the traffic that entered the device. Attack statistics do not generate log entries; the statistics are used by the Realtime Monitor module. For details on how attack statistics are displayed in the Realtime Monitor, see [“Viewing Attack Statistics” on page 724](#).

Ethernet Statistics

The device forwards statistics for Ethernet activity on the device. Ethernet statistics do not generate log entries; the statistics are used by the Realtime Monitor module. For details on how Ethernet statistics are displayed in Realtime Monitor, see [“Viewing Ethernet Statistics” on page 721](#).

Flow Statistics

The device forwards statistics for flows that entered the device. Flow statistics do not generate log entries; the statistics are used by the Realtime Monitor. For details on how flow statistics are displayed in the Realtime Monitor, see [“Viewing Flow Statistics” on page 723](#).

Protocol Distribution

The device forwards information on the protocol distribution of the traffic that entered the device. Protocol distribution information does not generate log entries; the information is used by the Realtime Monitor module. For details on how protocol distribution is displayed in the Realtime Monitor, see [“Viewing Traffic Distribution by Protocol” on page 716](#).

The device reports statistics generated by the following services:

- AH (Authentication Header)
- ESP (Encapsulating Security Payload)
- GRE (Generic Routing Encapsulation)
- ICMP (Internet Control Message Protocol)
- OSPF (Open Shortest Path First)
- TCP (Transmission Control Protocol)
- UDP (User Datagram Protocol)

You can also set the interval at which the NSM **Device Server** polls for policy statistics and protocol distribution events.

Atomic Updating Events

Devices running ScreenOS 5.0 and later support atomic updating, which enables the device to receive the entire modeled configuration (all commands) before executing those commands (instead of executing commands as they are received from the management system). Atomic updating also might cause the device to temporarily lose connection to NSM during the update process.

If the device cannot reconnect to the management system after processing the update, it automatically reboots (with the previously saved configuration) and reconnects to the management system. To prevent a device from rebooting or to configure the reboot timeout, open a device configuration and select **Report Settings > Events**, then configure the **Atomic Updating** options.

For details on **Atomic Updating**, see [“About Atomic Updating—ScreenOS Devices” on page 260](#).

Configuring SNMP Reporting Settings

Use SNMP settings to configure the Simple Network Management Protocol (SNMP) agent for the managed device. The SNMP agent provides a view of statistical data about the network and the devices on it, and notification of system events of interest. You can configure the SNMP settings at the device level, or skip this section and configure the GUI server to handle SNMP reporting; see [“Exporting to SNMP” on page 842](#).

In addition to configuring the SNMP reporting settings, you also must enable SNMP management service options on the interface through which the SNMP manager application communicates with the SNMP agent in the managed device.

To configure SNMP settings and enable the device to send SNMP traps, open a device configuration and select **Report Settings > SNMP**. Configure the following settings:

- System Name—The name of the device for which you are generating SNMP status.
- Contact Person—The name of the network administrator who manages the device. This contact information is useful when the SNMP community member needs to contact someone about the device.

- Location—The physical location of the device.
- Listen Port—The number of the port assigned to monitor SNMP traffic (listen and transmit SNMP traps).
- Trap Port—The number of the port assigned to transmit traps that have been generated by an SNMP alarm, threshold violation, or error.
- Enable Authentication Fail Trap—Specifies whether you want to generate a trap if a packet fails to be authenticated when attempting to enter the device. Select this option if the device sends SNMP messages through a VPN tunnel.

Next, configure SNMP communities. To send traps, the SNMP agent on the device requires that you define communities, their associated hosts, and assign permissions (read/write or read-only). You can create up to three (3) SNMP communities, with up to eight (8) hosts in each community.

To create an SNMP community, click the **Add** icon under **Community Settings** and configure the following settings:

- Community name—The device uses the community name to authorize users attempting to enter the device.
- Access Mode—Defines read-write or read-only privileges for the community.
- Trap Mode—When enabled (On), enables the device to send an SNMP trap for illegal SNMP connections attempts to the device.
- Traffic—When enabled, the device can accept traffic from the source interface.
- Version—Defines the versions supported by the community (SNMPv1, SNMPv2c, or both SNMP versions, as required by the SNMP management stations). For backward compatibility with earlier ScreenOS releases that only support SNMPv1, security devices support SNMPv1 by default.
- Hosts—Define one or more hosts that are associated with the community. Click the **Add** icon, then specify the host IP address and netmask, the trap version for the host (if an SNMP community supports both SNMPv1 and SNMPv2c, you must specify a trap version for each community member), and the source interface.

Directing Logs to a Syslog Server

A managed device can generate syslog messages for system events at predefined severity levels and optionally for traffic that policies permit across a firewall. It sends these messages via UDP (port 514) to up to four designated syslog hosts running on UNIX/Linux systems. When you enable syslog reporting, you also specify which interface the devices use to send syslog packets.

You can configure the syslog server settings at the device level, or skip this section and configure the GUI server to handle syslog messages; see [“Exporting to the System Log” on page 841](#).

To send log entries to a Syslog server, click the **Syslog** option. NSM displays the **Syslog** dialog box. Enter appropriate data into the following fields. See [Table 98 on page 797](#).

Table 98: Syslog Settings for Log Entries

Field	Description
Enable Syslog Messages	Initiates the logging of system event messages to the Syslog server.
Port Number	Indicates the port number from where the messages are sent to the Syslog server.
Use Trust Zone Interface as Source IF for VPN	Specifies using the interface mapped to the Trust zone as the source of traffic for a VPN.
Include Traffic Log	Specifies that all traffic log events are included as part of the messages sent to the syslog server.
Config Host	Indicates the name of the host device.

Standalone IDP sensors running release 5.1 and later support the configuration of:

- The protocol to either UDP (the default) or TCP.
- The port. The default port is 514. However, you can set the port value within the range of 1 through 65535.

Directing Logs from DMI Devices

DMI devices send logs to NSM using one of the following modes:

- Event mode (default). The logs sent using this mode are control plane logs and typically consist of user process-generated logs, system daemons-generated logs, and so on. This mode does not require any additional configuration.
- Stream mode. The logs sent using this mode are data plane logs and typically consist of session logs, IDP logs, and so on. This mode requires additional configurations on NSM and the device.

Configuring the DMI Device for Stream Mode

To configure the DMI device to send the logs to NSM using stream mode:

1. Edit the `/var/netscreen/DevSvr/devSvr.cfg` file and set the `devSvr.enableSyslogOverUdp` parameter to **true**:


```
devSvr.enableSyslogOverUdp true
```
2. Restart the DevSvrMgr.
3. In the NSM navigation tree, select **Device Manager > Devices**.
4. Click the **Device Tree** tab, and then double-click the device for which you want to configure stream mode.
5. Click the **Configuration** tab. In the configuration tree, select **Security > Log**.
6. Configure the following parameters:

- **Mode**—Select **stream**.
 - **Format**—Select **sd-syslog**.
 - **Source Address**—Enter the address of the DMI device.
7. Select **Security > Log > Stream**.
 - a. Click the Add icon.
 - b. In the **Name** field, enter a unique name for the stream.
 - c. In the **Format** drop-down list, select **sd-syslog**.
 8. Select **Host**.
 - a. In the **Ipaddr** field, enter the IP address of the NSM server.
 - b. In the **Port** field, enter **5140**.

For HA setup, complete Step 8 for the primary DevSvr and the secondary DevSvr.
 9. Update the device.

Directing Data to a WebTrends Server

The managed device can send syslog reports to a Webtrends Syslog host. Webtrends offers a product called the Webtrends Firewall Suite that enables you to customize syslog reports to display the information you want in a graphical format.

To send log entries to a WebTrends server, click **WebTrends**. NSM displays the **WebTrends** dialog box. Enter appropriate data into the following fields.

Table 99: WebTrends Settings for Log Entries

Field	Description
Enable WebTrends Message	Directs NSM to forward a log to the WebTrends server.
WebTrends Host Name	The name of the WebTrends server.
Port	Specifies the port number through which the device sends the log to the WebTrends server.
Use Trust Zone Interface as Source IP for VPN	Directs the device to use the interface mapped to the trust zone as the location for the VPN over which the packets are forwarded to the WebTrends server.

To set severity levels for WebTrends destinations, click on **Log Settings** under **Report Settings** in the navigation tree. Then click the **WebTrends** tab and click the desired severity check box.

Managing Packet Data in Logs

Packet data can be stored on IDP sensors or included with log data sent to the NSM server. Use NSM to configure the IDP sensor to:

- Store packet data on the IDP sensor, which NSM can later retrieve. For IDP 4.1 and later, this option is the default setting and improves performance.
- Send packet data to NSM along with log data and store it on the NSM server. For IDP 4.0, this option is the default setting.



NOTE: This feature is available for IDP sensors only.

There is a maximum limit of how much packet data can be stored on the IDP sensor. When the limit is reached, the IDP sensor overwrites older packet data, purging it from the sensor. The maximum size for each log is 800KB. Packet data stored on the NSM server, whether retrieved after being stored on the IDP sensor or included with the log data, is stored permanently on the NSM server until or unless it is purged by the user.

To store the packet data on the IDP sensor, double-click an IDP sensor, select **Report Settings** in the navigation tree, and then uncheck **Include packet data in log**.

To send the packet data to NSM with the log data, double-click an IDP sensor, select **Report Settings** in the navigation tree, and then select **Include packet data in log** if it is not already selected.

To view a log with packet data, go to the main navigation tree and select Log Viewer, right-click the log containing the packet data, and then select **Show > Packet Data**. See [Figure 113 on page 799](#).

Figure 113: View Packet Data in a Log

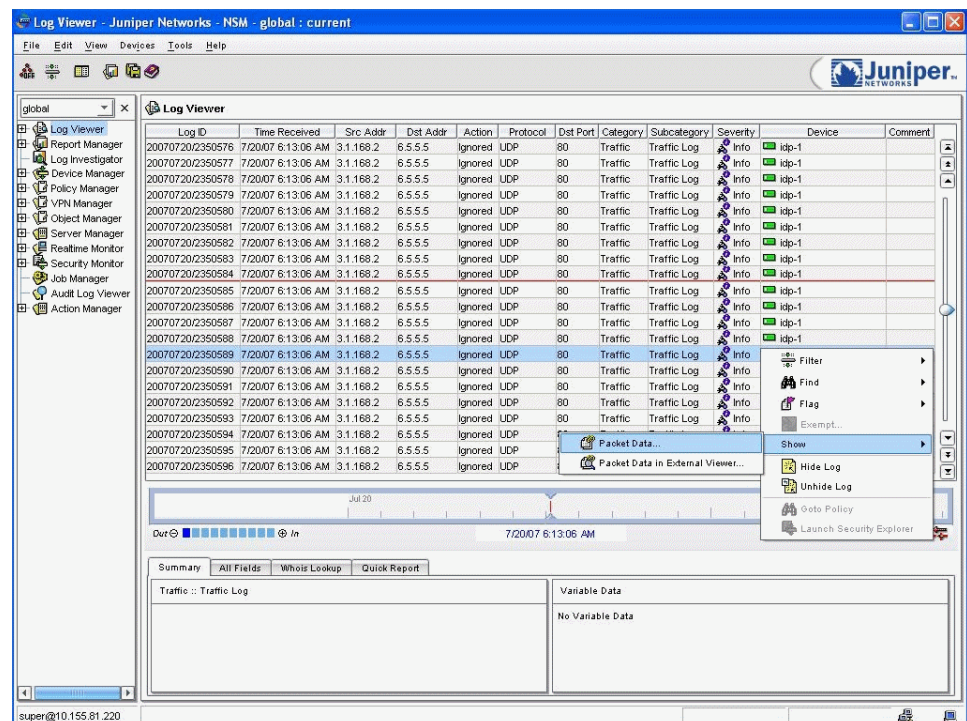
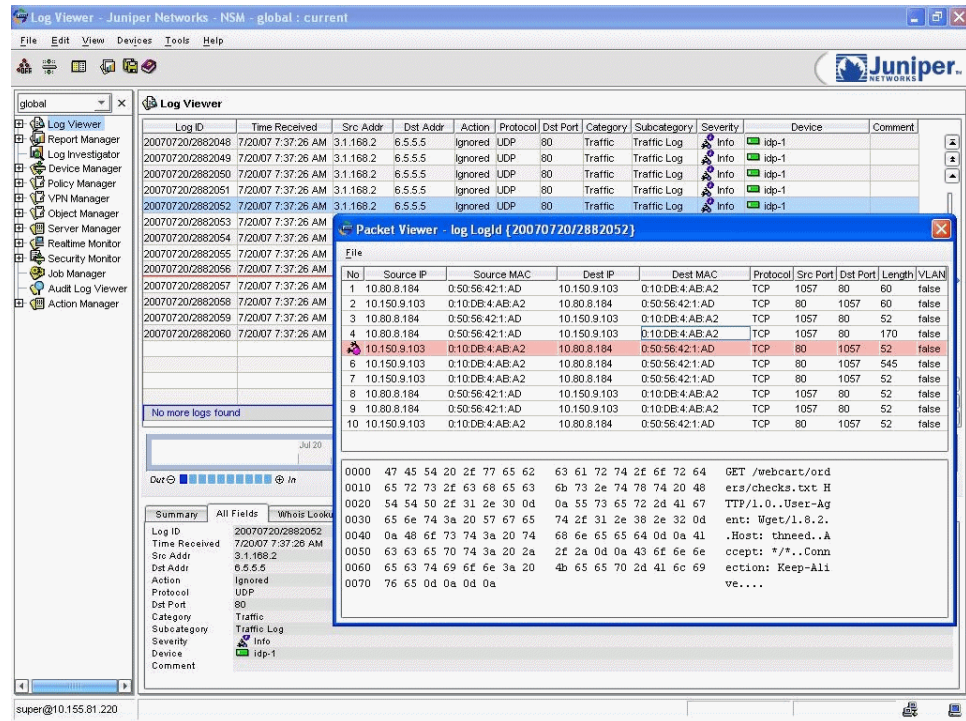


Figure 114 on page 800 provides an example of packet data.

Figure 114: Sample Packet Data



Using the Log Viewer

The Log Viewer displays log entries generated by a managed device when traffic matches a firewall or VPN rule, or when an event occurs that matches a predefined set of conditions. The main display of the Log Viewer displays summarized information about security events and alarms, while the detail panes provide more detailed information about a specific log entry.

This section provides details on the following Log Viewer functionality:

- **“Using Log Views” on page 801**—The Log Viewer includes several predefined views for critical severity attacks, configuration log entries, scans, and other important activity. This section describes how to use filters to create your own unique, customized log entry view, then save the custom view (with all its filters) for future use.
- **“Searching Log Entries” on page 808**—For networks that generate large numbers of log entries, it can be difficult to locate the exact log entries that detail the events you want to investigate. This section describes how to use the log timeline to find logs generated around a specific time, how to use the find utility to locate log entries with a specific value, and how to search by log ID to jump directly to a specific log entry.
- **“Filtering Log Entries by Event and Time” on page 812**—This section describes how to create custom filters based on event data or time. You can apply these filters to a Log Viewer column or cell to reduce the number of log entries that appear in the main display area, helping you to focus your investigations on a specific group of log entries.

- [“Filtering Log Entries by Range” on page 815](#)—This section describes how to create custom filters based on a user-defined range. You can apply these filters to a Log Viewer column or cell to reduce the number of log entries that appear in the main display area, helping you to focus your investigations on a specific group of log entries.
- [“Customizing Columns” on page 816](#)—The Log Viewer includes 46 columns of log entry information; however, each predefined view includes only a small subset of available columns. This section describes how to set viewable columns, change column display order, resize columns, and hide or unhide a specific column.
- [“Filtering Log Entries by Column” on page 818](#)—This section describes how to select one or more columns for filtering log entries and select filter settings for specified columns.
- [“Using Log Viewer Integration” on page 819](#)—This section describes how to use the Log Viewer integration to jump from a log entry directly to the responsible security policy or managed device configuration.
- [“Identifying Irrelevant Attacks” on page 820](#)—Irrelevant attacks are events that do not affect your network or that you do not consider important. For example, if you do not run an Apache Web server on your network, you do not need to worry about attacks against Apache Web servers. This section describes how to use your log entries to identify irrelevant attacks, then eliminate the attack object group that generated that attack from your security policy.

Using Log Views

The Log Viewer enables you to view and analyze logs generated by the managed devices in your network. For your convenience, many predefined views are included for critical severity attacks, authentication failures, configuration log entries, scans, and other important activities. Using filters, you can create your own unique, customized log entry view, then save the custom view (with all its filters) and manage them in folders for future use.

About Predefined Log Views

The Log Viewer provides several groups of predefined log views. By default, all predefined log views display the latest logs.

[Table 100 on page 801](#) lists and describes the EX Switch predefined log views.

Table 100: EX Series Switch Predefined Log Views

Log Type	Description
All-Switch-logs	Filters logs on devices whose device family name is junos-ex .
Chassis-logs	Filters CHASSISD logs on devices whose device family name is junos-ex .
Config-change-logs	Filters UI_COMMIT and UI_COMMIT_AT_COMPLETED logs on devices whose device family name is junos-ex .

Table 100: EX Series Switch Predefined Log Views (continued)

Log Type	Description
Topology-logs	Filter LLDPD logs on devices whose device family name is junos-ex .
User-interaction-logs	Filters UI logs on devices whose device family name is junos-ex .

[Table 101 on page 802](#) lists and describes the SSL/UAC Predefined log views.

Table 101: SSL/UAC Predefined Log Views

Log Type	Displays all logs filtered by
Admin	Category—Admin (13)
Cluster	SUBCATEGORY: SYS10061, SYS10062
Dynamic Policy Evaluation	Subcategory—AUT23523, AUT23524
Events	Category—Events (14)
Hardware	Subcategory—SYS24013, SYS24014, SYS24015, ERR24016, SYS24017, SYS24018, ERR24019, SYS24020, SYS24021, SYS24074, SYS24075, ERR24076, SYS24077, SYS24078, ERR24079
Host Checker	Subcategory—AUT22923, AUT22925
Network Ports	Subcategory—NET24462, NET24463,
Sensor Initiated Actions	Subcategory— SUBCATEGORY: IDP24101, IDP24102, IDP24103, IDP24104, IDP24105, IDP24106, IDP24107, IDP24108, IDP24109, IDP24190, IDP24191
Sensors	Category—(sensors)(15)
System Restarts	Subcategory—SYS10298, SYS10299, SYS10314, SYS24258, SYS24259
User	Category— User(12)
VLAN Assignments	Subcategory—EAM24459

[Table 102 on page 802](#) lists and describes the Predefined log views.

Table 102: Predefined Log Views

Log Type	Displays all logs filtered by
Critical	Severity—Critical

Table 102: Predefined Log Views (continued)

Log Type	Displays all logs filtered by
Alarm	Category—Alarm
IDP/DI	Category—Custom, Predefined
Screen	Category—Screen
Traffic	Category—Traffic
Info	Category—Info
Config	Category—Config
Self	Category—Self
Backdoor	Subcategory—Backdoor Detected (Traffic), Backdoor Dropped (Traffic)
PolicyLogViewer	Policy, Rule #, Rule Domain, Rule Domain Ver, Rulebase
Profiler	Category—Profiler
Scans	Subcategory—Distributed Port Scan, Distributed Port Scan in Progress, ICMP Sweep, ICMP Sweep In Progress, Network Scan , Network Scan In Progress, TCP Port Scan, TCP Port Scan In Progress, TSIG Session Rate Exceeded, UDP Port Scan, UDP Port Scan In Progress

Creating Custom Views and Folders

A custom view enables you to organize log entries in a format that is most helpful to you. Because the custom view is based on filters, incoming log entries that match the filter criteria are automatically displayed in the view. You do not need to reapply the view to new logs.

You might want to create views to help manage the following situations:

- **Workflow**—To help a team of security administrators work together to investigate and resolve incidents, create a view that filters on the flag column of the Log Viewer to indicate the status of each log entry and assignment.
- **Attackers**—To track the activities of a known attacker, create a view that filters on a specific source IP. The source IP address of an attack appears in the source address column, and the destination IP address of an attack appears in the destination address column.

- **Alarms**—To quickly access log entries generated by a policy rule that contains an alarm, create a view that filters on the alarm column. This method is useful when you are fine-tuning policies to distinguish between genuine attacks and false positives.
- **Devices**—To manage devices in multiple locations that use different investigation processes, create a separate view for each device at a specific location.

You can create and save custom views using one of the following methods:

- **Create New View**—In the navigation tree, select the Log Viewer module. From the **File** menu, select **New View**. In the **New View** dialog box, enter a name for the custom view, enter a name for the folder that you want to save the view in, and click **OK**. The new view is displayed in the navigation tree in the folder specified. By default, all new views are saved in the Others folder. In the main display area, you can then set the desired filters for the log entries.
- **Set Filters**—In the Log Viewer main display area, set the desired filters for the view. From the **File** menu, select **Save As**. In the **New View** dialog box, enter a name for the custom view, enter a name for the folder that you want to save the view in, and click **OK**. The new view is displayed in the navigation tree in the folder specified.

Creating Per-Session Views

Log views that you create on the fly, also called “transient” views, (views set from filters defined in the Report Manager), appear in the **Drill Down** folder under Log Viewer. These views remain in this folder until you log out of the UI.

Log Viewer Columns

The Log Viewer contains the columns in [Table 103 on page 804](#). When filtering by column, the filter affects all log entries.

Table 103: Log Viewer Columns

Column	Default	Meaning
Log ID	Default	The unique identifier ID for the log entry. The log ID comprises both a date and an incrementing integer.
Time Received	Default	The date and the time that the Log Viewer received the log entry.
Alert	Default	Indicates whether an alert flag was generated in response to the event that generated the log entry.
User Flag	Default	The UI assignable flag associated with the current log.
Src Addr	No	The source address of the packet that generated the log.
Dst Addr	Default	The destination device to which the packet associated with the log entry was targeted.
Action	Default	The device action performed on the packet / connection that generated the log.

Table 103: Log Viewer Columns (continued)

Column	Default	Meaning
Protocol	Default	The protocol used by the packet that generated the log..
Dst Port	No	The destination port of the packet that generated the log.
Rule #	Default	The rule in the rulebase in the policy in the specific version of an object that generated this log entry.
Nat Src Addr	Default	The source address that has undergone NAT, and is associated with the packet that generated the log entry.
Nat Dst Addr	Default	The destination address that has undergone NAT, and is associated with the packet that generated the log entry.
Details	No	A string that captures device log data that could not otherwise be stored in a column.
Category	Default	The type of log you are viewing. Can be expressed either as a category or a subcategory. A category is either admin, alarm, config, custom, event, implicit, info, predefined, profiler, screen, self, sensors, traffic, urlfiltering, or user. A subcategory is an attack type.
Sub Category	Default	Subcategory of the log entry you are viewing.
Severity	Default	Level of severity associated with the attack detected. Every attack has a default severity level although you can configure a different one.
Device	Default	The device that generated the log entry.
Comment	Default	Enables you to comment about relates to the generated log entry. To enter a comment, click in the cell and enter text.
App	Default	The application associated with this log.
Application Name	Default	Service application name
Bytes In	No	Number of bytes that comprised the log data entering the Log Viewer per session.
Bytes Out	No	Number of bytes that comprised the log data being transmitted from the Log Viewer per session.
Bytes Total	No	The sum of the number of bytes transmitted and received by the Log Viewer.
Dev Domain Ver	No	The version of the object that contained the device that generated the log entry.

Table 103: Log Viewer Columns (continued)

Column	Default	Meaning
Device Domain	No	Name of the domain in which the device resides. Note: The Log Viewer displays log entries for a single domain at a time. By default, when logged in as the super administrator, the Log Viewer displays log entries for managed devices in the global domain. To change the domain, apply a domain filter to view log entries for managed devices in a specific domain.
Device family	No	The family of the device that generated the log entry.
Dst Inf	No	The name of the packet-centric outbound interface.
Dst Zone	No	The destination zone associated with a traffic log entry. Log entry data for this column only appears after you update the security device.
Elapsed Secs	No	The number of seconds that have elapsed since the beginning of the current session.
Has Packet Data	No	Specifies if this log has associated packet data.
Nat Dst Port	No	A destination port that has undergone NAT and is associated with the packet that has generated the log.
Nat Src Port	No	A source port that has undergone NAT and is associated with the packet that has generated the log.
Packets In	No	The number of received packets for a given session on the current port.
Packets Out	No	Number of transmitted packets for a given session on the current port.
Packets Total	No	Aggregate number of both received and transmitted packets for a given session on the current port.
Policy	No	The policy in a specific domain version that generated the log.
Policy ID	No	The unique policy rule number that generated the log. This policy number is constant in both ScreenOS and NSM.
Roles	No	A role group to which the user belongs.
Rule Domain	No	The domain that contained the rule that generated this log.
Rule Domain ver	No	The domain version containing the rule that generated this log.

Table 103: Log Viewer Columns (continued)

Column	Default	Meaning
Rulebase	No	The policy rulebase of a specific domain version that generated this log.
Src Intf	No	The name of the packet-centric inbound interface.
Src Port	No	The TCP/UDP port number of the source device that generated the packet that generated the log entry.
Src Zone	No	Source zone associated with a traffic log entry. Log entry data for this column only appears after you update the security device. The source zone is the zone that is attempting to send the traffic through the security device.
Time Generated	No	The time the current log was generated.
User	No	The user associated with this log.

After you import a device configuration, log entries from that device begin to appear in the Log Viewer.



NOTE: The Policy ID is supported for new logs from devices running ScreenOS 6.3 and later, and Junos firewall devices. The Policy ID column remains empty for older logs.

Log Viewer Detail Panes

The Log Viewer contains additional panes that provide summary and detail information for log entry events. To see detailed information about a log entry, select the entry and view the detail panes at the bottom of the Log Viewer. The detail pane contains four tabs of information about the selected log record:

- **Summary** tab (default tab)—Details the event associated with the selected log entry. Within the summary tab, you can view the event description (right side) and the variable data (left side). Not all log entries contain variable data—only log entries generated by an attack provide variable data.
- **All Fields** tab—Provides a condensed view of data for the selected log entry (so you do not need to scroll horizontally).
- **Whois** tab—Enables you to perform a Whois lookup on an IP address to see what organization has registered a particular address.
- **Quick Reports** tab—Enables you to quickly generate a predefined report on a filter criteria in the Log Viewer.

Log Viewer Status Bar

The status bar of the Log Viewer summarizes the filters applied to log entries in the log entry list. The status bar displays the filter type description.

For example, [Figure 115 on page 808](#) shows the status bar for log entries that are filtered by Category and Severity:

Figure 115: View Category and Severity Filters










To view filter details, place the cursor over the filter type.

Navigating the Log Viewer

Using the side scroll bar, you can navigate through hundreds of log entries quickly and precisely, as shown in [Table 104 on page 808](#):

Table 104: Log Viewer Navigation Controls

Scroll Bar Components	Function
	Jump to top of log entry list.
	Page up within log entry list.
	Scroll up within log entry list.
	Use the slider to move up or down within log entry list. The farther you drag the slider from the center, the faster you scroll through the log entry list.
	Scroll down within log entry list.
	Page down within log entry list.
	Jump to bottom of log entry list.

Log entries higher in the list are older than log entries at the bottom of the list. To navigate through log entries based on a specific time, use the Log Timeline (for details, see [“Log Timeline” on page 809](#)).

Searching Log Entries

The Log Viewer can receive thousands or even millions of log entries each day. To quickly locate a specific log entry or logs, use the log searching tools in [Table 105 on page 809](#).

Table 105: Search Tools for Log Viewer

Tool	Description	Benefit
Log Timeline	A 14-day timeline that enables you to zoom to log entries for a specific day and time.	<ul style="list-style-type: none"> Specify an exact date and time, or use the timeline selection slider to move immediately to a specific day's log entries. Use the Tailing Logs feature to jump directly to incoming log entries. Timeline covers any 14-day period, in increments of days, hours, or minutes.
Flags	A symbol used to tag a specific log entry that you want to return to at a later point.	<ul style="list-style-type: none"> Gain greater control over identifying events. The flagged entry stands out from other entries, making it easier to locate quickly. You can filter on a flag setting.
Find Utility	A string search that searches for a log entry based on a character string in the reported event.	<ul style="list-style-type: none"> Locate a specific event quickly with minimal detail; for example, search using the timestamp or IP address field. Move quickly from one relevant event to another, avoiding scrolling.
Log ID Number	A value search that searches for a log entry based on the log ID number.	Locate a specific log entry immediately. Typically, you use a log ID search when you have previously viewed the log entry and need to find it again quickly.

The following sections detail each search tool.

Log Timeline

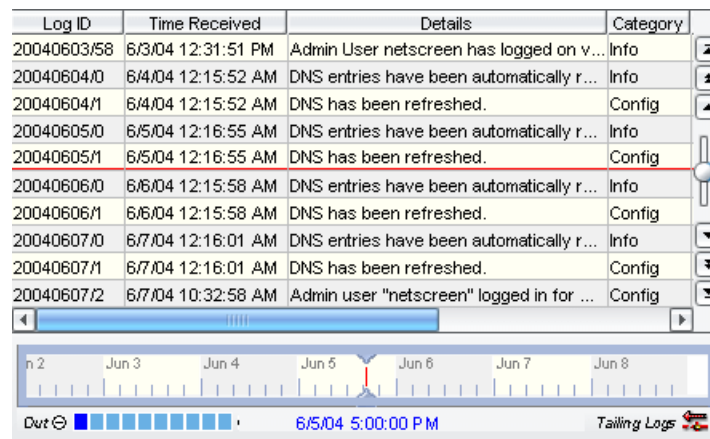
The log timeline is a powerful tool that enables you to quickly jump to a group of log entries generated during a specific time period.

The timeline consists of 4 components: the time slider, time entry, time blocks, and Tailing Logs. The following sections detail each component.

Using the Time Slider

The time slider marks the midpoint of the time interval selected for the timeline (for details on setting a time interval, see [“Using Time Blocks” on page 810](#)). You can move the time slider to the desired time using your mouse cursor: Click the slider (the vertical red line), then drag it to the area on the timeline that represents the date and time around which you want to view log entries. The log entry list automatically jumps to the selected date and time (shown by the horizontal red line). [Figure 116 on page 810](#) shows the time slider.


Figure 116: Log Viewer Time Slider



Using Time Entry

You can also enter a date and time into the log timeline directly. Enter the date and time, and click **Go** as shown in [Figure 117 on page 810](#).

Figure 117: Log Viewer Time Display



Using Time Blocks

To change the log timeline intervals, select a specific time block or use the **Out** and **In** buttons. From left to right, the time blocks are:

- 14 days
- 7 days
- 3 days
- 1 day
- 12 hours
- 6 hours
- 3 hours
- 1 hour
- 30 minutes
- 1 minute

Click the **Out** button to select the time block to the left of the currently selected time block. Click the **In** button to select the time block to the right of the currently selected time block. Alternatively, you can use the mouse wheel on your mouse to adjust the time interval.

Using Tailing Logs

To view arriving log entries, select **Tailing Logs**. The log entry list automatically jumps to the bottom of the list, where new log entries appear when they are received by the management system. As older log entries are moved up by arriving log entries, the view remains fixed at the bottom of the list.

Tailing Logs also works with filters, predefined views, and custom views.

To see configuration log entries as they arrive from a specific device:

1. Select **Predefined >7-Config** view from the Log Viewer . This view uses a predefined filter to display log entries with the “configuration” category only.
2. Set a custom filter to set log entries from a specific device (for details, see [“Filtering Log Entries by Event and Time” on page 812](#)). The view changes to display configuration log entries from that device.
3. Select Tailing Logs. The view jumps to the bottom of the log entry list, and remains there; as new configuration log entries for the device arrive, they appear at the bottom of the list.

Using Flags

Use a flag to mark a specific event with a severity or workflow marker. Applying a flag to a log entries helps the event stand out from other log entries. [Table 106 on page 811](#) shows the flag severity levels and corresponding filter types.

Table 106: Log Viewer Flags

Flag/Severity Level	Filter Type
High	Severity
Medium	Severity
Low	Severity
Closed	Information
False Positive	Information
Assigned	Information
Investigate	Information
Follow-Up	Information
Pending	Information

Within the Log Viewer, you can set a filter on one or more flags. Additionally, within Report Manager, you can generate a report that displays the count of all log entries that contain a specific flag.

Using the Find Utility

Use the Find utility to search for the next iteration (down) of a value in the Log Viewer. To use Find on a column or cell, right-click the column header or cell and select Find, then configure the search criteria. Select Negate to search for all log entries that do not contain the specified value.

Using Log ID Number

When you know the log ID number for the log entry, you can jump directly to the log entry. To locate a log entry by log ID number, from the **Edit** menu, select **Go To Log ID**. In the **Go To Log ID** dialog box, enter the log ID number, and click **OK**. The Log Viewer jumps to the specified log ID and highlights the log entry in the main display area.

Filtering Log Entries by Event and Time

An event-based or time-based filter is a criteria search for matching log entries. When you apply a filter to log entries, the Log Viewer filters out log entries that do not match the filter criteria. You can set multiple filters on any Log Viewer column (except Log ID and Details) or cell value.

When filtering by cell, the filter affects only the content in that cell's column. To set a cell filter, right-click a cell and select **Filter** to display the filter menu options:

- **Edit**—Use this option to set multiple filters for cell content at the same time. Select to display the **Filter** dialog box for that column, then select the columns you want to filter on.
 - To display only the selected content, click **OK**.
 - To display everything except the selected content, select the check box next to **Negate**, then click **OK**.
 - To clear filters for the selected content, click **Clear**.
- **Only This Value**—Displays only the content in the selected cell.
- **Not This Value**—Displays everything except the content in the selected cell.
- **Clear Filter**—Removes a current filter on the selected cell content. If no filter exists, this option is unavailable.
- **Clear All Filters**—Removes all filters on the current view.

When filtering by column, the filter affects all log entries. You can set an event-based filter using any log entry column that contains event data, and a time-based filter for the Time Generated and Time Received columns. Additionally, for all filters (cell or column), you can enable the Negate option to match all log entries that do not contain the specified filter criteria.

The following sections detail some common event-based and time-based filters used to manage log entries.

Setting a Category Filter

Apply a category filter to view log entries within a specific category or subcategory.

- To create a category filter, right-click the **Category** column header and select **Filter > Set Filter**. Select the categories you want to use as the filter criteria, then click **OK**. NSM applies the filter to all log entries and displays only the log entries that match the specified category.
- To create a subcategory filter, right-click the **Subcategory** column header and select **Filter > Set Filter**. Select the category first, then select the subcategories you want to use as the filter criteria, then click **OK**. NSM applies the filter to all log entries and displays only the log entries that match the specified subcategory.

Setting an Alert Filter

Apply an alert filter to view log entries that have an enabled or a disabled alert state. To create an alert filter, right-click the Alert column header and select **Filter > Set Filter**, then configure the alert filter settings:

- To display log entries that contain an enabled alert, select **On** and click **OK**.
- To display log entries that contain a disabled alert, select **Off** and click **OK**.

NSM applies the filter to all log entries and displays only the log entries that match the specified alert state.

Setting a Flag Filter

Apply a flag filter to view log entries that have a specified flag type. To create a flag filter, right-click the Flag column header and select **Filter > Set Filter**. Select the flag types that you want to use as the filter criteria, then click **OK**. NSM applies the filter to all log entries and displays only the log entries that match the specified flag type.



NOTE: The **Unflagged** option in the flag filter can be helpful when trying to locate log entries that do not have assigned flags. When setting the flag criteria, select **Unflagged** as the flag type; NSM then displays all log entries without flags.

Setting an Address Filter

Apply an address filter to view log entries that record events for a specific source or destination address, or source or destination NAT address. To create an address filter, right-click the Src Addr, Dst Addr, Nat Src Addr, or Nat Dst Addr column header and select **Filter > Set Filter**. Select **Click here to add address** and enter a valid IP address and click **OK**. For NAT addresses, enter the IP address that is translated and click **OK**.

NSM applies the filter to all log entries and displays only the log entries that match the specified IP address.

Setting a Protocol Filter

Apply a protocol filter to view log entries for events that use a specific protocol type. To create a protocol filter, right-click the **Protocol** column header and select **Filter > Set Filter**. Select the protocol types that you want to use as the filter criteria, then click **OK**. NSM applies the filter to all log entries and displays only the log entries that match the specified protocol types.

Setting a Domain Filter

The Log Viewer displays log entries for a single domain at time. By default, when you log in as the super administrator, the Log Viewer displays log entries for managed devices in the global domain. To change the domain apply a domain filter to view log entries for managed devices in a specific domain.

To create a domain filter, right-click the **Domain** column and select **Filter > Set Filter**. Select the domain for which you want to view log entries, then click **OK**. NSM applies the filter to all log entries and displays only the log entries that are generated from managed devices in the specified domain.

Setting a Time-Based Filter

Apply a time-based filter to view log entries generated or received within a specific time period. To create a time-based filter, right-click the **Time Generated** or **Time Received** column and select **Filter > Set Filter**:

- To filter on a specific start time, select **From** and configure the start date and time. When applied, this filter displays log entries for events that were generated or received after or at the specified start time.
- To filter on a specific end time, select **To** and configure the end date and time. When applied, this filter displays log entries for events that were generated or received before or at the specified end time.
- To filter on a time period, select **From and To**, then enter the start and end date and time. When applied, this filter displays log entries for events that were generated or received within the specified time period.

In this example, you want to view all critical severity log entries that have a Follow-Up flag assigned to them. Additionally, you want to limit the search to log entries generated by the Engineering NSRP cluster on your network.

1. In the Log Viewer, right-click the **Severity** column header and select **Filter > Set Filter**. Select the checkbox for Critical, then click **OK** to save and apply the filter to your log entries.
2. Right-click the **User Flag** column header and select **Filter > Set Filter**. Enable Follow-Up, then click **OK** to save and apply the filter to your log entries.
3. Right-click the **Device** column header and select **Filter > Set Filter**. Select **Engineering Cluster**, then click **OK** to save and apply the filter to your log entries.

Filtering Log Entries by Range

A range filter is a criteria search for matching log entries within a value range. You can set a range filter for the following columns:

- Bytes In
- Bytes Out
- Bytes Total
- Packets In
- Packets Out
- Packets Total
- Src Port
- Dst Port
- Nat Src Port
- Nat Dst Port
- Elapsed Seconds

The following sections detail some common range filters used to manage log entries.

Setting a Bytes In or Bytes Out Range Filter

To view log entries based on the number of bytes received or transmitted during the event, set a range filter on the **Bytes In** or **Bytes Out** column:

1. Right-click the **Bytes In** or **Bytes Out** column header and select **Filter > Set Filter**. The **Bytes In/Bytes Out** filter dialog box appears.



NOTE: By default, the Log Viewer does not display the **Bytes In** or **Bytes Out** column. To set a byte filter, you must first configure the Log Viewer to display these columns. For details on configuring column settings, see [“Customizing Columns” on page 816](#).

2. Set the range for bytes received (Bytes In) or transmitted (Bytes Out):
 - To filter on a minimum number of bytes only, select **From** and enter a value. When applied, this filter displays log entries for events that received or transmitted more than or equal to the specified minimum number of bytes.
 - To filter on a maximum number of bytes only, select **To** and enter a value. When applied, this filter displays log entries for events that received or transmitted fewer than or equal to the specified maximum number of bytes.
 - To filter on a range of bytes, select **From and To**, then enter the minimum and maximum values for the range. When applied, this filter displays log entries for events that received or transmitted a number of bytes within the specified range.

3. Click **OK** to apply the filter.

Setting a Port Number Range Filter

To view log entries based on a range of port numbers used in the event, set a range filter on the Dst Port or Src Port column:

1. Right-click the **Src Port** or **Dst Port** column header and select **Filter > Set Filter**. The Dst/Src Port filter appears.
2. Set the range for the port numbers:
 - To filter on a minimum port number only, select **From** and enter a value. When applied, this filter displays log entries for events that used a port greater than or equal to the specified minimum port number.
 - To filter on a maximum port number only, select **To** and enter a value. When applied, this filter displays log entries for events that used a port less than or equal to the specified maximum port number.
 - To filter on a range of port numbers, select **From and To**, then enter the minimum and maximum values for the range. When applied, this filter displays log entries for events that used ports within the specified range.
3. Click **OK** to apply the filter.

Customizing Columns

You can configure the Log Viewer to display specific columns.

Using Column Settings

The Log Viewer includes 46 columns of log entry information; however, each predefined view includes only a small subset of available columns. To view information in the other available columns, or to change the column display order, you can adjust the column settings for the view.

The more columns you configure to appear in the Log Viewer, the more information you can see at one time—and the more you must scroll from side to side to view all columns; setting fewer columns means less viewable information, but also less scrolling. Typically, you use fewer columns when you already have enough detail about the event or you are only interested in specific event data.

Use column selection in combination with filters to create a customized view of your log entries.

Hide, Unhide, and Move Columns

You hide, unhide, or move columns to display specific information using one of the following methods:

- When managing columns using the **Column Settings** dialog box:
 - To display hidden columns, select the columns and click **Show**.

- To hide columns, select the columns and click **Hide**.
- To reorder the column display sequence, select a column and click **Move Up** or **Move Down**.
- When managing columns in the main display area:
 - To hide a column, right-click the column header and select **Hide Column**. To unhide a hidden column, you must use the **Column Settings** dialog box.
 - To reorder the column display sequence, drag a column to a new location.
 - To change column width, drag the left or right edge of the column header to the desired width.

In this example, you want to view the following information in the Log Viewer:

- The attacks that attempt to enter your network.
- The source IP and port of the attacking computers.
- The destination IP and port on the target computers.
- The date and time of the attacks.
- The devices that detected the attack.
- The policies that matched the attack.

First, you configure the Log Viewer to display only the columns that contain the information you are interested in viewing. Then you set can filters on those columns to narrow your search.

To configure the column settings:

1. In the navigation tree, select the Log Viewer module.
2. From the **View** menu, select **Choose Columns**. NSM displays the **Column Settings** dialog box, listing all columns.
3. Select the following columns:
 - Time Received
 - Src Addr
 - Dst Addr
 - Src Port
 - Dst Port
 - Category
 - Subcategory
 - Device
 - Policy

Ensure all other columns are not selected, then click **OK** to apply your changes to the Log Viewer.

4. In the main display area, select the **Src Port** column header, then drag the column to the right of the Src Addr column.

To configure the column filters:

1. In the main display area, right-click the **Category** column header and select **Filter > Set Filter**. The **Category** filter dialog box appears.
2. Select the following categories: **Predefined**, **Custom**, and **Screen**. Click **OK** to apply your changes. The Log Viewer applies the filter to the log entries.

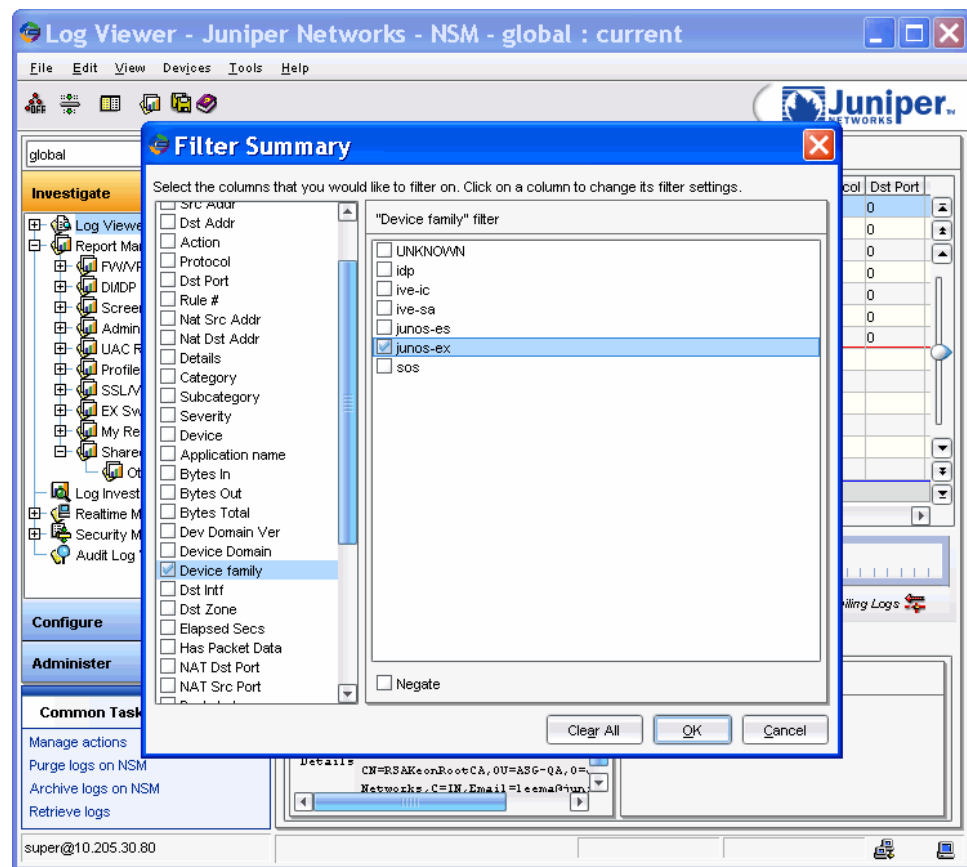
Filtering Log Entries by Column

Use the **Filter Summary** dialog box to select columns on which you want to filter log entries. To filter on one or more columns:

1. Select **View > Filter Summary**.
The **Filter Summary** dialog box is displayed.
2. From the **Filter Summary** dialog box, select a column on which you want to filter log entries.
3. Select the filter settings you wish to apply for the specified column, then click **OK**.
4. To select additional Log Viewer columns for filtering log entries, repeat Steps 1–3.
5. Before you exit NSM, save the Filter Summary changes that you made in Log Viewer.

[Figure 118 on page 819](#) shows the **Filter Summary** dialog box that you would use to configure filtering by Device family column.

Figure 118: Filter Summary Dialog Box



NSM applies the filter to all log entries and displays only the log entries that match the specified Log Viewer columns.

To clear a Log Viewer column that was selected for filtering log entries:

1. Select **View>Filter Summary**.

The **Filter Summary** dialog box is displayed.

2. • To clear a single column: Clear the column check box that you do not want to use for filtering log entries, then click **OK** .
- To remove all columns: Click the **Clear All** button.

Using Log Viewer Integration

The Log Viewer module is integrated with Security Policies and Device Manager modules. This integration enables you to jump from a log entry in the Log Viewer directly to the responsible security policy (jump to policy) or managed device (jump to device configuration).

[Jump to Policy](#)

To quickly edit a security policy rule from the Log Viewer, right-click a log entry and select **Goto Policy**. NSM opens a new UI window and displays the policy with the rule that generated the log entry.

- If the responsible rule exists within a rule group, the group is automatically expanded to reveal the rule.
- If the responsible rule exists within a VPN created by VPN Manager, the autogenerated rules appear.

Depending on the object version of the security policy, the rule might appear as read/write or read-only.

“Object version” refers to a specific modeled configuration; each time you install a modeled configuration (this includes security policies) on a managed device using NSM, the management system creates a new object version using the install date and time. NSM uses database snapshots to detect differences between the running configuration (installed on the physical device) and modeled configuration. Database snapshots also enable you to view previous object versions. For details on database snapshots, see [“Automatic Policy Versioning” on page 566](#).

Other options for archiving and restoring logs and configuration data are also available. For more information, refer to the *Network and Security Manager Installation Guide*.

When using the **Goto Policy** option in the Log Viewer, NSM compares the object version of the managed device to the current object version. If the responsible rule exists in a security policy that has the same object version as the security policy installed on the managed device, you can edit the rule.

However, if the responsible rule exists in a security policy that has a different object version from the security policy installed on the managed device, you cannot edit the rule. This typically occurs when you install a security policy on a managed device, then edit that policy in the NSM UI, but do not update the device with the new policy changes. Because the responsible rule exists in a policy that belongs to a previous object version, you cannot make changes to it.

[Jump to Device Configuration](#)

To quickly configure a parameter on an individual device from the Log Viewer, double-click a device in the Device column. NSM displays the device configuration for the device, enabling you to make changes to the device.

Identifying Irrelevant Attacks


Your log entries are a valuable tool in helping you identify irrelevant attacks. Irrelevant attacks are events that do not affect your network or that you do not consider important. Typically, you want to identify irrelevant attacks to:

- Reduce the number of log entries and increase system performance.
- Isolate log entries for harmless attacks.

- Focus on log entries for attacks to which you are actually vulnerable.

Select a log entry generated by a protocol anomaly or signature attack object, then view the Summary panel to see the attack description. An example is shown in [Figure 119 on page 821](#).

Figure 119: Viewing Summary Panel



Look carefully at the information about affected systems, and compare it with what you know about your network. Use the information in [Table 107 on page 821](#) to determine if the attack is relevant:

Table 107: Irrelevant Versus Relevant Attacks

Irrelevant Attacks	Relevant Attacks
<p>Attack target hardware you do not use.</p> <p>Example: Attacks that exploit Cisco routers do not affect Lucent routers.</p>	<p>Attack attempts to exploit vulnerabilities in the hardware you use in your network.</p>
<p>Attack target software you do not use.</p> <p>Example: Attacks that exploit Microsoft IIS Web servers do not affect Apache Web servers.</p>	<p>Attack attempts to exploit vulnerabilities in the software running on your network.</p>
<p>Attack target software versions you do not use.</p>	<p>Attack attempts to exploit vulnerabilities in the software versions running on your network.</p>

If the attack is irrelevant, you can remove the matching attack object group from the rule that triggered the log entry, or monitor the attack object group using custom severity setting.

Using the Log Investigator

The Log Investigator module enables you to investigate patterns and trends on your network using data gathered from your log entries. Log entries are generated by a security device when traffic matches a security policy rule, or when an event occurs that matches a predefined set of conditions. The Log Investigator uses the event data recorded in the log entry to identify the destination IP addresses and ports that are attacked most frequently, the services that are used to attack most frequently, and the source IP addresses that most frequently generate attacks.

When using NSM to manage large networks with multiple managed devices, you can potentially receive several hundred log entries in a single day (depending on how you have configured your devices for logging). The Log Investigator is a helpful tool for manipulating and correlating a large volume of log entry data so you can identify and analyze important activity that might threaten your network. By analyzing your data and then using that knowledge to proactively fine-tune your security policies, you can decrease risk while increasing security.

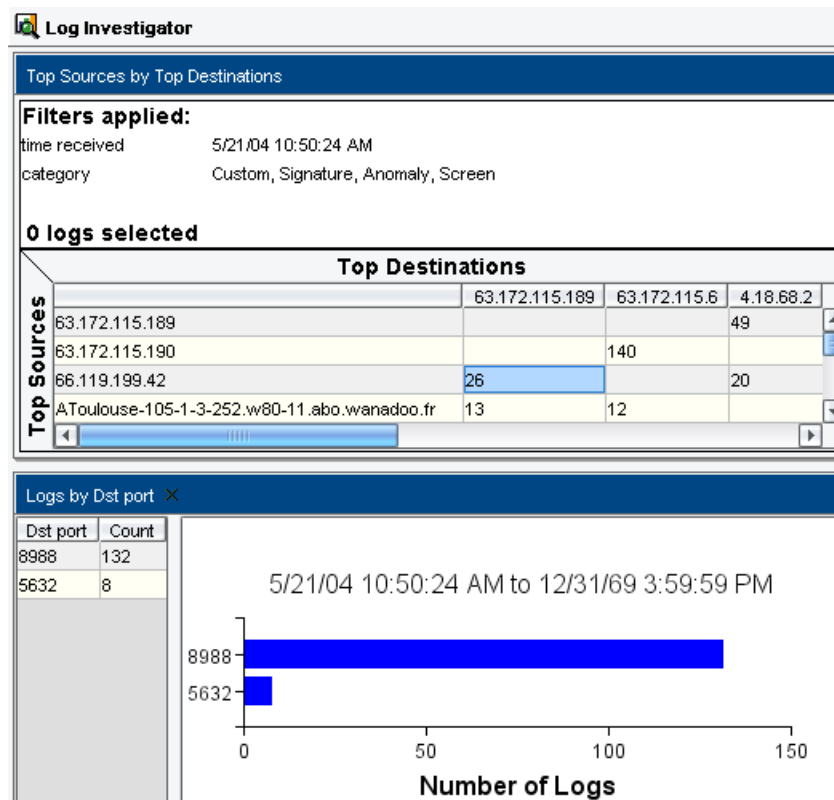
This section provides details on the following Log Investigator functionality:

- [“About the Log Investigator UI” on page 822](#)—The Log Investigator main display area includes a filter summary, a log entry matrix, and two detail panes that display detail information in table and chart format.
- [“Configuring Log Investigator Options” on page 824](#)—Configure the criteria the Log Investigator uses to create the matrix, including the time period, Left and Top Axes settings, the data point count (the number of data points the Log Investigator must collect before displaying data), and the maximum number of log entries you want the Log Investigator to use when collecting data.
- [“Setting Log Investigator Filters” on page 827](#)—As in the Log Viewer, you can set filters on log entry data so the Log Investigator displays only the information you want to see. Apply multiple filters to data for all log entry columns found in the Log Viewer.
- [“Investigating Log Entry Data” on page 829](#)—After you have configured the Log Investigator filters, time period, and data options, you are ready to begin investigating your log entry information. Within any cell the Log Investigator table, you can right-click and select an option to view specific details, including Destination Ports, Subcategories, and Time Period.
- [“Zoom Details” on page 831](#)—You can zoom in on specific details about activity between two data types, select a third data type for comparison, or display details about the event over time.
- [“Jumping to the Log Viewer” on page 832](#)—You can use the Log Viewer to see the corresponding log entries that are used in Log Investigator calculations.
- [“Excluding Data” on page 832](#)—You can configure the Log Investigator to exclude data for a cell, row, or column in the Log Investigator matrix.

About the Log Investigator UI

The main display of the Log Investigator is shown in [Figure 120 on page 823](#).

Figure 120: Log Investigator UI Overview



The Log Investigator contains the following UI components:

- **Filter Summary**—Displays the column filters currently applied to the Log Entry Matrix.
- **Selected Log Entries**—Displays the number of log entries currently selected in the Log Entry Matrix.
- **Left Axis**—The controlling axis for log entry data (the independent axis). The Log Investigator collects log entry data for the Left Axis setting, which determines data set that is used for Top Axis setting.
- **Top Axis**—The controlled axis for log entry data (the dependent axis). The Log Investigator collects log entry data for the Left Axis setting; for the Top Axis setting, the Log Investigator collects data that matches both the Left Axis and Top Axis setting.
- **Zoom Table**—Displays a table of log entry details. You can view Source, Destination, Destination Port, Attack Subcategories, or Time Period details for any cell, row, or column.
- **Zoom Chart**—Displays a chart of log entry details. You can view Source, Destination, Destination Port, Attack Subcategories, or Time Period details for any cell, row, or column.

Configuring Log Investigator Options

The first step in using the Log Investigator is to configure the basic criteria used to create the Log Investigator matrix. To change the default options, from the **View** menu, select **Set Log Investigator Options**. Use the **Log Investigator Options** dialog box to configure the desired settings (detailed below) and click **OK** to apply your changes.



NOTE: You can configure up to 20 Log Investigator sessions. To change this default number of sessions, edit the `devSvr.irMaxIndexCount` parameter in the `devSvr.cfg` file, which is located in the management system directory `/usr/netscreen/DevSvr/var/`.

The following sections detail each Log Investigator option.

Configuring a Time Period

The time period setting narrows the log entries included in your investigation based on a specified time interval or start time. Each log entry contains a timestamp that indicates the date and time the managed device generated the log entry (Time Generated). The Log Investigator compares the timestamp of a log entry to the specified time period setting, and eliminates those log entries that do not meet the time criteria.

First, you must specify a time duration. To specify a time interval for which you want to see log entries, set the number of weeks, days, hours, minutes, or seconds. Setting a longer interval time can help you identify broad trends in your network activity. Typically, you use a longer interface time to initially locate problems. After you have identified the issues you want to investigate, set a shorter time interval to eliminate irrelevant log entry data.

After you have determined the time interval, you must set the end or start time for the duration:

- To set the end time of the duration, select **Most Recent** (this is the default setting). The Log Investigator uses the current date and time as the end point for the time duration. For example, for a time interval of 5 hours, the Log Investigator collects data from log entries that have timestamps within the previous 5 hours.
- To set the start time of the duration, select **Start Time** and configure the start date and time. The Log Investigator uses the specified date and time as the start point for the time duration. For example, for a time interval of 5 hours and a start date of 5/12/04 8:00:00 AM, the Log Investigator collects data from log entries that have timestamps from the start date to the start date + 5 hours.

Typically, use **Most Recent** to investigate recurring activity or to monitor expected network changes. Use a start time when investigating past known events, such as a virus attack.

When using a large time interval, the number of matching log entries might exceed the capacity of the Log Investigator (100 log entries), causing a warning message to appear next to the Selected Logs indicator. If you do not make changes to the time interval filter,

the Log Investigator automatically clears the session, requiring you to create a new time filter.

For example, on Friday afternoon, you want to investigate attacks received by your network in the last seven hours. Configure the time period as shown in [Figure 121 on page 825](#).

Figure 121: Configure Time Period Filter

Time Period

Duration: 7 hours

☒ Most Recent

☐ Starting at: 5/21/04 12:00:00 AM

On Monday morning, you want to investigate attacks received by your network during the last work week. If Monday's data is 5/17/04, you configure the time period as show in [Figure 122 on page 825](#).

Figure 122: Changing Time Period Filter

Time Period

Duration: 5 days

☐ Most Recent

☒ Starting at: 5/10/04 12:00:00 AM

Configuring Axes

The Left Axis is the *independent* axis because it is the first data collected. The Top Axis is the dependent axis because it uses the Left Axis data as the data set.

The dependency occurs because the Log Investigator collects data that matches the Left Axis setting first; this data represents the data set for the entire log entry matrix. By default, the Left Axis is set to the data type Top Sources. After the Left Axis data set has been determined, the Log Investigator searches that data set for data that matches the Top Axis setting. By default, the Top Axis is set to the data type Top Destinations.

Because the Left Axis setting controls the initial data set, it is the most important axis setting. Typically, you should set the Left Axis to the data type you most want to investigate.

Setting the Data Type

You can change the data type for each axis. The data type defines the type of information that the Log Investigator attempts to locate in your log entries. For either axis, you can set the following data types:

- **Top Sources**—The IP address that generated the event.
- **Top Destinations**—The IP address that received the event.
- **Top Subcategories**—The attack subcategory detected in the event.
- **Top Destination Ports**—The port numbers on the Destination device that received the event. The port number can help you identify the service used in the event.

By default, the Left Axis uses the data type Top Sources and the Top Axis uses the data type Top Destinations. To change these settings, select the desired data type in the data point source menu.

Setting Data Points

A data point is a single data type field that matches the axis setting. By default, each axis collects 10 data points for each evaluation. These default settings create a Log Entry Matrix of 100 cells (the top 10 source IP addresses are correlated against the top 10 destination IP addresses, creating a 100-cell matrix). For example, a data point count of 6 for each axis would create a 36-cell matrix.

You can set the data point higher (maximum 40) or lower (minimum 5), depending on your investigation requirements. To change these settings, select the desired data type in the data point count field. The higher the data count, the larger the log entry matrix—and the more processing power required by the Log Investigator UI. Using large data counts can slow performance.

In this example, you swap the setting for the Left and Top Axes of the Log Investigator to see how each axis controls data.

Set the filter to Attacks, then configure the Left and Top Axes:

- To identify which of the most popular source addresses are generating attacks against the most popular destinations:
 - Select the Left Axis (the independent axis) as Top Sources.
 - Select the Top Axis (the dependant axis) as Top Destinations.

The Left Axis displays all attacks for the Top Source IP addresses, while the Top Axis displays the number of attacks for each of the Top Destinations attacked by the Top Sources.

- To identify which of the most popular destination addresses are receiving attacks from the most popular sources:
 - Select the Left Axis (the independent axis) as Top Destinations.
 - Select the Top Axis (the dependant axis) as Top Sources.

The Left Axis displays all attacks against the Top Destination IP addresses, while the Top Axis displays the number of attacks for each Top Source IP address that attacked a Top Destination.

Setting a Log Entry Limit

You can limit the number of log entries used in Log Investigator calculations. The NSM Device Server stores log entries from managed devices and the management system; when the GUI Server accesses a log entry to display its information in the UI, that log entry is placed in a log buffer. As the Log Investigator searches your log database for log entries that match the filter, time period, and data type criteria, it places all matching log entries in the log buffer.

To control the size of this buffer (the number of matching log entries), you can configure the Max Log Count for your investigations. The limit defines the number of matching log entries the Log Investigator accepts for its calculations.

You can set the following log entry limits:

- 100,000 log entries
- 200,000 log entries
- 400,000 log entries
- 600,000 log entries
- 800,000 log entries
- 1,000,000 log entries



NOTE: Setting a large buffer limit can degrade Log Investigator performance. The maximum buffer size of one million log entries uses all memory on the GUI Server and is not recommended.

Setting Log Investigator Filters

Log Investigator filters operate much like Log Viewer filters: You set criteria for log entries and the Log Investigator filters out log entries that do not match the filter criteria. Using the **Filter Summary** dialog box, you can select and apply multiple filters to the Log Investigator matrix.

To set filters, from the **View** menu, select **Set Filter**. [Table 108 on page 827](#) details filter types:

Table 108: Log Investigator Filters

Filter Type	Sample Filters	Description
Time Filter	Time Generated	Identifies packets by the time when a packet is sent from a device.
Address Filters	Src Addr Src Intf Dst Addr Dst Intf	Identifies packets based on information about an address of a device from which the packet was sent or an address of a device to which the packet was sent.
Direction Filters	Packets In Packets Out Packets Total	Identifies packets based on the direction they are heading to or from a specified device.

Table 108: Log Investigator Filters (continued)

Filter Type	Sample Filters	Description
Device Filters	Device Device Domain Ver Device family	Identifies device names, device versions, and device types.
Transmission Type Filter	Has Packet Data	Identifies transmissions if they are seen as packets.
Port Filters	Src Port Dst Port NAT Src Port NAT Dst Port	Identifies packets based on the port on a device from which they were transmitted or on the port on a device to which they were transmitted.
Policy Filters	Policy Rule # Policy ID	Identifies packets based on whether they meet the conditions of a policy or a rule.
Alarm Filters	User Flag Severity	Identifies the severity level of a generated alarm.
Miscellaneous Filters	Details Protocol Category Alert Roles User Application name	Various



NOTE: For a complete list of log entry columns available for filtering, see “Log Viewer Columns” on page 804.

After you have set a filter, the **Filter Summary** displays a list of all filters applied to the log entry data and the Log Investigator matrix displays values for matching log entries.

Example: Setting Filters in the Log Investigator

In this example, the Left Axis is set to Top Sources and the Top Axis is set to Top Destinations (these are the default settings). To set a filter that displays all attack category log entries generated by the Top Sources and received by the Top Destinations:

1. From the **View** menu, select **Set Filter** to display the **Filter Summary** dialog box.
2. In the filter list on the left, select **Category**, then select the following categories in the right: Predefined, Custom, and Screen.
3. Click **OK** to save and apply your changes.

To view the number of attacks between a specific source-destination pair, locate the Source Address 63.172.115.190 and Destination Address 63.172.115.6, then find the cell where the two addresses intersect. The Log Investigator displays 140 log entries for this Source-Destination pair, as shown in [Figure 123 on page 829](#).

Figure 123: View Log Investigator Results

Top Sources by Top Destinations

Filters applied:

time received5/21/04 10:50:24 AM

categoryCustom, Signature, Anomaly, Screen

0 logs selected

Top Sources

Top Destinations		
	63.172.115.189	63.172.115.6
63.172.115.189		
63.172.115.190		140
66.119.199.42	26	
A Toulouse-105-1-3-252.w80-11.abo.wanadoo.fr	13	12
blix.interl.net	26	
mail.wirthware.com	52	
164.164.94.253	13	14
modem-221.babbelas.dialup.pol.co.uk	13	
www.rno.org	26	5
cs6668187-121.austin.rr.com	26	

This high value (140) reflects the number of attack log entries that have occurred between these two IP addresses.

Investigating Log Entry Data

After you have configured the Log Investigator options and set filters as desired, you are ready to begin investigating your log entry data.

Using Rows and Columns

Each row or column in the Log Entry matrix represents events for a single data type. When selecting a row or column, you are evaluating how the data type (source, destination, subcategory, or destination port) for that axis relates to the other axis during a specific time period. Typically, reviewing a row or column in the matrix helps you analyze all events for a single data type.

For example, to investigate a sudden drop in performance on a specific destination, set the Left Axis to Top Sources and the Top Axis to Top Destinations, then select the column for the destination IP address. For each cell that displays a high number of events received by that destination, locate the corresponding source IP address. You might determine that destination 1 is receiving a large number of events from sources A, B, and C. This activity could be a harmless event, such as multiple users attempting to contact a single application server at the same time. You could eliminate the bottleneck by adding another application server to the network or restricting access to the existing server.

Using Cells

Each cell in the Log Entry matrix represents events that occur at the intersection of two data types. When selecting a cell, you are evaluating the events that occurred between those two specific data types (source, destination, subcategory, or destination port) during a specific time period. Typically, reviewing a cell in the matrix helps you analyze all events that occur between a data type pair.

For example, to investigate a sudden drop in network performance, set the Left Axis to Top Sources and the Top Axis to Top Destinations, then review the log entry matrix to locate a large number for a location pair. You might identify that source A is sending an unusually large number of transmissions to destination 1. This activity could be a harmless event, such as an employee archiving multiple large files before leaving work; however, this activity might be the result of a denial-of-service attack triggered by an internal trojan. You probably need to get more details, such as destination ports used and attack subcategories for the events before you can resolve the issue.

Table 109 on page 830 details the benefits of each type of Log Investigator analysis.

Table 109: Log Investigator Analysis

Data Type A (Left Axis)	Data Type B (Top Axis)	Benefit
Multiple Rows	Multiple Columns	View all network activity for specific data types. No cells or columns are selected (default view). Useful for analyzing events for multiple data types, such as multiple destinations and multiple sources. To focus on a specific data type pair, select the intersection cell.
Multiple Rows	One Column	View network activity for a single data type. A single column is selected. Useful for analyzing network performance issues, such as multiple sources generating traffic to a single destination.
One Row	Multiple Column	View network activity for a single data type. A single row is selected. Useful for analyzing attack traffic, such as one source generating traffic to multiple destinations.

Table 109: Log Investigator Analysis (continued)

Data Type A (Left Axis)	Data Type B (Top Axis)	Benefit
One Row	One Column	View specific activity between two specific data types. A single cell is selected. Useful for analyzing event traffic between two network components.

Zoom Details

You can zoom in on specific details about activity between two data types. You can select a third data type for comparison, or display details about the event over time. To get details, right-click a cell, row, or column and select **Zoom in** to see the list of available data types. Because the **Zoom in** menu is dynamic, it contains all data types not currently used for the Left or Top Axis of the Log Investigator matrix. Alternatively, you can select time as the third data type.

Details appear in the Zoom area, which contains two panes:

- The Zoom table (left pane)
- The Zoom chart (right pane)

The table and chart use the same information to generate values.

In the following example, the Left Axis is set to Top Sources and the Top Axis is set to Top Destination (these are the default settings); the filter is set to attacks (for details on setting the filter, see [“Example: Setting Filters in the Log Investigator” on page 829](#)).

To view the service ports on the destination device used by the attacks, right-click a cell that contains a nonzero value and select **Zoom In > Dst Port**. In the Zoom area:

- The left pane displays a table of service ports listed in descending order (the port accessed by the most attacks is listed first). The left column lists the Destination Port Number and the right column lists the number of attacks received by that port number. Because services are mapped to specific port numbers, you can use the port number to identify the service used in the attack.
- The right pane displays a chart using the same information.

In the following example, the Left Axis is set to Top Sources and the Top Axis is set to Top Destination (these are the default settings); the filter is set to attacks (for details on setting the filter, see [“Example: Setting Filters in the Log Investigator” on page 829](#)).

To view the individual attacks (the attack subcategories) against the destination device, right-click a cell that contains a nonzero value and select **Zoom In > Subcategory**. In the Zoom area, the left pane displays a table of attack subcategories listed in descending order (the attack found in the most number of log entries is listed first); the right pane displays a chart using the same information.

In the following example, the Left Axis is set to Top Sources and the Top Axis is set to Top Destination (these are the default settings); the filter is set to attacks (for details on setting the filter, see the example [“Example: Setting Filters in the Log Investigator” on page 829](#)).

To view the time period over which the attacks occurred, right-click a cell that has a nonzero value and select **Zoom In > Time**. In the Zoom area, the left pane displays a table of attacks listed in order (the oldest attack is listed first); the right pane displays a chart using the same information.

Jumping to the Log Viewer

The Log Investigator uses log entry data for calculations, and does not display the actual log entries. However, you can use the Log Viewer to see the log entries that are used in Log Investigator calculations.

To see corresponding log entries, right-click a cell, row, or column from the Log Investigator matrix or the Zoom table and select **View in Log Viewer**. A new UI window displays the log entries in the Log Viewer.

Excluding Data

You can manually configure the Log Investigator to exclude data for a cell, row, or column in the Log Investigator matrix. You might want to exclude:

- Irrelevant values (such as values from sources or destinations no longer in production)
- Abnormally high or low values (to establish a baseline)
- Specific data type (source, destination, destination port, subcategory)
- High values (when investigating events that generate lower values)

To exclude a specific attack from the Log Investigator calculations, right-click the attack cell and select **Exclude**. To help you keep track of excluded values, the **Filter Summary** area displays a list of values you have manually excluded.

Using the Audit Log Viewer

The Audit Log Viewer monitors administrative events that occur when a NSM administrator makes changes to a domain. Use the Audit Log Viewer to track changes to your managed device configurations. You can view audit-log entries for all managed devices in the all domains you have access to, or you can view entries for the devices in a single domain. When the disk space reaches the defined limits, old audit log entries are purged.

The Audit Log Viewer appears as one of the modules in the NSM UI. Select the Audit Log Viewer to display the audit log entry table, device view, and target view, as shown in [Figure 124 on page 833](#).

Figure 124: Audit Log Viewer UI Overview

Time Generated	Admin Name	Admin Login Domain	Command	Operation Outcome	Targets	Device Name
2/3/06 12:34:27 AM	super	global	modify	success	0/systemprefs/0	
2/3/06 12:34:27 AM	super	global	lock.systemprefs	success		
2/3/06 12:34:28 AM	super	global	sys_logout	failure		
2/6/06 11:19:56 AM	super	global	sys_login	success		
2/6/06 11:20:00 AM	super	global	select.sigpack	success		
2/6/06 11:20:00 AM	super	global	select.sigpack	success		
2/6/06 11:20:00 AM	super	global	firmwareList	success		
2/6/06 11:20:00 AM	super	global	firmwareList	success		
2/6/06 11:23:41 AM	super	global	select.sigpacksel	success		
2/6/06 11:24:38 AM	super	global	lock.systemprefs	success		
2/6/06 11:24:38 AM	super	global	modify	success	0/systemprefs/0	
2/6/06 11:24:39 AM	super	global	lock.systemprefs	success		

Target Name	Table	Domain ID
0	systemprefs	0

Device Name	Table	Domain ID



NOTE: Summaries appear in the Target Name or Device Name tables only for logs with targets or devices.

Audit Log Table

The audit log table contains the columns of information shown in [Table 110 on page 833](#).

Table 110: Audit Log Information

Column	Description
Time Generated	The time the object was changed. The Audit Log Viewer displays log entries in order of time generated by Greenwich Mean Time (GMT).
Admin Name	The name of the NSM administrator who changed the object.
Admin Login Domain	The name of the domain (global or subdomain) that contains the changed object.
Working Domain	<p>The domains from which a user is allowed to view audit logs. The values are:</p> <ul style="list-style-type: none"> empty— Audit-log entries created prior to this NSM release that do not have targeted objects or devices. These logs can be viewed by all NSM users. Subdomain— Audit-log entries for all managed devices associated with a subdomain. Users of the subdomain or global domain can view these audit-log entries. SYSTEM— Audit-log entries generated by the system. These audit-log entries can be viewed by all NSM users. global— Audit-log entries from the global domain. Users of the global domain can view all audit-log entries.
Command	The command applied to the object or system, for example, sys_logout or modify .
Authorization Status	The final access-control status of activities is either success or failure.

Table 110: Audit Log Information (continued)

Column	Description
Targets	<p>For changes made to a device configuration or object, the Audit Log Viewer displays the object type, an object name, and object domain.</p> <p>For changes made to a device, the Audit Log Viewer displays the device name, object type, and device domain.</p> <p>For changes made to the management system, such as administrator login or logout, the Audit Log Viewer does not display target or device data.</p>
Devices	<p>For changes made to a device configuration or object, the Audit Log Viewer displays the object type, an object name, and object domain.</p> <p>For changes made to a device, the Audit Log Viewer displays the device name, object type, and device domain.</p> <p>For changes made to the management system, such as administrator login or logout, the Audit Log Viewer does not display target or device data.</p>
Miscellaneous	Additional information that is not displayed in other audit log columns.

For changes to devices or device configurations, you can use the Target View or Device View to get more details about those changes.

Managing the Audit Log Table

NSM provides multiple ways to manage the data in your Audit Log table. The following sections describe these data-management options.

Select Audit Log Table

Use the **Set Audited Activities** option in the **Edit** menu to select read/write or read-only auditable activities. By default, all read/write activities are unchecked to avoid overloading the system's storage capacity. The settings for selected activities are saved in the GUI Server and are used to decide whether to create audit-log entry for an activity.

Sort Audit Log Table

Sort one column at a time by clicking on the Time Generated, Admin Name, Command, or Operation Outcome column. Hold down the Ctrl key and click on a column header to sort the selected column. Hold down the Shift key, and click on a column header to clear sorting on the selected column. The arrow icon on the table column header indicates descending or ascending order.

Filtering Audit Log Entries on One or More Columns

When you apply a filter to audit log entries, the Audit Log Viewer filters out log entries that do not match the filter criteria. You can set filters based on one or multiple columns. To select the columns on which you want to filter audit log entries:

1. Select **View >Set Filter**.
2. From the **Filter Summary** dialog box, select a column on which you want to filter log entries.



NOTE: For the Targets and Devices columns, you can specify filtering based on domain, category, or object, or the complete domain/category/object row value.

3. Select the filter settings you wish to apply for the specified column, then click **OK**.
4. To select additional Audit Log Viewer columns for filtering log entries, repeat Step 2 and Step 3.

The types of filters that you create are shown at the bottom of the Window next to **Filtered On**. Filters are added or removed immediately.

NSM applies the filter to all log entries and displays only the log entries that match the specified Audit Log Viewer columns.

To clear Audit Log Viewer columns that are selected for filtering log entries:

1. Select **View>Set Filter**.
2.
 - To clear a single column: Clear the column check box that you do not want to use for filtering log entries, then click **OK**.
 - To remove all columns, click **Clear All**.

Filtering Audit Log Entries On Column Field Values

You can set filters based on Audit Log Viewer column field values.

To set a column field filter, right-click a column field and select **Filter** to display the filter menu options:

- Time-based column filter— To create a time-based filter, right-click a field in the Time Generated column and select **Filter**. The following filter menu options are displayed:
 - Edit— Displays the **Filter** dialog box for the column, then you select the start and end time and dates you want to filter on.
 - To filter on a specific start time, select **From** and configure the start date and time. When applied, this filter displays log entries for events that were generated or received after or at the specified start time.
 - To filter on a specific end time, select **To** and configure the end date and time. When applied, this filter displays log entries for events that were generated or received before or at the specified end time.

- To display audit log entries for objects generated within the start and end date and time, click **OK**.
- **>= This Value**—Displays log entries for events that were generated at or after the time specified in the selected row cell.
- **<= This Value**—Displays log entries for events that were generated before or at the time specified in the selected row cell.
- **Clear Filter**—Removes the filter for the selected column.
- **Clear All Filters**—Removes the filter for all columns.
- Single or multiple value column filter— To create a filter based on a column field for the Admin Name, Admin Login Domain, Working Domain, Command, Authorization Status, Targets, Devices, or Miscellaneous column, right-click a field in the column and select **Filter**. The following filter menu options are displayed:
 - **Only This Value**—Displays only the content in the selected row cell.
 - **Not This Value**—Displays everything except the content in the selected row cell.
 - **Clear Filter**—Removes a current filter on the selected column.
 - **Clear All Filters**—Removes all filters for all columns.

Log Size and Data Migration

When the disk space usage reaches the defined threshold, the oldest audit log entries are purged before new entries are recorded. Configure the threshold in the **guiSvr.cfg** file.

You can migrate old audit log data into the latest version of the audit log.

Target View and Device View

The Audit Log Viewer also contains two detail views:

- **Target View**—For a change made to the device configuration, such as changing an IP address or renaming the device, select the audit log entry for that change in the Audit Log table, then view the Target View to see details about that change.
- **Device View**—For a change made to the device itself, such as adding the device, autodetecting a device, or rebooting a device, select the audit log entry for that change in the Audit Log table, then view the Device View to see details about that change.

To see additional details for an target or device audit log entry, double-click the entry in the Target View or Device View. For targets, NSM displays the configuration screen that the change was made in and marks the changed field with a solid green triangle. For devices, NSM displays the **Job Manager** information window for the job task.

Setting a Start Time for Audit Log Entries

By default, the Audit Log Viewer displays audit log entries in order of time generated by Greenwich Mean Time (GMT). To configure the Audit Log Viewer to display log entries for events that occurred after a specific time, configure the **Log By Time** option.

From the **View** menu, select **Go To Log By Time** to display the **Log By Time** dialog box. Select a date and time, then click **OK** to save and apply the time change to the Audit Log Viewer. The Audit Log table now displays only the audit log entries that were generated on or after the date and time you specified.

Managing Log Volume

Security administrators have different requirements for the number of log entries they need to retain. As directed by their corporate security policy, some administrators must keep all log entries, resulting in large numbers of log entries that the administrator might not have time to review, but needs to store.

To manage log volume, you can use the NSM UI to both archive and purge logs.



NOTE: Excessive logging creates additional traffic on your network. It is recommended that you balance your logging needs with the performance needs of your management system.

You can also export your log records to other formats for use in other applications. For details on how to forward logs, see [“Forwarding Logs” on page 840](#).

Automatic Device Log Cleanup

In NSM, logs are gathered and stored by the Device Server. In a given deployment, the Device Server can be deployed on the same machine as the GUI Server or on a separate machine. When determining disk space requirements for NSM, you must consider a log management strategy that optimizes the performance of your management system.

The `/usr/netscreen/DevSvr/var/devSvr.cfg` file contains log cleanup parameters that you can use to manage log disk space.

- **storageManager.alert**—If you configure this parameter, the Device Server triggers a warning when the available disk space falls below the configured value. The default value is 1500 MB. The user receives an e-mail alert about the low availability of free disk space on the Device Server.



NOTE: Use the Server Manager node in the NSM UI to configure e-mail notification. Refer to [“Configuring Servers” on page 745](#) for more information.

- **storageManager.minimumFreeSpace**— With this parameter you can configure a minimum free disk space value. The default is 1000 MB. If the available disk space falls below the configured value, the Device Server begins to purge logDb data starting from the earliest date in the logs directory until the available disk space reaches the configured value. However, if the free disk space is lower than the configured value of the **storageManager.Threshold** prior to the current day of log data, the Device Server shuts down automatically.

- **storageManager.threshold**—With this parameter, you can define a threshold for available disk space which if breached, causes the Device Server to automatically shut down. The default value is 800 MB.

If the Device Server fails to restart because of a lack of sufficient free disk space, an indicative error message appears in the console window advising you to either back up your data, or create additional disk space in order to restart the server.

You can change the parameters for managing disk space on the Device Server by editing the Device Server configuration file. For more information on managing disk space on the Device Server, refer to the *Network and Security Manager Installation Guide*.

Archiving Logs

You can archive and retrieve log entries to and from a storage device using standard Unix commands.

Logs reside on the Device Server in the following directory:

```
/usr/netscreen/DevSvr/var/logs
```

We recommend using the following commands to archive your logs:

- The **tar** command
- The **scp** (Secure Copy) command
- The **ftp** (File Transfer Protocol) command

For full descriptions and options for each command, see the man pages.



NOTE: You do not need to stop the processes on the Device Server before archiving.

Log Archival Mechanism

All managed device logs are stored in **/usr/netscreen/DevSvr/var/logs** that contains logs and associated files organized into subdirectories for each day. To archive the logs for a specified day, archive the entire directory for that day.

Each directory is named according to the YYYYMMDD convention, indicating the day contained in the directory. Do not attempt to archive the current day's files. You can automate archival using cron. To archive logs:

1. Use **scp** to copy all directories in **/usr/netscreen/DevSvr/var/logs/** to a remote archival location.
2. Remove the directories from the Device Server machine.

You can analyze the archived logs later by restoring them to the logs directory on the Device Server. The restored logs are then available in the **Log Viewer** and **Log Investigator** just as they were before archival.

1. Use scp to copy directories from the remote archival location to `/usr/netscreen/DevSvr/var/logs/`.
2. Analyze the logs using the NSM UI.
3. Remove directories when finished with analysis.

Setting Log Storage Limits

You can specify the number of days that NSM stores logs; as well as purge or archive a specified log based on the following configurable criteria:

- [“Date Limits” on page 839](#)
- System-Wide Retention Policy
- Obsolete Logs
- Required Disk Space

Date Limits

If you use a date to purge or archive logs, the limit is based only on the calendar date.

System-wide Retention Policy

The system administrator can specify the maximum number of days the system stores logs. One configuration throughout NSM is permitted at a time. Users can specify how the retention policy is triggered and when it is scheduled.

Obsolete Logs

As logs become obsolete, the user can archive before the system purges the logs. You have the option of purging the logs directly or archiving them first. If you select the archive option, NSM archives all the logs from the selected date.

Required Disk Space

After you define the number of logs and the number of days you want archived, NSM estimates the disk space required for storing the logs. In calculating the estimated required disk space, NSM uses the average size of logs per day and indicates to the user how the estimate was reached or if there was not data available to provide an estimate.

Archive Location

The location of the archive is user-configurable from the **Disk and Log Management** dialog box. The options are **Local** and **Remote**:

- **Local**—To archive logs locally, specify the directory location for file storage in the Archive Location field.

- **Remote**—To archive logs remotely, specify the IP address, username, password, and the protocol (scp and sftp). The path on the remote server is stored in the user's preferences. SCP and SFTP work only with trusted hosts.

File Name

One log archive location is applicable throughout the system. The following naming convention is used for storing log files:

```
YYYYMMDD_<No>.tar.gz
```

where **YYYYMMDD** is the date of the file containing the log and **No** is the archive number if multiple archivals exist for the same date.



NOTE: If the archive process fails; for instance, if the host is not preconfigured, or if there is not enough disk space, the user is notified via e-mail, and NSM creates an error log entry. However, the selected logs are purged even if archival fails.

Define Location

Before you archive a log, you must first define a location for the archived logs. To do this, open NSM and select **Server Manager > Servers > Device Server > Disk and Log management**. Enter the path in the **Archive Location** text box indicating where you want NSM to store the archived files.

Forwarding Logs

You can forward your log records for use in other applications using one of the following methods:

- **Action Manager**—Use the Action Manager, a node on the main UI, to configure the management system to forward logs generated within a specific domain or subdomain in NSM.
- **log2Action** utility—A command-line utility located on the NSM Device Server.



NOTE: You can also forward logs based on specific rules in a security policy. See [“Configuring Firewall Rules” on page 492](#) for more information.

Sending E-mail Notification of Downed Device

You can configure NSM to send you an e-mail notification when a device goes down so that even if you do not have access to NSM, you will be informed of device status.

1. In the main navigation tree, select **Action Manager > Action Parameters**, then double-click the row that lists all your action parameters.
2. Enter the default e-mail address in the EMail section for the 'From' e-mail address.
3. Click the **Add** icon to open the **New Add/Edit EMail Address** dialog box.
4. Enter the default 'To' e-mail address for all log actions in the current domain, then click **OK**. Repeat this step if additional default 'To' e-mail addresses are required.
5. In the main navigation tree, click **Action Manager > Device Log Action Criteria**, then click the **Add** icon.
6. Click the **Category** drop-down list box and select Info, select the **Device Disconnect** subcategory, then click **OK** to save the changes.
7. Click the **Actions** tab and check **SMTP Enable**.
8. Configure the target e-mail addresses for this rule, if they differ from the defaults you configured in the previous steps.

Using the Action Manager to Forward Logs by Domain

Use the **Action Manager** node to configure the management system to perform actions (such as syslog, export, or alarm) on log data based on the criteria you specify. These actions occur for all the managed devices in a specific domain or subdomain.

To enable the management system to export logs, you must configure the following:

- **Action Parameters**—These settings define the default log export settings for the management system, and determine how the system handles qualified log entries (log entries that match specified log criteria).
- **Device Log Action Criteria**—The criteria specifies the category and severity of the log entries you want to export. When a log entry meets the specified criteria, it is considered *qualified*, and NSM performs the specified actions defined in the criteria.

Configuring Action Parameters

From the **Action Manager**, select **Action Parameters** to define the default log export settings for the management system. To enable the management system to export qualified logs to the system log, SNMP, CSV, XML or e-mail, configure the export settings for each format as detailed in the following sections. For every log action criteria, you can specify and edit multiple system log and SNMP servers with their respective system log server facilities and SNMP communities.

Exporting to the System Log

For exporting to the system log, configure the IP address and the server facility for all of multiple syslog servers to which you want to send qualified logs. NSM uses the specified server when exporting qualified log entries to the system log.

To actually export logs to a system log server, you must select "Syslog Enable" using the **Actions** tab in the **Device Log Action Criteria** node.

Exporting to SNMP

For exporting to SNMP, configure the following SNMP settings:

- **SNMP Manager**—Specify the IP addresses of the SNMP servers to which the GUI Server sends SNMP traps.
- **SNMP Community**—Specify SNMP community names that provide the desired combination of both read and write access from the SNMP server.

NSM uses this information when exporting qualified log entries to SNMP. This setting defines the SNMP settings for the management system. To actually export logs to the specified SNMP servers and community, you must select “SNMP Enable” using the **Actions** tab in the **Device Log Action Criteria** node.



NOTE: The NSM MIB files resides in the `/usr/netscreen/DvrSvr/Utils` directory on the management system. The file names are `jnx.nsm-traps.mib` and `jnx-smi.mib`.

Exporting to CSV

For exporting to CSV, configure the following CSV settings:

- **File Path**—The directory and filename that you want log entries exported to in .CSV format.
- **Print Header**—When selected, column headers are exported to .CSV format.

NSM uses this information when exporting qualified log entries to CSV. This setting defines the CSV settings for the management system. To export logs to CSV, you must select “CSV Enable” from the **Actions** tab in the **Device Log Action Criteria** node.

Exporting to XML

For exporting to XML, configure the directory and filename to which you want to send qualified logs in XML format. NSM uses this information when exporting qualified log entries to XML; each log becomes an XML record, which you can open in most Web browsers.

This setting defines the XML settings for the management system. To actually export logs to XML, you must select “XML Enable” from the **Actions** tab in the **Device Log Action Criteria** node.

Exporting to E-mail

For exporting to e-mail, configure the following e-mail and SMTP settings:

- **SMTP Enable**—Enables the SMTP server to send e-mail.
- **Default ‘From’ Email Address**—Some servers require a valid “from” e-mail address to relay mail.
- **Default ‘To’ Email Addresses**—The e-mail address that receives e-mail alarms. You can specify multiple “to” e-mail addresses.

NSM uses the preceding information when exporting qualified log entries to e-mail. These settings define the e-mail and SMTP settings for the management system.



NOTE: After editing your e-mail settings, you must restart the Device Server for your changes to take effect

Setting Device Log Action Criteria

A **Device Log Action Criteria** instance defines the criteria for a qualified log; each instance contains two criteria settings (category and severity), and multiple action settings for logs that meet the criteria settings. For example, to only export critical severity attack logs to XML, you create a device log action criteria instance that specifies the log category as predefined, the severity as critical, and the action as XML. For each log entry that matches the criteria, NSM exports the log as XML, using the default XML settings configured in the **Actions** tab.

To add a new **Device Log Action Criteria** instance, use the **Add** button, then configure the following settings.

Selecting Category

In the Category list, select a category of log entry for the criteria. Some categories contain subcategories; however, to set an action based on a subcategory, you must first select a category.

For details on each category and subcategory, see [“Log Entries” on page 919](#).

Selecting Severity

In the **Severity** tab, select the severities for the criteria.

Selecting Actions

In the **Actions** tab, select the actions (SNMP, syslog, XML, CSV, Email, and Script) you want the management system to take for logs that meet the criteria. You can enable multiple actions.

When you enable the Email/SMTP and Script actions, you can also configure the following additional settings:

- **Email Action**—To direct the management system to e-mail qualified log records, specify the From and To e-mail addresses:
 - **From Email Address**—The e-mail address that the server uses to send e-mail. Some servers require a valid “from” e-mail address to relay mail.
 - **To Email Addresses**—The e-mail address that receives e-mail alarms. You can specify multiple “to” e-mail addresses.
- **Script Action**—To direct the management system to send qualified log records to a script, you must configure the following:
 - **Script Enable**

- **Script To Run**—Select the script you want to run from the Script To Run list. For a script to appear in the list, the script must be located in the appropriate directory on the NSM Device Server.

Scripts for the global domain must appear in the `/usr/netscreen/DevSvr/var/scripts/global/` directory.

Scripts for subdomains must appear in `/usr/netscreen/DevSvr/var/scripts/global/subDomainName/` where *subDomainName* is the name of the subdomain. The *subDomainName* directory must be created manually.

The `/usr/netscreen/DevSvr/lib/scripts/` directory contains two sample scripts you can use or modify, `sample.sh` and `sample.pl`.

- **Action Upon Script Failure**—Specify the error handling for the script:
 - **Skip**. Directs the system to skip any log for which the script had an error.
 - **Retry**. Directs the system to try the action again for the same log. When using this filter, you must also specify:
 - **Retry Count**. Specifies the maximum number of retries to attempt before moving on to the next log record.
 - **Retry Interval** (in seconds). Specifies the number of seconds until the action is tried again.

Using the log2action Utility to Export Logs

You can also use a command line utility on the Device Server to export logs. To export to XML, CSV, SNMP, Syslog, e-mail, or script format using the **log2action** utility:

1. Log in to the NSM Device Server as root.
2. Change to the utility directory by typing: `cd /usr/netscreen/DevSvr/utlils`.
3. Specify the common filters, format, and format-specific filters for the format you want to export to:

```
sh devSvrCli.sh --log2action<common_filters>--action <format><format_options>
```

The **log2action** utility exports all log records to the specified format. After executing the action, the system generates an exit status code of 0 (no errors) or 1 (errors).

The following sections detail common filters, actions, and required and optional format-specific filters.

Using Filters

The **log2action** utility generates data for a maximum of 100,000 logs.



NOTE: If you want to generate more than 100,000 logs, use the `matches-to-return` option to specify the number of logs that you want.

Because of the large volume of logs potentially generated, it is highly recommended that you specify filtering criteria when using the **log2action** utility. Without filtering, the action report generates data from the earliest date in the log database and stops providing output after 100,000 logs. In this case, it is possible that you may not get the action output of your most recent data. Specifying a time filter is recommended in this situation.

Using Time Filters

For example, if you wanted to view data in the logs of 20060317, run the following command:

```
./devSvrCli.sh --log2action --filter --log-id 20060317:0-20060317:4294967294 --action --xml
--file-path /tmp/newtest.xml
```

If you wanted to view data for all logs from 2006/03/15 to 2006/03/17, run the following command:

```
./devSvrCli.sh --log2action --filter --log-id 20060315:0-20060317:4294967294 --action --xml
--file-path /tmp/newtest.xml
```

Using Common Filters

To control which log records are exported, use common filters. Common filters are optional and must be used *before* the action command (**-action**).

Table 111 on page 845 shows the common filters.

Table 111: Common Filters

Option	Default	Multiple	Specifies	Format
--category	yes	yes	Category	<category> Specify one or more of the following values: admin, alarm, config, custom, events, implicit, info, predefined, profiler, screen, self, sensors, traffic, urlfiltering, user.
--device	yes	yes	Device name	<domain-path>:<device-name>
--device-family	yes	yes	Device type	<device family> idp, ive-ic, ive-sa, j/SRX Series, EX Series, m/MX Series, sos
--domain	yes	yes	Domain path	<global[/<subdomain-name>]
--dst-ip	yes	yes	Destination IP address	<a.b.c.d[/n]-<a.b.c.d>]>
--dst-port	yes	yes	Destination port	<[0-65535][-[0-65535]]>
--log-id	yes	no	From Log ID To Log ID	<<yyyymmdd>:[0-MAX][-<yyyymmdd>:[0-MAX]]>
--matches-to-return	yes	no	Number of log entries to match	<[1-4294967295]>

Table 111: Common Filters (continued)

Option	Default	Multiple	Specifies	Format
--rule	yes	no	Rule to match	<domain-path>: <policy-name>:<rulebase>:<rule number> where <rulebase> is one of the following values: fw, idp, honeypot, backdoor, synpro, vpn, mpolicy, tsig.
--severity	yes	yes	Severity	<severity> Specify one of the following values: none, info, device_warning_log, minor, major, device_critical_log, emergency, alert, critical, error, warning, notice, informational, or debug.
--src-ip	yes	yes	Source IP address	<a.b.c.d[/n]-<a.b.c.d>]>
--src-port	yes	yes	Source port	<[0-65535][-[0-65535]]>
--time-recv	yes	yes	Time received	<<yyyymmdd>:<hhmmss>>-<<yyyymmdd>:<hhmmss>>
--user-flag	yes	yes	User flag number	<[0-7]> High = 0 Medium = 1 Low = 2 Closed = 3 False Positive = 4 Assigned = 5 Investigate = 6 Follow-Up = 7 Pending = 8
--action	yes	no	Action, usually followed by format-specify filters	<action> (csv, e-mail, script, snmp, syslog, xml)

Some common filters support multiple entries, enabling you to specify more than one criteria. When using multiple entries for a common filter, you must use the common filter before each entry.



NOTE: Use the help command to print all relevant filter options:

```
sh devSvrCli.sh --help
```


Example: Specifying a Common Filter with Multiple Entries

To set a filter that displays all log entries for IDP and EX Series devices, type:

```
./devSvrCli.sh --log2action --filter --device-family idp,junos-ex --action --csv --file-path /tmp/moon.csv --include-header no
```

In the following example, a common filter option is specified to view logs based on the category option. To set a filter that displays log entries that indicate an implicit rule was matched or a configuration change occurred on the device, type:

```
./devSvrCli.sh --log2action --filter --category implicit,config --action --csv --file-path /tmp/sun.csv --include-header no
```



NOTE: When a filter option includes multiple entries, use a comma-separated list with no space between the entries, as shown in the preceding example.

Using Format-Specific Filters

To control how log records are exported, use format-specific filters. Some formats have required and optional format-specific filters. Use format-specific filters *after* the specified action.



NOTE: To see all format-specific filters for a format, type:

```
sh devSvrCli.sh --log2action --action -- format
```

Exporting to XML

The xml action directs the system to output logs using the XML format. To export:

1. Login to the Device Server as root, then change to the utility directory by typing: **cd /usr/netscreen/DevSvr/utls.**
2. To export to a file, type:

```
sh devSvrCli.sh --log2action --action --xml <file-path> <include-header>
```

The Device Server exports all log records to XML; each log record becomes an XML record, which you can open in most Web browsers.

Using XML Required and Optional Format-Specific Filters

You can use the following required and optional format-specific filters for exporting to XML:

CSV	Multiple	Required	Meaning
--file-path	No	Yes	Specifies where the system should direct the output. For example, myLogs.xml
--include-header	No	No	Specifies that the system should print the field name before each field.

Viewing XML Format Output

To view the XML schema file, type:

```
/usr/netscreen/DevSvr/lib/logActions/log.xsd
```

To export predefined and custom attack category log records to an XML file located in the **/usr** directory of the Device Server, use the **--category** common filter to specify the categories:

```
sh devSvrCli.sh --log2action --category predefined --category custom --action --xml --file-path /usr/MyXmlLogRecords/attacks.xml
```

Exporting to CSV

The csv action directs the system to output logs using the CSV format. To export:

1. Login to the Device Server as root, then change to the utility directory by typing: **cd /usr/netscreen/DevSvr/utlis.**
2. To export to a file, type:

```
sh devSvrCli.sh --log2action --action --csv <file-path> <include-header>
```

The Device Server exports all log records to CSV; each log record becomes an CSV record.

Using CSV Required and Optional Format-Specific Filters

You can use the following required and optional format-specific filters for exporting to CSV:

CSV	Multiple	Required	Meaning
--file-path	No	Yes	Specifies where the system should direct the output. For example, myLogs.csv

CSV	Multiple	Required	Meaning
--include-header	No	No	Specifies that the system should print the field name before each field.

Viewing CSV Format Output

CSV log files use this format:

```
Log Day Id, Log Record Id, Time Received (UTC), Time Generated (UTC), Device
Domain, Device Domain Version, Device Name, Device IpAddr, Category, Sub-Category,
Src Zone, Src Intf, Src Addr, Src Port, NAT Src Addr, NAT Src Port, Dst Zone,
Dst Intf, Dst Addr, Dst Port, NAT Dst Addr, NAT Dst Port, Protocol, Policy Domain,
Policy Domain Version, Policy, Rulebase, Rule Number, Policy ID, Action, Severity,
Is Alert, Details, User, App, URI, Elapsed Secs, Bytes In, Bytes Out, Bytes
Total, Packets In, Packets Out, Packets Total, Repeat Count, Has Packet Data, Var
Data Enum, Application name, Device family.
```

To print the column headers for log records when exporting to a CSV file, use the include-header option:

```
sh devSvrCli.sh --log2action --action --csv --include-header
```

```
sh devSvrCli.sh --log2action --action --csv --include-header --file-path
/usr/MyCSVLogRecords/logrecords.csv
```

Exporting to SNMP

The snmp action directs the system to output logs to an SNMP server in SNMP format. You must specify the SNMP community string and the SNMP server IP address that receives the exported log records.

To export:

1. Login to the Device Server as root, then change to the utility directory by typing: **cd /usr/netscreen/DevSvr/utlis.**
2. To export to a file, type:

```
sh devSvrCli.sh --log2action --action --snmp <community> <server>
```

The Device Server exports all log records to the specified SNMP community and server.

Using SNMP Required and Optional Format-Specific Filters

You can use the following required format-specific filters for exporting to SNMP:

SNMP	Multiple	Required	Meaning
--community	No	Yes	Specify SNMP community string. The community is an arbitrary string that the SNMP server is configured to recognize. For details on the community parameter, refer to section 3.2.5 of RFC 1098. You might need to ask your SNMP server administrator for the server community string.
--server	No	Yes	Specify SNMP manager IP address The value must be encoded as [IP FQDN:<port>]

The SNMP format has no optional format-specific filters.

Viewing SNMP Format Output

SNMP trap log entries use the following format:

```
<day id> <record id> <time received> <time generated> <device domain> <device
domain version> <device> <device ip> <category> <subcategory> <source zone> <source
interface> <source ip> <source port> <nat src ip> <nat src port> <destination
zone> <destination interface> <destination ip> <destination port> <nat dst ip>
<nat dst port> <protocol> <rule domain> <rule domain version> <policy> <rulebase>
<rulenum> <action> <severity> <isalert> <details> <user str> <application
str> <uri str> <elapsed secs> <bytes in> <bytes out> <bytes total> <packets in>
<packets out> <packets total> <repeat count> <has packet data> <var data enum>
<application name> <device family> <policy id> <var data>
```

To send log records to the public SNMP server at 192.168.1.15, use the --public and --server options:

```
sh devSvrCli.sh --log2action --action --snmp --community public --server 192.168.1.15
```

Exporting to E-mail

The e-mail action directs the system to output logs for an e-mail address in SMTP format. You must specify the recipient's e-mail address the exported log records, and you have the option of specifying the sender's email address.

To export:

1. Login to the Device Server as root, then change to the utility directory by typing: **cd /usr/netscreen/DevSvr/utlis.**
2. To export to a file, type:

```
sh devSvrCli.sh --log2action --action --email <sender> <recipient>
```

The Device Server exports all log records to the specified e-mail address for the recipient.



NOTE: You do not specify the SMTP server IP address in the `log2action` utility. The system uses the IP address configured for e-mail in the Log Actions area of the GUI Server (in the NSM UI). To configure this IP address, select **Server Manager > Servers** in the Administer panel of the main configuration tree. Then double-click the server name on the GUI Server panel and enter the SMTP Server IP address, and the From Email Address and To Email Address in the Email Notification tab. For details on configuring this value, see [“Exporting to E-mail” on page 842](#). You must configure the IP address before attempting to export logs to an e-mail address.

Using E-mail Required and Optional Format-Specific Filters

You can use the following required and optional format-specific filters for exporting to e-mail:

E-mail/SMTP	Multiple	Required	Meaning
<code>--recipient</code>	Yes	Yes	Specify the receiving e-mail address for the SMTP log records
<code>--sender</code>	No	No	Specify the sender e-mail address

Exporting to syslog

The syslog action directs the system to output logs to a syslog server in syslog format. You must specify the IP address of the syslog server that receives the exported log records and the syslog facility.

To export:

1. Login to the Device Server as root, then change to the utility directory by typing: `cd /usr/netscreen/DevSvr/utls.`
2. To export to a file, type:

```
sh devSvrCli.sh --log2action --action --syslog <server> <facility>
```

The Device Server exports all log records to the specified IP address for the syslog server.

Using Syslog Required and Optional Format-Specific Filters

You can use the following required format-specific filters for exporting to syslog:

Syslog	Multiple	Required	Meaning
--server	No	Yes	Specify syslog server IP address as [IP FQDN[:<port>]]. Examples: <ul style="list-style-type: none"> 192.168.1.25:7889 syslog.server@mycompany.com:7889
--facility	Yes	Yes	Specifies the facility that receives syslog messages. For details on the facility parameter, refer to section 4.1.1 of RFC 3164. The syslog severity, also used to calculate the overall syslog message priority, is automatically set to alert.

The syslog format has no optional format-specific filters.

Viewing Syslog Format Output

Syslog messages use the following format:

```
<day id>, <record id>, <timeReceived>, <timeGenerated>, <devicedomain>,
<devicedomainVersion>, <deviceName>, <deviceIpAddress>, <category>, <subcategory>,
<src zone>, <src intface>, <src addr>, <src port>, <nat src addr>, <nat src
port>, <dst zone>, <dst intface>, <dst addr>, <dst port>, <nat dst addr>, <nat
dst port>, <protocol>, <rule domain>, <rule domainVersion>, <policyname>,
<rulebase>, <rule number>, <policy id>, <action>, <severity>, <is alert>,
<details>, <user str>, <application str>, <uri str>, <elapsed>, <bytes in>, <bytes
out>, <bytes total>, <packet in>, < packet out>, < packet total>, <repeatCount>,
<hasPacketData>, <varData Enum>, <application name>, <device family>
```

Exporting to a Script

The script action directs the system to execute a script, use STDIN to pass log records formatted as XML to the script, and report output status. You must specify the name of the script that receives the exported log records (script must be located in the /usr/netscreen/DevSvr/lib/scripts/ directory).

To export:

1. Login to the Device Server as root, then change to the utility directory by typing: **cd /usr/netscreen/DevSvr/lib**.
2. To export to a file, type:

```
sh devSvrCli.sh --log2action --action --script <script-name> <error-handling>
```

The Device Server exports all log records to the specified script.

Using Script Required and Optional Format-Specific Filters

You can use the following required format-specific filters for exporting to a script:

Script	Multiple	Required	Meaning
<code>--script-name</code>	No	Yes	<p>Specify the script name. The script must be located in <code>/usr/netscreen/DevSrv/var/scripts/<domain>/<script-name></code></p> <p>For example:</p> <p><code>/usr/netscreen/DevSrv/var/scripts/global/<script-name></code></p> <p>or</p> <p><code>/usr/netscreen/DevSrv/var/scripts/global/<subdomain>/<script-name></code></p>
<code>--error-handling</code>	No	Yes	<p>Specifies error handling for the specified script. When using this filter, you must specify one of the following error-handling filters:</p> <ul style="list-style-type: none"> • <code>--skip</code> Directs the system to skip any log for which the script had an error. • <code>--retry</code> Directs the system to try the action again for the same log. When using this filter, you must also specify: • <code>--retry-interval</code> Specifies the number of seconds until the action is tried again. • <code>--num-retries</code> Specifies the maximum number of retries to attempt before moving on to the next log record.

The script format has no optional format-specific filters.

CHAPTER 20

Reporting

Use the Report Manager module in Network and Security Manager to generate and view reports summarizing log and alarms generated by the managed Juniper Networks devices in your network. You can use these reports to track and analyze log incidents, network traffic, and potential attacks.

This chapter contains the following sections:

- [About Reporting on page 855](#)
- [Report Types on page 857](#)
- [Working with Reports on page 862](#)
- [Setting Report Options on page 870](#)
- [Log Viewer Integration on page 872](#)
- [Using Reports on page 873](#)
- [Using the Watch List on page 877](#)

About Reporting

The Report Manager module in NSM is a powerful and easy-to-use tool that enables you to generate reports summarizing key log and alarm data originating from the managed devices in your network. The reports in Report Manager provide a useful complement to the monitoring and logging capabilities in NSM, enabling you to track and analyze network traffic, activities, and potential attacks.

Report Manager contains the following benefits for generating reports:

- [Report Type Groupings on page 855](#)
- [Graphical Data Representation on page 856](#)
- [Integration with Logs on page 856](#)
- [Central Access to Management Information on page 856](#)

Report Type Groupings

The reports in Report Manager are grouped together according to the type of data they provide:

- FW/VPN—Reports that summarize log and alarm data generated by the managed security devices in your network.
- DI/IDP—Reports that provide data on deep inspection (DI) and intrusion detection and prevention (IDP) attacks.
- Screen—Reports that provide data on Screen attacks detected by the firmware on the managed security devices in your network.
- Administrative—Reports specifically designed to help system administrators track and manage log incidents and security rules.
- UAC Reports—Reports that provide data on Unified Access Control (UAC) sessions.
- Profiler Reports—Reports that help system administrators investigate and analyze potential problems in the network and to resolve security incidents.
- AVT—Reports that help system administrators track the volume of application traffic in the managed network.
- EX Switch Report—A report that provides configuration data on EX Series switches.
- My Reports—All reports that you have saved or created as custom reports.
- SSL/VPN Reports—Reports that provide data on Infranet Controller.
- Shared Reports—All reports that you have saved or created that you want made accessible to others in a domain.

Grouping these reports by type enables administrators and operations staff interested in tracking and analyzing specific types of information to work only within the group of reports that they need.

For details on each of the specific reports per group, see [“Report Types” on page 857](#). For additional details on each report type group, refer to the *Network and Security Manager Online Help*.

Graphical Data Representation

You can use reports to view log data in both tabular and graphical form. The various depictions of the data make it easier to identify trends and potential areas of risk. You can also choose to view the data in either a horizontal bar graph or a pie chart.

Integration with Logs

Reports are also integrated with the Log Viewer and Log Investigator modules. By clicking a data point depicted in a report, you can quickly drill down to access and view the specific log entries presented in the report data. Refer to [“Log Viewer Integration” on page 872](#) for more information about how you can use reports and log entries together to further analyze network events and attacks.

Central Access to Management Information

For network administrators and security analysts interested in tracking and identifying potential network trends and attacks, Report Manager provides a single graphical view into the network.

Report Types

Report Manager contains a set of predefined reports that you can use out of the box. We recommend that you use the predefined reports. First to familiarize yourself with how reporting works in NSM. You can then fine tune these reports by generating custom reports based on the predefined reports.

Predefined Reports

The predefined reports in Report Manager provide a summary of key log events and alarms generated by the devices in your network (such as Top Scan Sources or Top Attacks). Two reports (Logs By User-set Flag and Top Rules) provide administrative information useful if you are tracking incidents or optimizing your rules. For typical use cases describing each of these reports, see [“Using Reports” on page 873](#).

Report Manager groups predefined reports into the following categories:

- [Firewall/VPN Reports on page 857](#)
- [DI/IDP Reports on page 858](#)
- [Screen Reports on page 859](#)
- [Administrative Reports on page 860](#)
- [UAC Reports on page 860](#)
- [Profiler Reports on page 861](#)
- [AVT Reports on page 861](#)
- [SSL/VPN Reports on page 862](#)
- [EX Series Switches Report on page 862](#)

Firewall/VPN Reports

[Table 112 on page 857](#) lists and describes reports in NSM that provide information related to your network’s firewalls and VPNs.

Table 112: Firewall and VPN Reports

Report	Description
Top Alarms	The total number of alarms generated by the managed security devices in your network, excluding traffic alarms
Top Traffic Alarms	The total number of traffic alarms generated by the managed security devices in your network
Top Traffic Log	The total number of traffic log entries generated by the managed security devices in your network, within filter constraints
Top Configuration Logs	The total number of configuration log entries generated by the managed security devices in your network, within filter constraints

Table 112: Firewall and VPN Reports (continued)

Report	Description
Top Information Logs	The total number of information log entries generated by the managed security devices in your network, within filter constraints
Top Self Logs	The total number of Self log entries generated by the managed security devices in your network, within filter constraints

DI/IDP Reports

Table 113 on page 858 lists and describes reports in NSM that provide DI and IDP information.

Table 113: DI/IDP Reports

Report	Description
Top 100 Attacks (last 24 hours)	Those attacks that are detected most frequently within the last 24 hours.
Top 100 Attacks Prevented (last 24 hours)	Those attacks that are prevented most frequently within the last 24 hours.
Top 20 Attackers (All Attacks - last 24 hours)	20 IP addresses that have most frequently been the source of an attack during the last 24 hours.
Top 20 Attackers Prevented (All Attacks - last 24 hours)	20 IP addresses that have most frequently been prevented from attacking the network during the last 24 hours.
Top 20 Targets (last 24 hours)	20 IP addresses that have most frequently been the target of an attack during the last 24 hours.
Top 20 Targets Prevented (last 24 hours)	20 IP addresses that have most frequently prevented attacks during the last 24 hours.
All Attacks by Severity (last 24 hours)	Number of attacks by severity level (set in attack objects).
All Attacks Prevented by Severity (last 24 hours)	Number of attacks by severity level (set in attack objects).
All Attacks Over Time (last 7 days)	All attacks detected during the last 7 days.
All Attacks Prevented Over Time (last 7 days)	All attacks prevented during the last 7 days.
All Attacks Over Time (last 30 days)	All attacks detected during the last 30 days.
All Attacks Prevented Over Time (last 30 days)	All attacks prevented during the last 30 days.

Table 113: DI/IDP Reports (continued)

Report	Description
Critical Attacks (last 24 hours)	All attacks categorized as “critical” detected during the past 24 hours.
Critical Attacks Prevented (last 24 hours)	All attacks categorized as “critical” prevented during the past 24 hours.
Critical Thru Medium Attacks (last 24 hours)	All attacks categorized as either “critical” or “medium” detected during the past 24 hours.
Critical Thru Medium Attacks Prevented (last 24 hours)	All attacks categorized as either “critical” or “medium” prevented during the past 24 hours.
Top 50 Scan Sources (last 7 days)	50 IP addresses that have most frequently performed a scan of a managed device.
Top 50 Scan Targets (last 7 days)	50 IP addresses that have most frequently been the target of a scan over the last 7 days.
Profiler - New Hosts (last 7 days)	New Hosts listed in the Profiler over the last 7 days.
Profiler - New Ports (last 7 days)	New Ports listed in the Profiler over the last 7 days.
Profiler - New Protocols (last 7 days)	New Protocols listed in the Profiler over the last 7 days.
Top IDP Rules	The total number of log entries generated by specific rules in your IDP policies. You can use the Top Rules report to identify those rules that are generating the most log events. This enables you to better optimize your rulebases by identifying those rules that are most and least effective. You can then modify or remove those rules from your security policies.

Screen Reports

When the firmware on your device identifies an attack, it generates a log event. These events are totaled and summarized for your review in the reports shown in [Table 114 on page 859](#).

Table 114: Screen Reports

Report	Description
Top Screen Attacks	The most common attacks detected by the firmware on your security device
Screen Attacks by Severity	The number of attacks detected by the firmware on your security device according to severity level

Table 114: Screen Reports (continued)

Report	Description
Screen Attacks over Time	A summary of when attacks are detected by the firmware on your security device
Top Screen Attackers	Where attacks originate from most frequently
Top Screen Targets	Which hosts on your network are the most frequent targets of attackers for firewall attacks

Administrative Reports

[Table 115 on page 860](#) lists and describes reports in NSM that provide information specifically for administrators.

Table 115: Administrative Reports

Report	Description
Logs by User-set Flag	The total number of log entries that were flagged by an administrator in the Log Viewer according to the predefined flag type set. You can flag log events as High, Medium, Low, Closed, False Positive, Assigned, Investigate, Follow-Up, or Pending. You can use the Logs by User-set Flag report to quickly identify log events of specific interest.
Top FW/VPN Rules	The total number of log entries generated by specific rules in your ScreenOS/DI policies. You can use the Top Rules report to identify those rules that are generating the most log events. This enables you to better optimize your rulebases by identifying those rules that are most and least effective. You can then modify or remove those rules from your security policies.

UAC Reports

[Table 116 on page 860](#) lists and describes those reports in NSM that provide information about Unified Access Control (UAC) session logs.

Table 116: UAC Reports

Report	Description
Time graph of UAC session logs	Total number of UAC session logs over the last 7 days.
Top 20 Destinations	20 destination IP addresses that have most frequently appeared on UAC logs over the last 7 days.
Top 20 Enforcers (devices) for UAC logs	20 Infranet enforcer devices that have most frequently appeared on UAC logs over the last 7 days.
Top 10 auth failures for user@realm	Ten user authentication failures that have mostly frequently appeared on UAC logs over the last 24 hours.

For more specific information describing each report, refer to the *Network and Security Manager Online Help*.

Profiler Reports

[Table 117 on page 861](#) lists and describes those reports in NSM that provide information about Profiler session logs.

Table 117: Profiler Reports

Report	Description
Top 10 Peers by Count	Ten source and destination IP addresses that appeared most frequently in the Profiler logs.
Top 10 Peers with maximum hits	Ten source and destination IP addresses that had highest number of hits in the Profiler logs.

For more specific information describing each report, refer to the *Network and Security Manager Online Help*.

AVT Reports

[Table 118 on page 861](#) lists and describes those reports in NSM that provide information about application volume tracking.

Table 118: AVT Reports

Report	Description
Top 10 Applications by Volume	Ten applications with highest volume in bytes in the past 24 hours.
Top 10 Application Categories by Volume	Ten application categories with highest volume in bytes in the past 24 hours.
Top 5 Applications by Volume over Time (last 1 hour)	Five applications with highest volume in bytes in the past hour.
Top 5 Application Categories by Volume over Time (last 1 hour)	Five application categories with highest volume in bytes in the past hour.
Top 5 Source by Volume over Time (last 1 hour)	Five source IP addresses with the highest volume in bytes in the past hour.
Top 5 Destination by Volume over Time (last 1 hour)	Five destination IP addresses with the highest volume in bytes in the past hour.

SSL/VPN Reports

Table 119 on page 862 lists and describes those reports in NSM that provide information about SSL/VPN session logs.

Table 119: SSL/VPN Reports

Report	Description
Top 10 active users based on total bytes	Ten users with highest number of Bytes Out and Bytes Out in SSL/VPN logs.
Top 10 auth failures for user@realm	Ten users with highest number of authentication failures in SSL/VPN logs during the last 24 hours.

For more specific information describing each report, refer to the *Network and Security Manager Online Help*.

EX Series Switches Report

Table 120 on page 862 lists and describes a report in NSM that provides information about EX Series session logs.

Table 120: EX-Switch Reports

Report	Description
Top Configuration changes (last 1 week)	The ten EX Series switches with the highest number of configuration changes during the last seven days.

For more specific information describing each report, refer to the *Network and Security Manager Online Help*.

My Reports

Once you are comfortable using reports, you can create your own custom reports to provide the exact information that your network security needs require. My Reports are associated with a specific user across domains.

Shared Reports

You can also allow others to use your custom reports by creating them as a shared report. Shared Reports are associated with domains. Subject to user-defined access control settings, shared reports are available to all other users in the domain.

Working with Reports

You can use the Report Manager to perform the following actions:

- [Generating a Predefined Report on page 863](#)
- [Creating a Custom Report on page 863](#)

- [Deleting Reports on page 864](#)
- [Organizing Reports in Folders on page 864](#)
- [Generating Reports Automatically on page 864](#)
- [Exporting Reports to HTML on page 869](#)

Generating a Predefined Report

To generate a predefined report, click that report from the Report Manager. The report is generated according to its default report settings. You can set different report options to further tailor your reports to your specific needs. See ["Setting Report Options" on page 870](#) for more information.

Creating a Custom Report

Both system administrators and read-only administrators can create custom reports based on their own reporting requirements.

You can also use an existing predefined report as the basis for a custom report by generating that report and saving it as custom report. You can use the same process to copy reports.



NOTE: All System Administrators, including those assigned a Read-Only role, can create and run their own reports.

Example: Creating a Custom Report

Suppose, for example, that you are a security administrator responsible for monitoring and protecting the corporate DMZ network. A Top Attacks report comes predefined in IDP, but the report displays attacks on the entire network, and you are interested only in the DMZ.

To create a custom report based on a predefined report:

1. In the objects component, configure a network object called Corporate DMZ Network and add all the IP addresses located in the DMZ.
2. Using the **Top Attacks** report, use **Save As** to rename the report "Top DMZ Attacks". You can also click the **Save As** icon on the toolbar or use the Ctrl-S keyboard shortcut.
3. In **Columns for Report**, select **Destination Address**.
4. In the **Log Filter** tab, select to filter on only those destination addresses defined in the Corporate DMZ network object.
5. In the **General** tab, under **Other Parameters** use the **Save Report In** field to select **My Reports** and then edit Others to "My DMZ Reports".
6. Click **OK**.

NSM creates the new report and displays it in a new folder called My DMZ Reports folder under My Reports.



NOTE: You cannot create a subfolder under the first level of custom report folders.

Deleting Reports

If a custom report no longer serves your informational needs, you can delete it. You cannot delete predefined reports.

To delete a custom report, select the custom report you want to delete and select **Delete** from the **File** menu. You can also right-click the report and select **Delete**, or use the Ctrl-D keyboard shortcut.

Organizing Reports in Folders

If you have a large number of custom reports, it is good practice to organize them in folders. You can organize your reports in folders under the My Reports and Shared Reports.



NOTE: You cannot create folders under any of the Predefined Reports folders.

To create a report folder, you need the appropriate permissions —Create Catalog Object, Delete Catalog Object, and Edit Catalog Object.

After you have created a report folder, you can later edit or delete it. You cannot, rename a folder. You can achieve the same result by saving the same report under a new folder name, and then deleting the previous folder. For more information about editing and deleting a report folder, refer to the *Network and Security Manager Online Help*.

Generating Reports Automatically

You can generate scheduled log-based reports automatically by using the **guiSvrCli.sh** command line utility located on the NSM GUI Server. The utility works with scripts that enable you to generate and send the reports using e-mail or FTP.



NOTE: You cannot use **guiSvrCli.sh** to generate reports defined as a My Report.

To create scheduled log-based reports, perform the following steps:

1. Run the report using the **guiSvrCli.sh** utility to provide the login information and the report name to be created.
2. Create or edit the action script to define what is done with the report after it is run. Sample scripts are included to show how to send reports as e-mail and as FTP data.
3. Set up a cron job to execute the **guiSvrCli.sh** utility with the proper parameters.



NOTE: The preceding configuration must be done from the GUI Server console, not the UI. You can verify the status of an executed report in the Job Manager.

Running Reports Using the guiSvrCLI.sh Utility

The `guiSvrCli.sh` utility is located in the `/usr/netscreen/GuiSvr/utls` directory on the GUI Server. Use the following syntax to generate reports:

```
export NSMUSER=global/<user name>; export NSMPASSWD=<password>;
/usr/netscreen/GuiSvr/utls/guiSvrCli.sh --generate-reports --report global:<Report
Folder>:"report-name" --script ftp.sh
```

The `export NSMUSER` and `NSMPASSWORD` statements set the username and password used to generate the report. This user must have appropriate rights on the NSM management system. You need to create a specific account for this purpose.

Syntax to Generate a Report for Critical Attacks for the Last 24 hours

Use the following syntax to generate report named "Critical Attacks Last 24 hrs" under the shared report category.

```
export NSMUSER=global/<user name>; export NSMPASSWD=<password>;
/usr/netscreen/GuiSvr/utls/guiSvrCli.sh --generate-reports --report global:shared:"Critical
Attacks Last 24 hrs" --script ftp.sh
```

Syntax to Generate Report for the Major Attacks on the Device the Previous Day

Use the following syntax to generate a report named "Major Attacks Yesterday" under the shared report category for the subdomain IDP.

```
export NSMUSER=global/<user name>; export NSMPASSWD=<password>;
/usr/netscreen/GuiSvr/utls/guiSvrCli.sh --generate-reports --report global:shared:"Critical
Attacks Last 24 hrs" --script ftp.sh
```

Syntax to Generate a Report for All Major Attacks

Use the following syntax to generate a report named "ALL_MAJOR" under the shared report category for the subdomain IDP.

```
export NSMUSER=global/<user name>; export NSMPASSWD=<password>;
/usr/netscreen/GuiSvr/utls/guiSvrCli.sh --generate-reports --report global:shared:ALL_MAJOR
--script ftp.sh
```

You can generate any of the predefined reports by specifying "system" in the `<Report Folder>` field, or any user-created report other than those defined in "My Reports" by specifying "shared" in the `<Report Folder>` field. You must reference the report with the full domain path, using a colon to separate domain elements and the report folder.

The script parameter refers to the script, located in the `/usr/netscreen/GuiSvr/var/scripts` directory, that is run on completion of the report generation.

Creating and Editing Action Scripts



NOTE: Sample scripts enabling you to e-mail and FTP the report results are available in `/usr/netscreen/GuiSvr/lib/scripts` for your convenience. To use these scripts, we recommend that you first copy them to `/usr/netscreen/GuiSvr/var/scripts`, and then change the permissions on the scripts so that they are both writable and executable. You can then customize the scripts to your needs.

Transferring Reports to an FTP Server

The `ftp.sh` script is used to transfer the report to an FTP server. The following portion of the script needs to be edited:

```
#####
#      CONFIGURABLE PARAMETERS
#####
# Remote hostname or IP address
remote_host="localhost"
# login for ftp account
userid="ftp"
# password for ftp account
passwd="ftp"
# pick reports from this directory prefix
local_dir_prefix="/usr/netscreen/GuiSvr/var"
```

E-mailing Reports

To e-mail reports, you must configure two scripts:

- **email.sh**—This script is called in the `guiSvrCli.sh` utility and defines how the reports are to be included in the e-mail message
- **Email.pl**—This script is called by `email.sh` and configures the actual SMTP parameters.

You can attach or embed the report in the e-mail by uncommenting a specific line in `email.sh`. You can also deliver multiple reports in separate mail messages or in a single collated one.

```
#####
#      CODE
#####
dir=`dirname $0`
cd $dir
# Each report in $1 is mailed as an attachment.
/usr/bin/perl -w email.pl -d $1 -o /dev/null
# Each report in $1 is embedded into an email and mailed.
/usr/bin/perl -w email.pl -d $1 -embed -o /dev/null
# Each report in $1 is attached to a collated message and the collated message
# is mailed
/usr/bin/perl -w email.pl -d $1 -collate -o /dev/null
# Each report in $1 is embedded into a collated message and the collated message
```

```
# is mailed
#/usr/bin/perl -w email.pl -d $1 -collate -embed -o /dev/null
# Don't change this line
exit $?
The required SMTP settings are provided in email.pl.
#####
# CONFIGURABLE PARAMETERS SECTION (CHANGEABLE)
#####
# From address, don't delete \ before the @ separator
my $from_addr = "user\@localhost.localdomain";
# To address, don't delete \ before the @ separator
my $to_addr = "remote-user\@somewhere.net";
# Email server: not required if sendmail is configured for mail transport
my $email_server = "everywhere.net";
# Subject
my $subject = "Reports are here!";
# Body text for emails with reports as attachments
my $body_text_attach = "Attached reports";
# Body text for emails with reports embedded
my $body_text_embed = "Embedded reports";
# Mail transport agent
my $send_mail_prog = "/usr/sbin/sendmail -t";
# Directory prefix for report directory
my $prefix = "/usr/netscreen/GuiSvr/var";
# Report extension type
my $type = "html";
# Mail output file: Capture sent email in this file, /dev/stdout for screen
my $mail_ofile = "/dev/stdout";
#####
```



NOTE: The `email.pl` and `emailReports.sh` scripts only do MIME formatting of the reports. The actual mailing is done by `sendmail` or other Mail Transfer Agent (MTA). If you require authentication to send these reports, you must configure the authentication parameters as part of the MTA configuration. If you are using the FTP script to send you reports, you will also need to add values for the remote host, userid and password for the FTP account in the `ftp.sh` file.

Using Cron with Scheduled Reports

The actual scheduling is done using the `cron` application. This utility executes scripts at specific times. It is configured through a file called **crontab**. To edit the file, use the command **crontab -e**. This command invokes the `vi` editor and opens the crontab table.

Entries in the table consist of a command set and a schedule. The command to run the report is the same as described above.

The timing of the job is determined by a string of numbers preceding the script. There are five places and they represent, in order:

- Minute (0-59)
- Hour (0-23)

- Day of Month (1-31)
- Month (1-12)
- Day of Week (0-6) (Sunday = 0)

Use an asterisk(*) to mark a place that is not part of the given schedule. For example, to run a script every Tuesday night at 11:05 PM, use "5 23 * * 2 Script".

In this example, perform the following steps to generate a predefined report and FTP it to a server every Monday at 12:01 in the morning:

1. Change to the utility directory by typing **cd /usr/netscreen/GuiSvr/utls**
2. Create a shell script called **reportscript.sh**:
 - a. Set the NSMUSER environment variable with an NSM domain/user pair, for example:

export NSMUSER=domain/user
 - b. Set the NSMPASSWD environment variable with an NSM password. The command for setting environment variables depends on your operating system and shell, for example:

export NSMPASSWD=password
 - c. Specify a **guiSvrCli** command string

**/usr/netscreen/GuiSvr/utls/guiSvrCli.sh --generate-reports --report global:system:
\"report-name\" ' --script ftp.sh**
3. Make the script executable. Make sure the person who creates the cron job can run the script.

4. Run the crontab editor:

```
crontab -e
```

5. Add the following line:

```
0 0 * * 1 /usr/netscreen/GuiSvr/utls/reportscript.sh
```

Running Xdb Audit Log Exporter Tools with High Availability

The audit log contains a log entry for all the changes in NSM administrator. Using the **xdbAuditLogExporter.sh** tool you can export the logs into CSV and syslog formats.



NOTE: The Audit Log contains a log entry for all the changes in NSM administrator. You can export this data into CSV and syslog format types. Using `xdbAuditLogExport.sh` tool you can export the logs into csv and syslog formats.

To export audit log in HA environment:

1. Use SSH or Telnet to connect to your NSM server and log in.
2. Change to the utils directory by entering:
`cd /usr/netscreen/GuiSvr/utils.`
3. Run `guiSvrCli.sh` tool to export log in CSV or syslog format.

Followings are some example to export the data

- To export auditlog to csv file run the following command:
`$ sh guiSvrCli.sh --export_audit_log --action --csv audit.log`
- To send auditlog to specific syslog server run the following command:
`$ sh guiSvrCli.sh --export_audit_log --action --syslogs 10.205.1.202`
- To send audit log with different filter option to csv file run the following command:
`$ sh guiSvrCli.sh --export_audit_log --action --csv test1.csv --filter --domain global`
`$sh guiSvrCli.sh --export_audit_log --action --csv test2.csv --filter --action-field sys_login`
`$sh guiSvrCli.sh --export_audit_log --action --csv test4.csv --filter --admin abc`
`$ sh guiSvrCli.sh --export_audit_log --action --csv test2.csv --filter --time 01/09/2012::12:30:00-01/09/2012::16:00:00`



NOTE: The audit log CSV files are always stored in /tmp location.

Exporting Reports to HTML

After you create your reports, you can export them into HTML. For example, if you want to share information with other security experts about the attacks that you are noticing in your network, use the following process to export the report onto disk:

1. Select **Export Reports** from the **File** menu. Alternatively, you could right-click in the chart window, and use the “Export reports in HTML” option.
2. Select the **Top Attacks** report check box.

3. Click **Browse** to save the file onto CD or to any other location on your desktop.
4. Click **Export**. The report is exported onto your CD.

NSM saves the report in several file formats (such as .png, .html, and .gif) that you can later display in any Web browser.

Setting Report Options

By default, each report in NSM provides information based on data available from the current day in a horizontal bar chart. You can configure the duration, number of data points, and appearance of each report by using the **Set Report Options** selection in the **View** menu.



NOTE: You can also access the **Set Report Options** dialog by right-clicking the chart on each report.

Use the **Set Report Options** selection in the **View** menu to tailor your reports to display only the specific information that you want. You can configure the following options in each report:

- Report title
- Report type
- Columns for the report
- Time period
- Data point count
- Chart type

You can also access the **Set Report Options** dialog by right-clicking the chart on each report.

Naming a Report

You can enter a name for a new report or rename an existing report in the **General** tab of the **Set Report Options** window. You can also configure the name of the report displayed in the report graph by editing its title.

Setting the Report Type

You can create two types of reports:

- **Time-Based**—Displays activity over time. For example, the Attacks Over Time report is a time-based report that measures the top attacks recorded in log records over a specified period.
- **Count-Based**—Displays total current activity to date. For example, the Top Scan Targets report is a count-based report that displays the total number of scans currently recorded against a specified number of destination IP addresses.

Configuring Report Source Data

You can configure a report to one of the following log record columns: Action, Alert, Src Addr, or Policy. Select the report source data by checking the appropriate check boxes in the **Columns For Report** selection area. The data that you choose for the columns in your report appear in the Y-axis of the graph.

Configuring a Report Time Period

You can configure a report to display all available data from either a specific date and time or during a specific time interval.

For example, if you suspect that your network was attacked on September 15 at 6:00 PM, you could set the Starting At Time Period Duration report field in the options on a Top Screen Attacks report to that time, then generate the report.

If you are not sure of the exact date or time of the attack, but know it occurred during the past 2 days, set the Duration field in the Time Period Duration report options on a Top Screen Attacks report to two days, then generate the report.



NOTE: The data that you can display in each report is limited by the amount of log information available.

Configuring the Data Point Count

Typically, the top 50 occurrences of each data type are displayed in each report. You can configure a report to display more or fewer data points depending upon the level of detail you need. For example, if you want to obtain a more precise view of the top occurrences of events, you would configure a lower data point count (such as 25).



NOTE: The minimum data point count that you can configure in all reports is 5; the maximum data point count is 200.

Configuring the Chart Type

By default, each report depicts information in a horizontal bar chart. You can also configure the report to depict information using a pie chart.

Sharing Your Custom Report

Use the **Save Report In** pull-down menu and select the **Shared Reports** option to specify that you want to share your report across all domains.

Modifying Report Filters

You can also use report filters to reduce the amount of unwanted or unnecessary log information compiled in each report. This makes it easier for you to focus on only the log

data of interest to you. You can specify criteria to filter your log data on any of the columns that you have chosen to base the report.

For example, you are a security administrator that typically reviews the “Attacks by Severity” report. You notice that critical attacks are on the rise. To track this more closely, you can modify the log filter on the “Attacks by Severity” report so that the report only displays critical attacks. To do this, select the “Attacks by Severity” report, and use **Set Report Options** to access the **Log Filter** tab. In the **Log Filter** tab, select to filter on attacks, and unselect all attacks except for those that are critical.

Configuring Report Processing Warnings

Each time you generate a report, it must perform a scan operation on a certain set of log records in the log database. The total number of log records that a report operation requires can have an adverse impact on your overall management performance. To prevent extraordinarily lengthy report operations from impacting your overall system performance, you can use the **Preferences** tool to configure NSM to display a warning message before a report is to scan a certain threshold number of log records.



NOTE: This setting also applies to the Dashboard and Log Viewer.

For example, set a warning message threshold at around 1,000,000 logs. To do this, use the **Preferences** option in the **Tools** menu and select **Reports**. In the **New Preference Settings** dialog box, click in the **Enable Warnings** check box and use the up and down arrows to specify 1,000,000 as the number of **Maximum Records to Filter**.

After this preference is applied, a warning appears each time a report is set to perform an operation requiring 1,000,000 log records to be scanned.

Saving Your Report Settings

After you have defined your custom reports, you can save the report settings as a custom report. Saved reports are organized under the tree node named “Custom Reports”.

Log Viewer Integration

Report Manager uses log data as the basis of all the information presented in each report. Because of this, we recommend that you consider requirements for reporting as you decide how many log entries you want to maintain and store.

Viewing Logs From Report Manager

One key benefit of Report Manager’s tight integration with log entries is the ability to quickly access the source log data presented in each report. To view the source log entries in the Log Viewer for more detailed information about the report data, right-click a data point in any report and select **Log Viewer** from the **View** menu. The source log entries will appear in the Log Viewer.

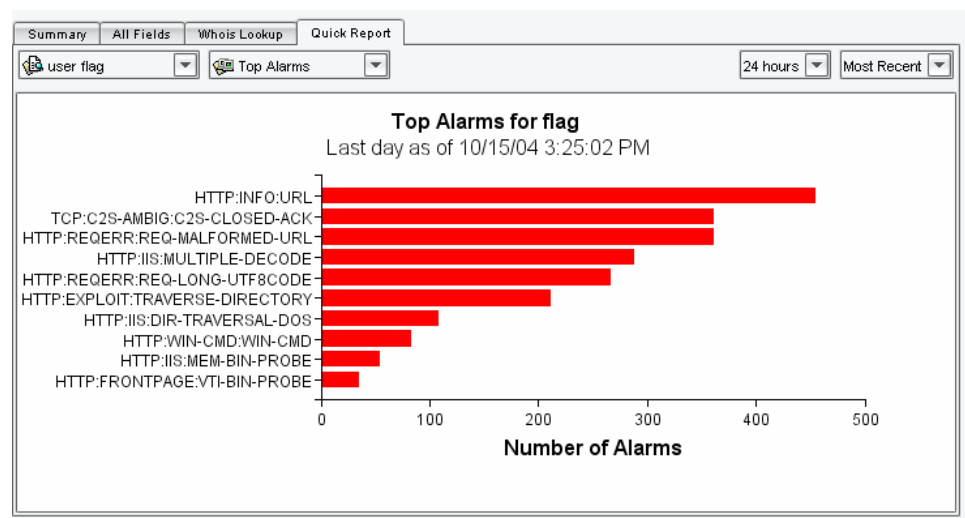


NOTE: You cannot save the view generated in the Log Viewer for use in a later UI session.

Generating Quick Reports

You can generate a Quick Report from data displayed in the Log Viewer or Log Investigator. Use the **Quick Report** tab located at the bottom of the Log Viewer or Log Investigator module to display a count-based custom report called a Quick Report, as shown in Figure 125 on page 873.

Figure 125: Generating A Quick Report



From the Quick Report screen, you can further set report options using the pull-down menus provided to define the report. You can then save the report as a custom report.

Using Reports

The following examples describe typical use cases for the reports in NSM.

Example: Using Administrative Reports to Track Incidents

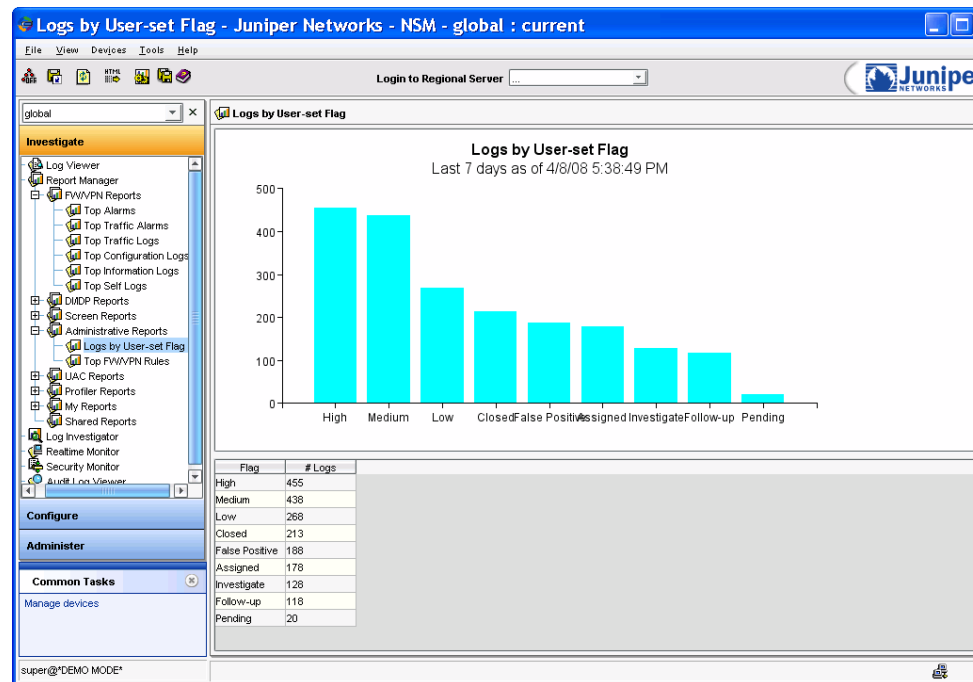
In this example, firewall administrators use the Log Viewer to monitor and investigate log events. They are specifically interested in configuration changes that are causing outages sporadically throughout the network. When they encounter a configuration log that seems out of the ordinary, they flag the log by using the predefined flag type "Investigate". To flag a log entry, right-click on the log and select **Flag > Investigate** from the drop down menu.

After completing their investigation, they change the flag to either "Closed" or "Assigned" for further investigation. During normal operations, firewall administrators investigate over 200 log entries per day.

You are a network manager interested in the progress of the investigation. To help track the progress, you generate a "Logs by User-set Flag" report, as shown in [Figure 126 on page 874](#).

By setting the duration of the report to one week, you can determine the total number of log entries flagged for investigation, total closed, and total assigned for further analysis.

Figure 126: Logs by User-Set Flag Report

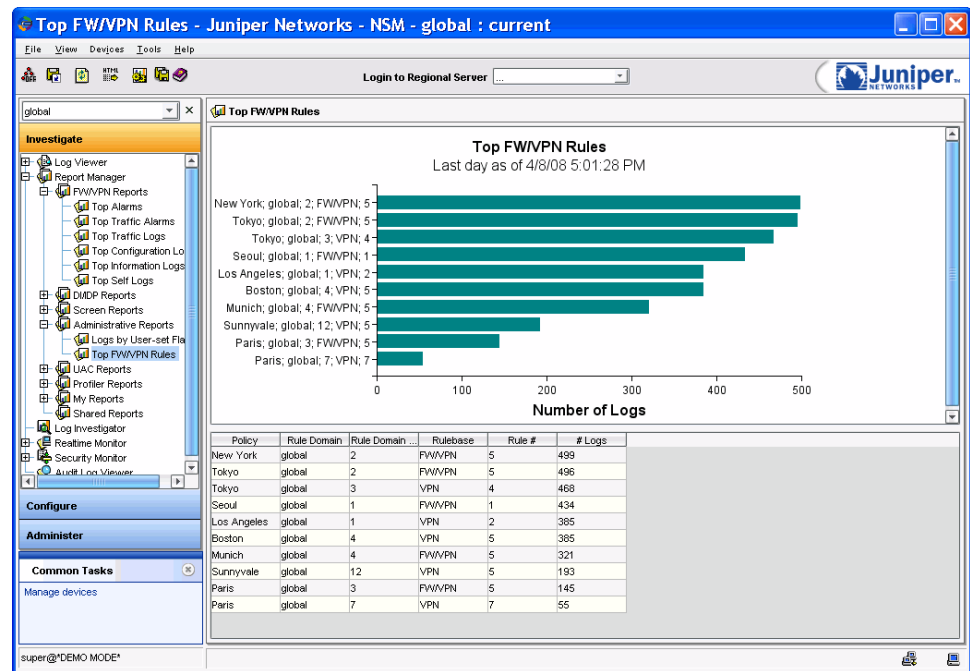


Example: Using Administrative Reports to Optimize Rulebases

In this example, you are a security administrator responsible for implementing new rules to your firewall rulebase. After you have updated the new security policy on the managed security devices in your network, you want to know the effect of the new rules on network traffic.

You configure a "Top FW/VPN Rules" report to start at the same date and time that the new rulebase settings were updated in the network. You also set the report data point count to 100. In this way, you can get an indication for the top 100 rules that are generating log events. [Figure 127 on page 875](#) shows the Top FW/VPN Rules report.

Figure 127: Top FW/VPN Rules Report



By identifying the new rules that you implemented in the network, you can track how effective the new rules are. If you find that a specific rule that is permitting too much traffic, you may want to redefine it to be more strict. If you find that a specific rule is not generating any log events, you may want to check it again to verify that you configured it correctly; perhaps you configured an IP address incorrectly.

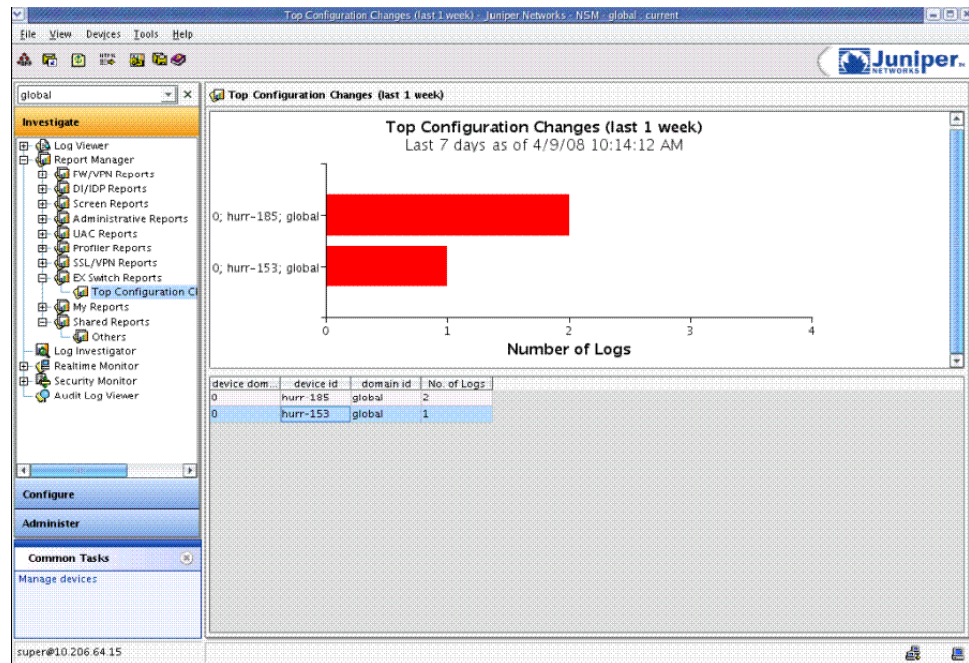
Regular review of the "Top FW/VPN Rules" report can help you to update and optimize the rulebases implemented in your security policies.

Example: Using EX Switch Reports to Track Configuration Changes

In this example, you are a switch administrator responsible for configuring all the managed switches in your network. You routinely update your switch configurations after hours. To track those switches that have undergone the most configuration changes, you generate a "Top Configuration Changes" report each night.

During the week, you can generate a similar report to track switches that have undergone the most configuration changes committed during the past seven days.

[Figure 128 on page 876](#) shows the Top Configuration Changes report.



Example: Using SSL/VPN Reports to Track Authentication Failures

In this example, you are a security administrator responsible for tracking users on the SSL/VPN security devices on your network. You routinely watch for unauthorized users attempting to access your network by tracking authentication failures. To keep watch for potential hackers, you can generate a "Top 10 auth failures for user@realm" report each night.

Example: Using Screen Reports to Identify Attack Trends

In this example, you are a security administrator in the network operations center responsible for tracking potential network attacks. You daily generate and track an "Attacks By Severity" report.

Over time, you notice that the number of critical attacks has increased 20 percent. To verify this, you generate an "Attacks over Time" report for the past 30 days.

The report indicates a recent increase in attacks detected by your firewall. You can generate "Top Attacks", "Top Attackers", and "Top Targets" reports to further investigate the nature and assess the risk of these attacks.

For details on generating and configuring these reports, refer to the *Network and Security Manager Online Help*.

Example: Using DI Reports to Detect Application Attacks

In this example, you are a security analyst responsible for tracking potential deep inspection attacks. You routinely generate an "Attacks By Severity" report daily to track and identify potential attacks.

One day, you notice a significant increase in the number of critical attacks detected by the deep inspection rules you have implemented in your Security Policy. You then generate a "Top Attackers" report for the last day.

The report indicates an IP address as the top attacker for all the DI attacks that you have been tracking. You recognize the IP address as an external server that is running a service using a nonstandard protocol. Although the traffic is not malicious, it happens to match a malicious signature anomaly that you have configured in your DI policy. You can then revise your policy rules to reclassify this traffic.

For details on generating and configuring these reports, refer to the *Network and Security Manager Online Help*.

Using the Watch List

NSM lets you create and configure both a destination and a source watch list. The Destination Watch List contains key hosts within the network against which a proportionally large number of logs is recorded. The Source Watch List contains key hosts outside the network that are sending a large number of log records and are therefore suspected or known sources of attacks on your network.

The watch lists are convenient ways to create a list of source or destination hosts to use as a filter in:

- Log Viewer—Includes logs with destination or source watch lists in a query filter.
- Log Investigator—Investigates logs with destination or source watch lists as data point sources.
- Report Manager—Includes custom reports for destination and source watch lists.

Access the **Destination Watch List** or **Source Watch List** from **Tools > Preferences**. For details about creating and configuring watch lists, refer to the *Network and Security Manager Online Help*.

PART 5

Appendixes

- [Glossary on page 881](#)
- [Unmanaged ScreenOS Commands on page 907](#)
- [SurfControl Web Categories on page 909](#)
- [Common Criteria EAL2 Compliance on page 917](#)
- [Log Entries on page 919](#)

APPENDIX A

Glossary

- [Network and Security Manager \(NSM\) Term Definitions on page 881](#)

Network and Security Manager (NSM) Term Definitions

A

Access List	A list of network prefixes that are compared to a given route. If the route matches a network prefix defined in the access list, the route is either permitted or denied.
Access-Challenge	An additional condition required for a successful Telnet login by an authentication user via a RADIUS server.
Action (Deep Inspection)	A DI action is performed by a security device when the permitted traffic matches the attack object specified in the rule. Deep Inspection actions include <i>drop connection</i> , <i>drop packet</i> , <i>close client</i> , and so on.
Action (firewall)	A firewall action is performed by a security device when the device receives traffic that matches the direction, source, destination, and service. Firewall actions include <i>permit</i> , <i>deny</i> , <i>reject</i> .
Activate Device Wizard	The Activate Device wizard guides you through activating a modeled device in the NSM User Interface.
Add Device Wizard	The Add Device wizard guides you through importing or modeling a new device to the NSM User Interface.
Address Object	An address object represents a component of your network, such as a workstation, router, switch, subnetwork, or any other object that is connected to your network. Use address book objects to specify the network components you want to protect.
Address Shifting	A mechanism for creating a one-to-one mapping between any original address in one range of addresses and a specific translated address in a different range.
Address Spoofing	Address Spoofing is a technique for creating packets with a source IP address that is not the actual interface address. Attackers may use spoofed IP address to perform DDoS attacks while disguising their true address, or to take advantage of a trusted relationship between two hosts. To guard against spoofing attacks, configure a security device to check its own route table. If the IP address is not in the route table, the security device denies the traffic.

Adjacencies	When two routers can exchange routing information with one another, they are considered to have constructed an adjacency. Point-to-point networks have only two routers so those routers automatically form an adjacency. But point-to-multipoint networks are a series of several point-to-point networks. When routers pair in this more complex networking scheme, they are considered to be adjacent to one another.
Advanced Encryption Standard (AES)	AES is a 128-bit encryption key standard. Use AES in your VPNs when you need greater interoperability with other network security devices.
Advertisement	A method a router uses to announce itself to other devices on the network, transmitting basic information including IP address, network mask, and other data.
Aggregate State	A router is in an aggregate state when it is one of multiple virtual BGP routing instances bundled into one address.
Aggregation	The process of combining several different routes in such a way that only a single route advertises itself. This technique minimizes the size of the routing table for the router.
Aggregator	An object used to bundle multiple routes under one common route generalized according to the value of the network mask.
Aggressive Aging	A mechanism to accelerate the timeout process when the number of sessions in the session table surpasses a specified high-watermark threshold. When the number of sessions in the table dips below a specified low-watermark threshold, the timeout process returns to normal.
Antivirus (AV) Scanning	A mechanism for detecting and blocking viruses in File Transfer Protocol (FTP), Internet Message Access Protocol (IMAP), Simple Mail Transfer Protocol (SMTP), Hypertext Transfer Protocol (HTTP)—including HTTP webmail—and Post Office Protocol version 3 (POP3) traffic. Juniper Networks offers an internal AV scanning solution.
APN	Access Point Name. An APN is an IE included in the header of a GTP packet that provides information on how to reach a network. It is composed of two elements: a network ID and an operator ID.
Application Layer Gateway (ALG)	On a security device, an ALG is a software component that is designed to manage specific protocols such as SIP or FTP. The ALG intercepts and analyzes the specified traffic, allocates resources, and defines dynamic policies to permit the traffic to pass securely through the security device.
Area	The most fundamental ordering method in the OSPF routing protocol. An OSPF area divides the internetwork into smaller, more manageable constituent pieces. This technique reduces the amount of information that each router must store and maintain about all the other routers. When a router in the area needs information about another device in or out of the area, it contacts a special router that stores this information. This router is called the Area Border Router (ABR) and contains all essential device information. In addition, the ABR area border router filters all information coming into the area to avoid bogging down other routers in the area with information they may not need.

Area Border Router	A router with at least one interface in Area 0 and at least one interface in another area.
Area Range	A sequence of IP addresses defined by a lower limit and upper limit that indicates a series of addresses of devices that exist within an area.
AS	See Autonomous System.
AS Number	The identification number of the local autonomous system mapped to a BGP routing instance. The ID number can be any valid integer.
AS Path Access List	An access list used by a BGP routing instance to permit or deny packets sent by neighbor routing instances to the current virtual routing instance.
AS Path Attribute Class	The BGP provides four classes of path attributes. Well-Known Discretionary, Optional Transitive, and Optional Non-Transitive.
AS Path String	A string that acts as an identifier for an AS path. It is configured alongside an AS Path access list ID.
Atomic Aggregate	An object used by a BGP router to inform other BGP routers that the local system selected a generalized route.
Atomic Configuration	Atomic configuration is a fail-safe feature in ScreenOS 5.x. For devices running ScreenOS 5.x, if the configuration deployment fails for any reason, the device automatically uses the last installed stable configuration. Additionally, if the configuration deployment succeeds, but the device loses connectivity to the management system, the device rolls back to the last installed configuration. This minimizes downtime and ensures that NSM always maintains a stable connection to the managed device.
Attack Objects	An attack object contains attack patterns for known attacks that attackers can use to compromise your network. Use attack objects in your firewall rules to enable your security devices to detect known attacks and prevent malicious traffic from entering your network.
Attack Protection	Attack Protection is defined by the DI Profile used in a firewall rule.
Audit Log Target	An Audit Log Target is a directive that was sent to a security device.
Audit Log Viewer	The Audit Log Viewer is a module of the NSM User Interface. The Audit Log Viewer records administrative actions. Each audit log includes the date and time the administrative action occurred, the NSM administrator who performed the action, and the domain (global or a subdomain) in which the action occurred.
Authentication	Authentication ensures that digital data transmissions are delivered to the intended receiver. Authentication also assures the receiver of the integrity of the message and its source (where or whom it came from). The simplest form of authentication requires a username and password to gain access to a particular account. Authentication protocols can also be based on secret-key encryption, such as DES, or on public-key systems using digital signatures.

Authentication Header (AH)	See ESP/AH.
Authentication Server Objects	An authentication server provides authentication for NSM administrators and RAS users on your network. Use authentication servers objects to set a default authentication server for the global domain and each subdomain, or access an external RADIUS or SecurID system to provide authentication.
Autonomous System (AS)	An AS is a set of routers set off from the rest of the network and governed by a single technical administration. This router group uses an interior gateway protocol (IGP) or several IGPs and common metrics to route packets within the group. The group also uses an exterior gateway protocol (EGP) to route packets to other ASs. Each AS has a routing plan that indicates what destinations are reachable through it. This plan is called the Network Layer Reachability Information (NLRI) object. BGP routers generate and receive NLRI updates periodically.
Autonomous System Boundary Router	A router that connects an AS running one routing protocol to another AS running a different protocol.
Autonomous System Path	A list of all the autonomous systems that a router update has traveled through in the current transmission.

B

Bastion Host	A bastion host is a hardened system that is configured with the minimal software to support a single network service.
BGP Neighbor	(Also known as a BGP Peer). BGP is a the Border Gateway Patrol dynamic routing protocol. A BGP neighbor is another device on the network that is running BGP. There are two types of BGP neighbors: internal neighbors, which are in the same autonomous system, and external neighbors, which are in different autonomous systems. A reliable connection is required between neighbors and is achieved by creating a TCP connection between the two. The handshake that occurs between the two prospect neighbors evolves through a series of phases or states before a true connection can be made. See Connection States.
Border Gateway Protocol (BGP)	An inter-autonomous system routing protocol. BGP routers and autonomous systems exchange routing information for the Internet.
Broadcast Network	A network that connects many routers together and can send, or broadcast, a single physical message to all the attached routers. Pairs of routers on a broadcast network are assumed to be able to communicate with each other. Ethernet is an example of a broadcast network. On broadcast networks, the OSPF router dynamically detects its neighbor routers by sending Hello packets to the multicast address 224.0.0.5. For broadcast networks, the Hello protocol elects a Designated Router and Backup Designated Router for the network.

C

CIDR (Classless Inter-Domain Routing)

An IP addressing scheme in which a single IP address is used to designate multiple unique IP addresses. A CIDR address includes an IP address and an IP network prefix.

Table 121: CIDR Translation

192.168.0.1/24	192.168.0.1	192.168.0.254	254
192.168.0.1/25	196.168.0.1	192.168.0.126	126
192.168.0.1/26	192.168.0.1	192.168.0.62	62
192.168.0.1/27	192.168.0.1	192.168.0.30	30
192.168.0.1/29	192.168.0.1	192.168.0.9	6
192.168.0.9/29	192.168.0.9	192.168.0.14	6
192.168.0.10/30	192.168.0.10	192.168.0.11	2
10.0.0.0/8	10.0.0.1	10.255.255.254	16777214
10.0.1.17/28	10.0.1.17	10.0.1.30	14

Circuit-Level Proxy

Proxy or Proxy Server is a technique used to cache information on a Web server and acts as an intermediary between a Web client and that Web server. This proxy holds the most commonly and recently used content from the World Wide Web to provide quicker access to content for users and to increase server security.

Classless Routing

Support for interdomain routing, regardless of the size or class of the network. Network addresses are divided into three classes, but these are transparent in BGP, giving the network greater flexibility.

CLI

The CLI is the command line interface.

Cluster List

A list of paths recorded as a packet travels through a BGP route reflector cluster.

Community

A community is a grouping of BGP destination. By updating the community, you automatically update its member destinations with new attributes.

Confederation

An object inside a BGP AS that is a subset of routing instances in the AS. By grouping devices into confederations inside a BGP AS, you reduce the complexity associated with the matrix of routing connections, known as a mesh, within the AS.

Configlet

A configlet is a small, static configuration file that contains information on how a security device can connect to NSM.

Configuration Group A collection of configuration statements whose inheritance can be directed in the rest of the device configuration. The same group can be applied to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.

CRC Errors CRC errors indicate the number of packets generating a cyclic redundancy code error processed through the security device over the selected interface.

D

Data Encryption Standard DES is a 40- and 56-bit encryption algorithm developed by the National Institute of Standards and Technology (NIST). DES is a block encryption method originally developed by IBM. It has since been certified by the U.S. government for transmission of any data that is not classified top secret. DES uses an algorithm for private-key encryption.

Data Encryption Standard-Cipher Block Chaining (DES-CBC) DES-CBC is used to encrypt single DES keys.

Default Route A "catch all" routing table entry that defines the forwarding of traffic for destination networks that are not explicitly defined in the routing table. The destination network for the default route is represented by the network address 0.0.0.0/0.

Delta A delta is a difference, or discrepancy. Example: the differences between the configuration running on the physical device and the difference between the configuration in NSM are known as deltas.

Demilitarized Zone A DMZ is an area between two networks that are controlled by different companies. A DMZ ethernet can be external or internal; external DMZ ethernets link regional networks with routers.

Denial of Service (DoS) Attack A DoS attack is designed to disrupt a network service. Typically, an attacker sends a flood of information to overwhelm a service's system resources, causing the server to ignore valid network requests. Other DoS attacks can cause the service process to crash.

Device Administrator A device administrator is the person who uses WebUI or CLI to manage a single security device.

Device Discovery Rules Sets of rules that define subnets or ranges of IP addresses to scan for EX Series devices in your network.

Device Editor A set of NSM screens used for displaying and editing the configuration of a device.

Device Monitor The Device Monitor displays information about individual devices, their configuration and connection status, and memory usage.

Device Server The Device Server is the component of the NSM management system that handles communication between the GUI Server and the device, collects data from the managed devices on your network, formats configuration information sent to your managed device, and consolidates log and event data.

DHCP (Dynamic Host Configuration Protocol)	DHCP is used to dynamically assign IP addresses to networked computers.
Directive	A directive is a command sent by NSM to your managed devices. Directives include importing, updating, rebooting, and so on. When you send a command to a device or group of devices, NSM creates a job for that command and displays information about that job in the Job Manager.
Distributed Denial of Service (DDoS) Attack	A DoS attack (typically a flood) from multiple source points. A DDoS attacks is more effective than a DoS attack, as it is no longer one computer against one server in an effort to overwhelm the server.
DM (Data Model)	A Data Model is an XML file that contains configuration data for an individual device. The DM is stored in the Device Server; when you create, update, or import a device, the GUI Server edits the Abstract Data Model (ADM) to reflect the changes, then translates that information to the DM.
DMI	Device Management Interface—A common, secure management interface used by all device families added to NSM in release 2008.1 and later releases. DMI is based on a common protocol and device-specific schemas for configuration, inventory management, logging, and status monitoring. DMI schemas can be updated without the need to upgrade NSM.
DNS	The Domain Name System maps domain names to IP addresses.
Domain	A domain is a logical grouping of devices, their policies, and their access privileges. A domain can contain devices, templates, objects, policies, VPNs, administrators, activities, authentication servers, groups—a representation of the all or a subset of the physical devices and functionality on your network. The domain above a domain is the parent domain, and the domain below a domain is the child domain. Domains at the same level are considered peer domains.
Domain Menu	The Domain Menu is the pull-down menu above the navigation tree where domains and subdomains are selected.
Dynamic Routing	A routing method which adjusts to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages populate the network, directing routers to rerun their algorithms and change their routing tables accordingly. There are two common forms of dynamic routing, including Distance Vector Routing and Link State Routing.

E

Encryption	Encryption is the process of changing data into a form that can be read only by the intended receiver. To decipher the message, the receiver of the encrypted data must have the proper decryption key. In traditional encryption schemes, the sender and the receiver use the same key to encrypt and decrypt data. Public-key encryption schemes use two keys: a public key, which anyone may use, and a corresponding private key, which is possessed only by the person who created it. With this method, anyone may send a message encrypted with the owner's public key, but only the owner has the private key necessary to decrypt it. PGP (Pretty Good Privacy) and DES (Data Encryption Standard) are two of the most popular public-key encryption schemes.
Equal Cost Multipath (ECMP)	Equal cost multipath assists with load balancing among two to four routes to the same destination or increases the effective bandwidth usage among two or more destinations. When enabled, security devices use the statically defined routes or dynamically learn multiple routes to the same destination through a routing protocol. The security device assigns routes of equal cost in round robin fashion. Default: disabled
ESP/AH	AH and ESP are IP level security headers that were originally proposed by the Network Working Group focused on IP security mechanisms known as IPSec. The term IPSec refers to packets, keys, and routes associated with ESP and AH headers. The IP Authentication Header (AH) provides authentication. The IP Encapsulating Security Header (ESP) provides confidentiality to IP datagrams.
Ethernet	Ethernet is a local area network (LAN) technology invented at the Xerox Corporation, Palo Alto Research Center. Ethernet is a best-effort delivery system that uses CSMA/CD technology. Ethernet can be run over a variety of cable schemes, including thick coaxial, thin coaxial, twisted pair, and fiber optic cable. Ethernet is a standard for connecting computers into a local area network (LAN). The most common form of Ethernet is called 10BaseT, which denotes a peak transmission speed of 10 Mbps using copper twisted-pair cable.
Export Rules	When you have two or more virtual routers on a security device, you can configure export rules that define which routes on one virtual router are allowed to be learned by another virtual router. See also Import Rules.
External Neighbors	Two BGP routers that are peers that reside in two different autonomous systems.
Extranet	An extranet connects two or more intranets. If an intranet as a company's internal Web site enables users inside the company to communicate and exchange information, an extranet connects that virtual space with another company's intranet, thus enabling these two (or more) companies to share resources and communicate over the Internet in their own virtual space. This technology greatly enhances business to business communications.

F

Filters	A filter organizes log entries based on administrator specifications.
Firewall	A firewall device that protects and controls incoming and outgoing traffic on network connections. Firewalls protect internal servers from damage (intentional or otherwise) and enable authorized external access.

G

G-PDU	A G-PDU is a user data message. It consists of a T-PDU plus a GTP header.
Gateway	Also called a router, a gateway is a program or a special-purpose device that transfers IP datagrams from one network to another until the final destination is reached.
GBIC	A Gigabit Interface Connector (GBIC) is the kind of interface module card used on some security devices for connecting to a fiber optic network.
GGSN	Gateway GPRS Support Node.
GI Interface	The interface between a GSN and an external network or the Internet.
Global Domain	A domain is a logical grouping of devices, their policies, and their access privileges. The global domain is the top level, or root domain, that contains all subdomains.
GMT (Greenwich Mean Time)	GMT is the Greenwich, England mean solar time. GMT is also known as Universal Time and is used for calculating time worldwide.
Gn Interface	The interface between two GSNs within the same PLMN.
Gp Interface	The interface between two GSNs located in different PLMNs.
GPRS	General Packet Radio Service. A packet-based technology that enables high-speed wireless Internet and other data communications. GPRS provides more than three to four times greater speed than conventional GSM systems. Using a packet data service, subscribers are always connected and always online so services are easy and quick to access.
Group Expression Objects	A group expression object represents a statement that sets conditions for authentication requirements, enabling you to combine multiple external user objects. You can create group expressions using the operator OR, AND, or NOT to combine user objects, user group objects, or other group expressions.
Groups	A group organizes previously-created devices into user-defined groups that make it easier for you to configure and manage devices in your domain. Groups enable you to execute certain NSM operations on multiple security devices at the same time.
GRX	GPRS Roaming Exchange.
GSM	Global System for Mobile Communications.
GTP	GPRS Tunneling Protocol.

GTP Tunnel	A GTP tunnel in the GTP-U plane is defined for each PDP Context in the GSNs. A GTP tunnel in the GTP-C plane is defined for all PDP Contexts with the same PDP address and APN (for Tunnel Management messages) or for each MS (for messages not related to Tunnel Management). A GTP tunnel is identified in each node with a TEID, an IP address and a UDP port number. A GTP tunnel is necessary to forward packets between an external network and an MS user.
GTP-C Message	GTP-Control Message. Control plane messages are exchanged between GSN pairs in a path. The control plane messages are used to transfer GSN capability information between GSN pairs, to create, update and delete GTP tunnels and for path management.
GTP-PDU	A GTP Protocol Data Unit is either a GTP-C message or a GTP-U message.
GTP-U Message	GTP-User Data message. User plane messages are exchanged between GSN pairs or GSN/RNC pairs in a path. The user plane messages are used to carry user data packets, and signalling messages for path management and error indication.
GUI Server	The GUI Server manages the system resources and data that drives NSM functionality. The GUI Server contains the NSM databases, and centralizes information for devices, their configurations, attack and server objects, and policies.

H

H.323	The H323 Application Layer Gateway (ALG) lets you to secure Voice-over-IP (VoIP) communication between terminal hosts, such as IP phones and multimedia devices. In such a telephony system, gatekeeper devices manage call registration, admission, and call status for VoIP calls. Gatekeepers can reside in the two different zones, or in the same zone.
Hardened System	A hardened system is a secure server with all appropriate security patches and bug fixes; these systems are designed to resist penetration.
Hello Interval	The amount of time that elapses between instances of Hello Packets.
Hello Packet	A Hello packet is a message sent out to the current network to announce the presence of the current routing instance to the network. Hello packets aid in the discovery of neighbors and in a router being able to connect to other devices on the network. When an OSPF interface is created, the interface sends Hello packets to the network to announce itself.
Histogram	A histogram is a vertical graph that represents different amounts by thin, color-coded bands or bars. These bars represent a frequency distribution; heights of the bars represent observed frequencies.
HLR	Home Location Register.
Hold Time	In OSPF, the maximum amount of time between instances of initiating Shortest Path First (SPF) computations. In BGP, the maximum amount of time that elapses between message transmissions between a BGP speaker and its neighbor.

I

ICMP Flood	An ICMP flood contains ICMP pings so numerous that they overload a system with echo requests, causing the system to expend all its resources responding until it can no longer process valid network traffic. If you set a threshold to invoke ICMP flood attack protection when exceeded, ICMP flood attacks are recorded as statistics.
IE	Information Element.
IKE Proposal Objects	An IKE proposal is a set of encryption keys and authentication algorithms that is used to negotiate a VPN connection. An IKE proposal object is a representation of an IKE proposal in the NSM UI.
Import Rules	When you have two or more virtual routers on a security device, you can configure import rules on one virtual router that define which routes are allowed to be learned from another virtual router. If you do not configure any import rules for a virtual router, all routes that are exported to that virtual router are accepted. See also Export Rules.
IMSI	International Mobile Station Identity.
In-Device Policy Management	Mode of policy management done through the Device Editor on a specific device and not through the central NSM Policy Manager. If you select this method to manage policies on a J Series or SRX Series device, the NSM Policy Manager, Object Manager, and VPN Manager are all disabled for that device.
Infranet Controller	The policy management component of Juniper Networks UAC solution.
Infranet Enforcer	The policy enforcement point or firewall within a Juniper Networks UAC solution.
Internet Control Message Protocol (ICMP)	ICMP is a network-layer protocol that does not carry user data, but does encapsulate its messages in IP datagrams. ICMP provides a query and response system (with error-reporting) used to determine if another system on the network can receive and send data. An ICMP echo request is also known as a <i>ping</i> .
Internet Key Exchange	IKE is a method for exchanging keys for encryption and authentication over an unsecured medium, such as the Internet.
Internet Protocol (IP)	IP is an Internet standard protocol that defines a basic unit of data called a datagram. A datagram is used in a connectionless, best-effort, delivery system. The Internet protocol defines how information gets passed between systems across the Internet.

IP Address	<p>Each node on a TCP/IP network usually has an IP address. The IP address has a network number portion and a host number portion:</p> <ul style="list-style-type: none">• Class A, >32,768 nodes, address format: nnn.hhh.hhh.hhh• Class B, 256-32,768 nodes, address format: nnn.nnn.hhh.hhh• Class C, <256 nodes, address format: nnn.nnn.nnn.hhh <p>This address format is called decimal dot format. The \"n\" represents a digit of a network number and \"h\" represents a digit of a host number; for example, 128.11.2.30. If you are sending data outside of your network, such as to the Internet, you need to obtain the network number from a central authority, currently the Network Information Center. See also Subnet Mask.</p>
IP Gateway	<p>Also called a router, an IP gateway is a program or a special-purpose device that transfers IP datagrams from one network to another until the final destination is reached.</p>
IP Pool Objects	<p>An IP Pool object represents a range of IP addresses. Use IP Pool object to configure a DHCP server for your managed devices.</p>
IP Security (IPSec)	<p>IPSec is a security standard maintained by the Internet Engineering Task Force (IETF). The IPSec protocol suite provides everything you need for secure communications—authentication, integrity, and confidentiality—and makes key exchange practical even in larger networks. See also DES-CBC, ESP/AH.</p>
IP Sweep	<p>An IP sweep is similar to a port scan attack. Attackers perform IP sweeps by sending ICMP echo requests (or pings) to different destination addresses and wait for replies that indicate the IP address of a target. If a remote host pings 10 addresses in 0.3 seconds, the security device flags the event as an IP sweep attack and drops the connection to prevent replies.</p>
IP Tracking	<p>A mechanism for monitoring configured IP addresses to see if they respond to ping or ARP requests. You can configure IP tracking with NSRP to determine device or VSD group failover. You can also configure IP tracking on a device interface to determine if the interface is up or down.</p> <p>ISAKMP. The Internet Security Association and Key Management Protocol (ISAKMP) provides a framework for Internet key management and provides the specific protocol support for negotiation of security attributes. By itself, it does not establish session keys, however it can be used with various session key establishment protocols to provide a complete solution to Internet key management.</p>
J	
Job Manager	<p>The Job Manager is a module of the NSM User Interface. Job Manager tracks the progress of the command as it travels to the device and back to the management server.</p>
JSRP	<p>Junos Services Redundancy Protocol.—A process that controls chassis clustering of Junos devices.</p>

K

Keepalive	The amount of time in seconds that elapses between keepalive packets which ensures that the TCP connection between the local BGP router and a neighbor router is up. This value is equal to one-third of the hold time. The default is 60 seconds.
Key Management	The only reasonable way to protect the integrity and privacy of information is to rely upon the use of secret information in the form of private keys for signing and/or encryption. The management and handling of these pieces of secret information is generally referred to as “key management.” This includes the activities of selection, exchange, storage, certification, expiration, revocation, changing, and transmission of keys. Most of the work in managing information security systems lies in the key management.

L

Land Attack	During a Land Attack, attackers may send spoofed SYN packets that contain the IP address of the target as both the destination and source IP address to create an empty connection. These connections flood the target system, overwhelming it and causing a denial-of-service. You can configure security devices to block Land Attack and record Land Attack attempts.
Launch Pad	An otherwise blank user interface pane that provides access to commonly used functionality within the associated NSM module.
Link State	Link state routing protocols operate using an algorithm commonly called the Shortest Path First (SPF) algorithm. Instead of relying on rumored information from directly connected neighbors as in distance vector protocols, each router in a link state system maintains a complete topology of the network and computes SPF information based on the topology.
Link state Advertisement	Link State Advertisements (LSAs) are the conveyance that enables OSPF routers to make device, network, and routing information available for the link state database. Each router retrieves information from the LSAs sent by other routers on the network to construct a picture of the entire internetwork from which they distill path information to use in the routing table.
Load Balancing	Load balancing distributes workload to processors to improve the throughput of a concurrent connections.
Local Preference	To provide better information than the Multi-Exit Discriminator (MED) value provides for a packet's path selection, BGP provides an attribute known as the LOCAL_PREF or local preference value. You can configure the LOCAL_PREF attribute so that it has a higher value for prefixes received from a router that provides a desired path to be higher than prefixes heard on the router that provides a less desirable path. The higher the value, the more preferred the route. The LOCAL_PREF attribute is the metric most often used in practice to express preferences for one set of paths over another.
Lockout	Lockout is an object state during which the object cannot be edited.
Log	A Log is a grouping of log entries.
Log Category	A log category defines the log type (alarm, config, traffic, and so on).

Log ID	A log ID is a unique ID for the log entry, derived from the combination of the date and log number.
Log Investigator	The Log Investigator is a module of the NSM User Interface. The Log Investigator contains tools for analyzing your log entries in depth. Use the Log Investigator to manipulate and change constraints on log information, correlate log entries visually and rapidly, and filter log entries while maintaining the broader picture.
Log Viewer	The Log Viewer is a module of the NSM User Interface. The Log Viewer displays log entries that your security devices generate based on criteria that you defined in your security policies, on the Device Server, and in the device configuration. Logs appear in table format; each row contains a single log, and each column defines specific information for a log.
Loopback Interface	A logical interface that emulates a physical interface on the security device, but is always in the up state as long as the device is up. You must assign an IP address to a loopback interface and bind it to a security zone.
M	
Main Display Area	The main display area displays the content for the currently selected module or module contents.
Management System	The management system includes the GUI Server and Device Server. You can deploy the GUI Server and Device Server on separate servers; however, the combination of the two servers is known as the management system.
Mapped IP Address	A MIP is a direct one-to-one mapping of traffic destined for one IP address to another IP address.
MCC	Mobile Country Code.
MD5	Message Digest (version) 5 is an algorithm that produces a 128-bit message digest (or hash) from a message of arbitrary length. The resulting hash is used to verify authenticity.
Media Gateway Control Protocol (MGCP)	MGCP is a text-based, application layer protocol that can be used for call setup and call control. The protocol is based on a master/slave call control architecture: the media gateway controller (call agent) maintains call control intelligence, and media gateways carry out the instructions from the call agent.
Member AS	The name of the autonomous system being included in a BGP confederation.
Menu Bar	The menu bar is the upper section of the NSM UI. The menu bar contains accessible commands.
Metric	A value associated with a route that the virtual router uses to select the active route when there are multiple routes to the same destination network with the same preference value. The metric value for connected routes is always 0. The default metric value for static routes is 1, but you can specify a different value when defining a static route.
MNC	Mobile Network Code.

Modeling Modeling is the process of creating a non-deployed device configuration in the NSM UI.

Modules A module is a first-level element in the NSM navigation tree.

MS Mobile Station.

MSIN Mobile Subscriber Identification Number.

N

NAT Object A NAT object is a global object that contains references to device-specific NAT configurations, enabling multiple devices to share a single object. Use the Device Manager to configure NAT for each device, then create a global NAT object that includes the device-specific NAT configuration. Use global NAT objects in security policies and VPNs; when you update a device, that device automatically replaces the global NAT object with its device-specific NAT configuration.

NAT-Traversal (NAT-T) A method for allowing IPSec traffic to pass through NAT devices along the data path of a VPN by adding a layer of UDP encapsulation. The method first provides a means for detecting NAT devices during Phase 1 IKE exchanges, and then a means for traversing them after Phase 2 IKE negotiations are complete.

Navigation Tree The navigation tree displays the 11 NSM modules in the left pane of the NSM window.

Neighbor To begin configuring a BGP network, you need to establish a connection between the current device and a counterpart, adjacent device known as a neighbor or peer. While this counterpart device may seem like unneeded information at first, it is actually central to the way BGP works. Unlike RIP or OSPF, you now have to configure two devices, both the current router and its neighbor, for BGP to work. While this requires more effort, it enables networking to occur on a larger scale as BGP eludes deploying the limited advertising techniques inherent to interior networking standards.

NetScreen Redundancy Protocol (NSRP) NSRP is a proprietary protocol that provides configuration and run time object (RTO) redundancy and a device failover mechanism for security devices in a high availability (HA) cluster.

Network Address Translation (NAT) NAT is a standard for translating secure IP addresses to temporary, external, registered IP address from the address pool. NAT enables trusted networks with privately assigned IP addresses to access the Internet, eliminating the need to use a registered IP address for every machine in your network.

NSAPI Network Service Access Point Identifier.

NSGP NetScreen Gatekeeper Protocol.

NSM Administrator The NSM administrator is the person who uses NSM User Interface to manage their devices.

O

Object	Objects represent reusable information, such as network addresses, individual users and user groups, and commonly used configuration data. In NSM, objects are shared objects, meaning they are shared between the global domain and all subdomains. Objects are the building blocks of the NSM management system.
Object Manager	A module of the NSM User Interface that contains the objects used in your NSM system. An object is a re-usable, basic NSM building block that contains specific information; you use objects to create device configurations, policies, and VPNs. All objects are shared, meaning that they can be shared by all devices and policies in the domain.
OnSite Administrator	The person who installs a configlet using Rapid Deployment.
Open Shortest Path First (OSPF)	A dynamic routing protocol intended to operate within a single Autonomous System.

P

Packet Filtering	Packet filtering is a router/firewall process that uses access control lists (ACL) to restrict flow of information based on protocol characteristics such as source/destination IP address, protocol, or port used. Generally, packet-filtering routers do not track sessions except when doing NAT (which tracks the session for NAT purposes).
PDP	Packet Data Protocol.
PDP Context	A user session on a GPRS network.
PDU	Protocol Data Unit.
Peer	See Neighbor.
Ping of Death	The ping of death is an intentionally oversized or irregular ICMP packet that can trigger a Denial of Service condition, freezing, or other adverse system reactions. You can configure a security device to detect and reject oversized or irregular packet sizes.
PLMN	Public Land Mobile Network. A public network dedicated to the operation of mobile radio communications.
Point-to-Multipoint Network	A non-broadcast network where OSPF treats connections between routers as point-to-point links. There is no election of a designated router and no LSA generated for the network. A router in a point-to-multipoint network sends Hello packets to all neighbors with which it can directly communicate.
Point-to-Point Network	Joins two routers over a Wide Area Network (WAN). An example of a point-to-point network is two security devices connected via an IPSec VPN tunnel. On point-to-point networks, the OSPF router dynamically detects neighbor routers by sending Hello packets to the multicast address 224.0.0.5.

Point-to-Point Protocol over Ethernet	Allows multiple users at a site to share the same digital subscriber line, cable modem, or wireless connection to the Internet. You can configure PPPoE client instances, including the username and password, on any or all interfaces on some security devices.
Policy	A security policy is the combination of both firewall rulebases and all rules into a comprehensive plan that defines how the security device works on your network.
Port Address Translation (PAT)	The translation of the original source port number in a packet to a different, randomly designated port number.
Port Mapping	The translation of the original destination port number in a packet to a different, predetermined port number.
Port Mode	A feature supported on some security devices, port mode allows you to select one of several different sets of port, interface, and zone bindings on the device. Changing the port mode removes any existing configurations on the device and requires a system reset.
Port Scan	A port scan attack occurs when packets are sent out to different port numbers, for the purpose of scanning the available services in hopes that one port will respond. If a remote host scans 10 ports in 0.3 seconds, the security device flags this as a port scan attack and drops the connection.
Preference	A value associated with a route that the virtual router uses to select the active route when there are multiple routes to the same destination network. The preference value is determined by the protocol or origin of the route. The lower the preference value of a route, the more likely the route is to be selected as the active route.
Prefix	An IP address that represents a route.
Process Status	The process status displays information about processes on a security device.
Protocol	Protocols are predefined services (HTTP, SNMP, Telnet, and so on) that are enabled for the security device.
PT	Protocol Type.

R

RADIUS	Remote Authentication Dial-In User Service is a service for authenticating and authorizing remote access service (RAS) users.
RAS (remote access services)	RAS is the acronym for remote access services, which enable users to access services protected by your security devices. Typically, you use a VPN to enable RAS, then add RAS users to the VPN.
Real Time Streaming Protocol (RTSP)	RTSP is an application layer protocol for controlling the delivery of a stream of real-time multimedia content.

Realtime Monitor	The Realtime Monitor is a module of NSM User Interface. It contains the Device Monitor, the VPN Monitor, and the NSRP Monitor.
Receive Collisions	The number of collisions on the line detected by the Carrier Sense Multiple Access Collision Detection (CSMA/CD) protocol.
Redistribution	The process of importing a route into the current routing domain from another part of the network that uses another routing protocol. When this occurs, the current domain has to translate all the information, particularly known routes, from the other protocol. For example, if you are on an OSPF network and it connects to a BGP network, the OSPF domain has to import all the routes from the BGP network to inform all of its devices about how to reach all the devices on the BGP network. The receipt of all the route information is known as route redistribution.
Redistribution List	A list of routes the current routing domain imported from another routing domain using a different protocol.
Remote Setting Objects	A Remote Settings object defines the DNS and WINS servers that are assigned to L2TP RAS users after they have connected to the L2TP tunnel.
Report Manager	Report Manager is a module of the NSM User Interface. Use Report Manager to generate and view reports summarizing log and alarm originating from the managed security devices in your network. You can use these reports to track and analyze log incidents, network traffic and potential attacks.
Report Procedure Call (RPC)	The RPC is a protocol that one program can use to request a service from a program located in another computer in a network.
Role-Based Administration (RBA)	Role-based administration enables you to define strategic roles for your administrators and create domains to organize your network devices. Use role-based administration to create a security environment that reflects your current offline administrator roles and responsibilities.
Route Flap Damping	BGP provides a technique to block the advertisement of the route somewhere close to the source until the route becomes stable. This method is called flap damping. Route flap damping allows routing instability to be contained at an AS border router adjacent to the region where instability is occurring. The impact of limiting the unnecessary propagation is to maintain reasonable route change convergence time as a routing topology grows.
Route Map	Route maps are used with BGP to control and modify routing information and to define the conditions by which routes are redistributed between routing domains. A route map contains a list of route map entries, each containing a sequence number and a match and a set value. The route map entries are evaluated in the order of an incrementing sequence number. Once an entry returns a matched condition, no further route maps are evaluated. Once a match has been found, the route map carries out a permit or deny operation for the entry. If the route map entry is not a match, then the next entry is evaluated for matching criteria.
Route Redistribution	Route redistribution is the exporting of route rules from one virtual router to another.

Route Reflector	A router whose BGP configuration enables readvertising of routes between Interior BGP (IBGP) neighbors or neighbors within the same BGP AS. A route reflector client is a device that uses a route reflector to readvertise its routes to the entire AS. It also relies on that route reflector to learn about routes from the rest of the network.
Routing Information Protocol (RIP)	A dynamic routing protocol used within moderate-sized autonomous systems.
Routing Table	A list in a virtual router's memory that contains a realtime view of all the connected and remote networks to which a router is currently routing packets.
Rule	A rule is a statement that defines a specific type of network traffic. When traffic passes through the security device, the device attempts to match that traffic against its list of rules. If a rule is matched, the device performs the action defined in the rule against the matching traffic.
Rulebase	A rulebase contains rules. a rulebase provides a method of detecting and acting upon suspicious traffic. A NSM security policy can contain three rulebases: Zone, Global, and Multicast.
Run Time Object (RTO)	A code object created dynamically in memory during normal operation. Some examples of RTOs are session table entries, ARP cache entries, certificates, DHCP leases, and IPSec Phase 2 security associations (SAs).
S	
Scheduled Object	A schedule object defines a time interval that a firewall rule is in effect. You use a schedule object in your firewall rule to determine when a device enforces that rule.
Secure Access Device	A Juniper Networks SSL VPN appliance.
Secure Copy (SCP)	A method of transferring files between a remote client and a security device using the SSH protocol. The security device acts as an SCP server, accepting connections from SCP clients on remote hosts.
Secure Server Protocol (SSP)	For communication between the UI, the GUI Server, and the Device Server, NSM uses <i>SSP</i> , a modified version of TCP that is more reliable than ordinary TCP, requires less CPU and memory resources from servers, and reduces the number of acknowledgement packets on the network. SSP uses AES encryption and SH1 authentication for all connections.
Secure Shell (SSH)	A protocol that allows device administrators to remotely manage the device in a secure manner. You can run either an SSH version 1 or version 2 server on the security device.
Security Association	The security association combines the Security Parameters Index and a destination address. Required for both Authentication Header and Encapsulating Security Payload protocols. See also Security Parameters Index.

Security Device	A security device enables access to your network components and protects your network against malicious traffic. NSM can manage security devices running ScreenOS 5.x and ScreenOS 6.0.x. All devices from NetScreen-5XT to the NetScreen-5400 are supported, except the NetScreen-5, NetScreen-10, and NetScreen-1000. NSM also supports the NetScreen-5GT running ScreenOS 4.0-DIAL2. NSM can also manage vsys configurations, NSRP clusters, and extranet devices.
Security Parameters Index	The SPI is a hexadecimal value which uniquely identifies each tunnel. It also tells the security device which key to use to decrypt packets.
Security Policies	A security policy defines access to your network, including permitted services, users, and time periods. Use security policies to control the shape of your network traffic as it passes through the firewall, or log specific network events.
Security Zone	A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic via access policies.
Server Manager	The Server Manager is a module of the NSM User Interface. Server Manager contains server objects that represent your management system components. Use Server Manager to manage and monitor the individual server processes that comprise your NSM system.
Service Object	Service objects represent the IP traffic types for existing protocol standards. Security devices monitor and manage network traffic using these protocols. NSM includes predefined service objects for most standard services. You can also create custom service objects to represent services that are not included in the list of predefined service objects, or to represent a custom service running on your network.
Session Description Protocol (SDP)	SDP session descriptions appear in many SIP messages and provide information that a system can use to join a multimedia session. SDP might include information such as IP addresses, port numbers, times, dates, and information about the media stream.
Session Initiation Protocol (SIP)	SIP is an IETF (Internet Engineering Task Force)-standard protocol for initiating, modifying, and terminating multimedia sessions over the Internet. Such sessions might include conferencing, telephony, or multimedia, with features such as instant messaging and application-level mobility in network environments.
SGSN	Serving GPRS Support Node.
SHA-1	Secure Hash Algorithm-1, an algorithm that produces a 160-bit hash from a message of arbitrary length. (It is generally regarded as more secure than MD5 because of the larger hashes it produces.)
Shared Objects	A shared object is an object that can be shared across domains.
Short Frame	A short frame contains less than 64 bytes of data.

Signaling Message	Any GTP-PDU except the G-PDU. GTP signalling messages are exchanged between GSN pairs in a path. The signaling messages are used to transfer GSN capability information between GSN pairs and to create, update and delete GTP tunnels.
Source Interface-Based Routing (SIBR)	SIBR allows the security device to forward traffic based on the source interface (the interface on which the data packet arrives on the security device).
Source Route	The source route is a option in the IP header. An attacker can use the source route option to enter a network with a false IP address and have data sent back to the attacker's real address.
Stateful Inspection	A firewall process that checks the TCP header for information on the session's state. The process checks whether it is initializing (SYN), ongoing (SYN/ACK), or terminating (FIN). A stateful inspection firewall tracks each session flowing through it, dropping packets from unknown sessions that appear to be part of an ongoing or illegal sessions. All security devices are stateful inspectors.
Static Routing	User-defined routes that cause packets moving between a source and a destination to take a specified path. Static routing algorithms are table mappings established by the network administrator prior to the beginning of routing. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.
Status Bar	The status bar is the lower section of the NSM UI. The status bar displays supplemental information.
Subdomain	A subdomain is a domains under the global domain.
Subinterface	A subinterface is a logical division of a physical interface that borrows the bandwidth it needs from the physical interface from which it stems. A subinterface is an abstraction that functions identically to an interface for a physically present port and is distinguished by 802.1Q VLAN tagging.
Subnet Mask	A subnet mask enables you to define subnetworks. For example, if you have a class B network, a subnet mask of 255.255.255.0 specifies that the first two portions of the decimal dot format are the network number, while the third portion is a subnet number. The fourth portion is the host number. If you do not want to have a subnet on a class B network, you would use a subnet mask of 255.255.0.0. A network can be subnetted into one or more physical networks which form a subset of the main network. The Subnet Mask is the part of the IP address which is used to represent a subnetwork within a network. Using Subnet Masks enables you to use network address space which is normally unavailable and ensures that network traffic does not get sent to the whole network unless intended. See also IP Address.
Super Admin(istrator)	The super administrator is the default administrator for all domains. The superadmin has immutable powers. You cannot change or delete permissions for the super administrator; you can, however, change the password for the super administrator.

SYN Attack	A SYN attack occurs when SYN packets overwhelm a network by initiating so many connection attempts or information requests that the network can no longer process legitimate connection requests, resulting in a Denial of Service.
syslog	A protocol that enables a device to send log messages to a host running the syslog daemon (syslog server). The syslog server then collects and stores these log messages locally.
T	
T-PDU	A T-PDU is the payload that is tunnelled in the GTP tunnel.
Tear Drop Attack	A Tear Drop Attack occurs when the first and second parts of a fragmented packet overlap, the server attempting to reassemble the packet can crash. If the security device sees this discrepancy in a fragmented packet, it drops the packet.
TEID	Tunnel Endpoint Identifier. The TEID uniquely identifies a tunnel endpoint in the receiving GTP-U or GTP-C protocol entity. The receiving end side of a GTP tunnel locally assigns the TEID value the transmitting side has to use. The TEID values are exchanged between tunnel endpoints using GTP-C messages.
Templates	<p>A template is a device configuration that you can define once and then use for multiple devices. You can specify most device configuration values in a template. In a template, you can define only those configuration parameters that you want to set; you do not need to specify a complete device configuration.</p> <p>The software remembers static routes until you remove them. However, you can override static routes with dynamic routing information through judicious assignment of administrative distance values. To do this, you must ensure that the administrative distance of the static route is higher than that of the dynamic protocol.</p>
TID	Tunnel Identifier.
TLS	Transport Layer Security—a cryptographic protocol that provides secure communication between the NSM UI and the NSM GUI Server.
Toolbar	The toolbar is the upper section of the NSM UI. The toolbar contains icons that relate to accessible commands.
Transmission Control Protocol/Internet Protocol (TCP/IP)	A set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks. A communications protocol which enables computers with different operating systems to communicate with each other. Controls how data is transferred between computers on the Internet.
Triple DES (3DES)	3DES is a more powerful version of DES in which the original DES algorithm is applied in three rounds, using a 168-bit key. DES provides a significant performance savings but is considered unacceptable for many classified or sensitive material transfers.
Trojan	A trojan is a program with hidden functionality. Trojans often install a remote administration program (known as a backdoor) that enables attackers to access the target system.

Trunk Port	A trunk port enables a switch to bundle traffic from several VLANs through a single physical port, sorting the various packets by the VLAN identifier (VID) in their frame headers.
Trust Zone	One of two predefined zones that enables packets to be secured from being seen by devices external to your current domain.
Tunnel Interface	A tunnel interfaces is the opening, or doorway, through which traffic to or from a VPN tunnel passes. A tunnel interface can be numbered (that is, assigned an IP address) or unnumbered. A numbered tunnel interface can be in either a tunnel zone or security zone. An unnumbered tunnel interface can only be in a security zone that contains at least one security zone interface. The unnumbered tunnel interface borrows the IP address from the security zone interface.
Tunnel Zone	A tunnel zone is a logical segment that hosts one or more tunnel interfaces. A tunnel zone is associated with a security zone that acts as its carrier.
Tunneling	A method of data encapsulation. With VPN tunneling, a mobile professional dials into a local Internet Service Provider's Point of Presence (POP) instead of dialing directly into their corporate network. This means that no matter where mobile professionals are located, they can dial a local Internet Service Provider that supports VPN tunneling technology and gain access to their corporate network, incurring only the cost of a local telephone call. When remote users dial into their corporate network using an Internet Service Provider that supports VPN tunneling, the remote user as well as the organization knows that it is a secure connection. All remote dial-in users are authenticated by an authenticating server at the Internet Service Provider's site and then again by another authenticating server on the corporate network. This means that only authorized remote users can access their corporate network, and can access only the hosts that they are authorized to use.
U	
UDP Flood	A UDP flood is an attack using multiple UDP packets. An attacker can send UDP packets to slow the target system to the point that it can no longer handle valid connections. You can configure the security device with a threshold to invoke UDP flood attack protection; when UDP packet flow exceeds this threshold, the device records the UDP flood attack as a statistics.
Universal Resource Locator (URL)	A URL is a standard method of specifying the location of an available electronic resource. Also known as a location or address, a URL specifies the location of files on servers. A general URL has the syntax protocol://address. For example, http://www.srl.rmit.edu.au/pd/index.html specifies that the protocol is http and the address is www.srl.rmit.edu.au/pd/index.html .
Universal Unique Identifier UUID)	The UUID is a 128-bit number assigned to any object within a Distributed Computing Environment (DCE) cell which is guaranteed to be unique.
Untrust Zone	One of two predefined zones that enables packets to be seen by devices external to your current domain.
User	A user is a person using the network your security devices are protecting. NSM supports two types of users: local users and external users.

User Datagram Protocol (UDP)	UDP is a protocol in the TCP/IP protocol suite that enables an application program to send datagrams to other application programs on a remote machine. UDP provides an unreliable and connectionless datagram service and does not guarantee delivery or duplicate detection; it does not use acknowledgments, or control the order of arrival.
User Interface (UI)	The NSM graphical User Interface (UI) is used to control the NSM system. Using the UI, you can configure NSM administrators, add devices, edit policies, view reports, and so on.
User Object	User objects represent the users of your managed devices. You can include user objects or groups in security policies or VPNs to permit or deny access to individuals or groups.
V	
View	A view is an admin-defined subset of column settings and filters in the Log Viewer.
Virtual Chassis	Stacked EX Series devices functioning as one logical EX Series switch or an SRX cluster represented in NSM as a virtual chassis..
Virtual IP Address	A VIP address maps traffic received at one IP address to another address based on the destination port number in the packet header.
Virtual Link	A logical path from a remote OSPF area to the backbone area.
Virtual Local Area Network (VLAN)	A VLAN is a logical rather than physical grouping of devices that constitute a single broadcast domain. VLAN members are not identified by their location on a physical subnetwork but through the use of tags in the frame headers of their transmitted data. VLANs are described in the IEEE 802.1Q standard.
Virtual Private Network (VPN)	A VPN is an easy, cost-effective and secure way for corporations to provide telecommuters and mobile professionals local dial-up access to their corporate network or to another Internet Service Provider (ISP). Secure private connections over the Internet are more cost-effective than dedicated private lines. VPNs are possible because of technologies and standards such as tunneling, screening, encryption, and IPSec.
Virtual Router (VR)	A virtual router is the component of ScreenOS that performs routing functions. By default, a security device contains two virtual routers: Untrust-VR and Trust-VR.
Virtual Security Device (VSD)	A VSD is a single logical device composed by a set of physical security devices.
Virtual Security Interface (VSI)	A VSI is a logical entity at layer 3 that is linked to multiple layer 2 physical interfaces in a VSD group. The VSI binds to the physical interface of the device acting as master of the VSD group. The VSI shifts to the physical interface of another device in the VSD group if there is a failover and it becomes the new master.
Virtual System (VSYS)	A virtual system is a subdivision of the main system that appears to the user to be a standalone entity. Virtual Systems reside separately from each other. Each one can be managed by its own Virtual System Administrator.

VPN Manager VPN Manager is a module of the NSM User Interface. Use VPN Manager to design a system level VPN and automatically set up all connections, tunnels, and rules for all devices in the VPN.

W

WebTrends A product offered by NetIQ that allows you to create customized reports based on the logs generated by a security device. When you use WebTrends, you can display the information you need in a graphical format.

Windows Internet Naming Service (WINS) WINS is a service for mapping IP addresses to NetBIOS computer names on Windows NT server-based networks. A WINS server maps a NetBIOS name used in a Windows network environment to an IP address used on an IP-based network.

WinNuke Attack A WinNuke attack can crash any computer on the Internet running Windows by introducing a NetBIOS anomaly that forces Windows to restart. You can configure the security device to scan any incoming Microsoft NetBIOS Session Service packets, modify them, and record the event as a WinNuke attack.

Worm A worm is a self-replicating attack program. Worms differ from typical viruses in that they are completely automatic—no attacker interaction is required. When the worm locates a vulnerable target, it immediately and automatically infects the new host with its malicious code. The newly infected host repeats the process and attempts to infect more hosts.

X

XAuth A protocol composed of two components—remote VPN user authentication (username plus password) and TCP/IP address assignments (IP address, netmask, DNS server, and WINS server assignments).

Z

Zone A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or a logical entity that performs a specific function (a function zone).

APPENDIX B

Unmanaged ScreenOS Commands

Network and Security Manager (NSM) is designed for system-level management, enabling multiple administrators to manage their devices from one central location using the majority of CLI commands available in ScreenOS. However, a small number of device commands are unmanaged from the NSM UI.

Most unmanaged commands are useful only when performing device administration on a specific device, and do not affect management capabilities (although future versions of NSM may support these commands). To use an unmanaged device command, you must connect locally to the Juniper Networks security device.

[Table 122 on page 907](#) details each unmanaged command.

Table 122: Unmanaged Commands for Firewall/VPN Devices

common-criteria	This command disable all internal commands. Only the root administrator can set this command. If someone other than the root administrator tries to set this command, the security device displays an error message.
envvar	These commands define environment variables. Security devices use environment variables to make special configurations at startup.
gate	This command checks the number of gates on a security device, how many are in use, and how many are still available. Gates are logical access points in the firewall for FTP and similar applications. Security devices create the gates, then convert a gate for each new session when data traffic occurs.
ike	These commands define the Phase 1 and Phase 2 proposals and the gateway for an AutoKey IKE (Internet Key Exchange) VPN tunnel, and specify other IKE parameters.
intervlan-traffic	These commands configure inter-VLAN traffic through a security device. It is possible to configure a virtual system (vsys) with two trusted interfaces, such that traffic can enter the vsys through one interface and exit through the other without undergoing any security services such as authentication or encryption. This is known as inter-VLAN traffic.

Table 122: Unmanaged Commands for Firewall/VPN Devices (continued)

set interface <name> dhcp client settings lease <minute>	This command configures settings for DHCP client lease time.
set log audit-loss-mitigation	This command configures logging to mitigate message loss due to memory limitations on a security device. Used for common criteria only.
set mac	This command configures a static Media Access Control (MAC) address for a security device interface.
timer	These commands display timer settings, or configure a security device to automatically execute management or diagnosis at a specified time. All timer settings remain in the configuration script after the specified time has expired.
user	These commands create, remove, or display entries in the internal user authentication database.
vr nsrp-config-sync	This command unsets synchronization for a specific virtual router in an NSRP cluster.

APPENDIX C

SurfControl Web Categories

SurfControl servers maintain a database of millions of sites organized into about 40 categories. [Table 123 on page 909](#) contains a list of the categories maintained by SurfControl and a description of the URLs in each category.

Table 123: SurfControl Web Categories

Adult/Sexually Explicit	<ul style="list-style-type: none">• Adult products including sex toys, CD-ROMs, and videos• Adult services including video conferencing, escort services, and strip clubs• Erotic stories and textual descriptions of sexual acts• Explicit cartoons and animation• Online groups, including newsgroups and forums, that are sexually explicit in nature• Sexually-oriented or erotic full or partial nudity• Depictions or images of sexual acts, including animals or inanimate objects used in a sexual manner• Sexually exploitative or sexually violent text or graphics• Bondage, fetishes, genital piercing• Nudist sites that feature nudity• Erotic or fetish photography, which depicts nudity <p>NOTE: <i>We do not include sites regarding sexual health, breast cancer, or sexually transmitted diseases (except in graphic examples).</i></p>
Advertisements	<ul style="list-style-type: none">• Banner Ad Servers
Arts and Entertainment	<ul style="list-style-type: none">• Television, movies, music and video programming guides• Downloadable (non-streaming) movie, video or sound clips• Discussion forums on television, movies, music and videos• Online magazines and reviews on the entertainment industry• Celebrity fan sites• Horoscopes• Online greeting cards• Jokes, comics, comic books, comedians or any site designed to be funny or satirical• Circuses, theatre, variety magazines, and radio• Broadcasting firms and technologies (satellite, cable)• Book reviews and promotions, publishing houses, and poetry• Museums, galleries, artist sites (included sculpture, photography)

Table 123: SurfControl Web Categories (continued)

Chat	<ul style="list-style-type: none"> • Web-based chat
Computing and Internet	<ul style="list-style-type: none"> • Reviews, information, buyer's guides of computers, computer parts and accessories, and software • Computer/software/Internet companies, industry news and magazines • Personal storage or backup • Pay-to-Surf sites • Freeware, shareware, and software downloads • Clipart, fonts and animated gif pages • Downloadable mobile phone/ PDA games, themes, graphics, and ringtones • Online photo albums/ digital photo exchange
Criminal Skills	<ul style="list-style-type: none"> • Advocating, instructing, or giving advice on performing illegal acts such as phone, service theft, evading law enforcement, lock-picking, fraud, and burglary techniques • Plagiarism/cheating, including the sale of research papers
Drugs, Alcohol and Tobacco	<ul style="list-style-type: none"> • Recipes, instructions or kits for manufacturing or growing illicit substances, including alcohol, for purposes other than industrial usage • Glamorizing, encouraging, or instructing on the use of or masking the use of alcohol, tobacco, illegal drugs, or other substances that are illegal to minors • Alcohol and tobacco promotional Web sites • Information on "legal highs": glue sniffing, misuse of prescription drugs or abuse of other legal substances • Distributing alcohol, illegal drugs, or tobacco free or for a charge • Displaying, selling, or detailing use of drug paraphernalia <p>NOTE: We do not include sites that discuss medicinal drug use, industrial hemp use, or public debate on the issue of legalizing certain drugs. Nor do we include sites sponsored by a public or private agency that provides educational information on drug use.</p>
Education	<ul style="list-style-type: none"> • Educational institutions, including pre-, elementary, secondary, and high schools; universities. • Educational sites: pre-, elementary, secondary, and high schools; universities. • Distance education and trade schools, including online courses. • Online teacher resources (lesson plans)
Finance and Investment	<ul style="list-style-type: none"> • Stock quotes, stock tickers, and fund rates • Online stock or equity trading • Online banking and bill-pay services • Investing advice or contacts for trading securities • Money management/investment services or firm • General finances and companies that advise thereof • Accountancy, actuaries, banks, mortgages, and general insurance companies
Food and Drink	<ul style="list-style-type: none"> • Recipes, cooking instruction and tips, food products, and wine advisors • Restaurants, cafes, eateries, pubs, and bars • Food/drink magazines, reviews

Table 123: SurfControl Web Categories (continued)

Gambling	<ul style="list-style-type: none"> • Online gambling or lottery web sites that invite the use of real or virtual money • Information or advice for placing wagers, participating in lotteries, gambling, or running numbers • Virtual casinos and offshore gambling ventures • Virtual sports leagues and sports picks and betting pools <p>NOTE: <i>Casino/Hotel/Resort sites that do not feature online gambling or provide gaming tips are categorized under Travel.</i></p>
Games	<ul style="list-style-type: none"> • Game playing or downloading; game hosting or contest hosting • Tips and advice on games or obtaining cheat codes ("cheatz") • Journals and magazines dedicated to game playing
Glamour and Intimate Apparel	<ul style="list-style-type: none"> • Lingerie, negligee or swimwear modeling • Model fan pages; fitness models/sports celebrities • Fashion or glamour magazines online • Beauty and cosmetics • Modeling information and agencies
Government and Politics	<ul style="list-style-type: none"> • Government services such as taxation, armed forces, customs bureaus, emergency services. • Local government sites • Political debate, canvassing, election information and results • Local, national, and international political sites
Hacking	<ul style="list-style-type: none"> • Promotion, instruction, or advice on the questionable or illegal use of equipment and/or software for purpose of hacking passwords, creating viruses, gaining access to other computers and/or computerized communication systems. • Sites that carry malicious executables or viruses • Sites that provide instruction or work-arounds for our filtering software • Cracked software and information sites • Pirated software and multimedia download sites • Sites that provide or promote parasites, including Spyware, Adware and other unsolicited commercial software

Table 123: SurfControl Web Categories (continued)

Hate	<ul style="list-style-type: none"> • Advocating or inciting degradation or attack of specified populations or institutions based on associations such as religion, race, nationality, gender, age, disability, or sexual orientation • Promoting a political or social agenda that is supremacist in nature and exclusionary of others based on their race, religion, nationality, gender, age, disability, or sexual orientation • Holocaust revisionist/denial sites • Coercion or recruitment for membership in a gang* or cult** • Militancy, extremist • Flagrantly insensitive or offensive material <p>NOTE: We do not include news, historical, or press incidents that may include the above criteria (except in graphic examples).</p> <p>**A gang is defined as: a group whose primary activities are the commission of felonious criminal acts, which has a common name or identifying sign or symbol, and whose members individually or collectively engage in criminal activity in the name of the group.</p> <p>**A cult is defined as: a group whose followers have been deceptively and manipulatively recruited and retained through undue influence such that followers' personalities and behavior are altered. Leadership is all-powerful, ideology is totalistic, and the will of the individual is subordinate to the group. Sets itself outside of society.</p>
Health and Medicine	<ul style="list-style-type: none"> • General health such as fitness and wellbeing • Medical information about ailments, conditions, and drugs • Medical reference • Medical procedures, including elective and cosmetic surgery • Alternative and complementary therapies • Prescription medicines • Hospital, medical insurance • Dentistry, optometry, and other medical-related sites • General psychiatry and mental wellbeing sites • Promoting self-healing of physical and mental abuses, ailments, and addictions • Psychology, self-help books, and organizations
Hobbies and Recreation	<ul style="list-style-type: none"> • Recreational pastimes such as collecting, gardening, kit airplanes • Outdoor recreational activities such as hiking, camping, rock climbing • Tips or trends focused on a specific art, craft, or technique Online publications on a specific pastime or recreational activity • Online clubs, associations or forums dedicated to a hobby • Traditional (board, card) games and their enthusiasts • Animal/ pet related sites, including breed-specific sites, training, shows and humane societies
Hosting Sites	<ul style="list-style-type: none"> • Web sites that host business and individuals' web pages (such as GeoCities, earthlink.net, AOL)
Job Search and Career Development	<ul style="list-style-type: none"> • Employment agencies, contractors, job listings, career information • Career searches, career-networking groups
Kid's Sites	<ul style="list-style-type: none"> • Child oriented sites and sites published by children

Table 123: SurfControl Web Categories (continued)

Lifestyle and Culture	<ul style="list-style-type: none"> • Homelife and family-related topics, including parenting tips, gay/lesbian/bisexual (non-pornographic sites), weddings, births, and funerals • Foreign cultures, socio-cultural information
Motor Vehicles	<ul style="list-style-type: none"> • Car reviews, vehicle purchasing or sales tips, parts catalogs • Auto trading, photos, discussion of vehicles including motorcycles, boats, cars, trucks and RVs • Journals and magazines on vehicle modification, repair, and customizing • Online automotive enthusiast clubs
News	<ul style="list-style-type: none"> • Newspapers online • Headline news sites, newswire services, and personalized news services • Weather sites
Personals and Dating	<ul style="list-style-type: none"> • Singles listings, matchmaking and dating services • Advice for dating or relationships; romance tips and suggestions
Photo Searches	<ul style="list-style-type: none"> • Sites that provide resources for photo and image searches
Real Estate	<ul style="list-style-type: none"> • Home, apartment, and land listings • Rental or relocation services • Tips on buying or selling a home • Real estate agents • Home improvement and inspection sites
Reference	<ul style="list-style-type: none"> • Personal, professional, or educational reference • Online dictionaries, maps, and language translation sites • Census, almanacs, and library catalogues • Topic-specific search engines
Religion	<ul style="list-style-type: none"> • Churches, synagogues, and other houses of worship • Any faith or religious beliefs, including nontraditional religions such as Wicca and witchcraft
Remote Proxies	<ul style="list-style-type: none"> • Remote proxies or anonymous surfing • Web-based translation sites that circumvent filtering • Peer-to-peer sharing
Sex Education	<ul style="list-style-type: none"> • Pictures or text advocating the proper use of contraceptives • Sites relating to discussion about the use of the Pill, IUDs and other types of contraceptives • Discussion sites on how to talk to your partner about diseases, pregnancy and respecting boundaries <p>NOTE: Not included in the category are commercial sites that sell sexual paraphernalia. These sites are typically found in the Adult category.</p>
Search Engines	<ul style="list-style-type: none"> • General search engines (Yahoo, AltaVista, Google)

Table 123: SurfControl Web Categories (continued)

Shopping	<ul style="list-style-type: none"> • Online auctions • Department stores, retail stores, company catalogs and other sites that allow online consumer shopping • Online downloadable product warehouses; specialty items for sale • Freebies or merchandise giveaways
Sports	<ul style="list-style-type: none"> • Team or conference web sites • National, international, college, professional scores and schedules • Sports-related online magazines or newsletters
Streaming Media	<ul style="list-style-type: none"> • Streaming media files or events (any live or archived audio or video file) • Internet TV and radio • Personal (non-explicit) webcam sites • Telephony sites that allow users to make calls via the Internet
Travel	<ul style="list-style-type: none"> • Airlines and flight booking agencies • Accommodation information • Travel package listings • City guides and tourist information • Weather bureaus • Car Rentals
Usenet News/Forums	<ul style="list-style-type: none"> • Newsgroups • Opinion or discussion forums • Weblog (blog) sites
Usenet News/Forums	<ul style="list-style-type: none"> • Newsgroups • Opinion or discussion forums • Weblog (blog) sites
Violence/Offensive	<ul style="list-style-type: none"> • Portraying, describing or advocating physical assault against humans, animals, or institutions • Depictions of torture, mutilation, gore, or horrific death • Advocating, encouraging, or depicting self-endangerment, or suicide, including through eating disorders or addictions • Instructions, recipes or kits for making bombs or other harmful or destructive devices • Excessive use of profanity or obscene gesticulation • Sites promoting terrorism • Excessively violent sports or games • Offensive or violent language or satire <p>NOTE: We do not block news, historical, or press incidents that may include the above criteria (except in graphic examples).</p>

Table 123: SurfControl Web Categories (continued)

Weapons	<ul style="list-style-type: none">• Online purchasing or ordering information, including lists of prices and dealer locations• Any page or site predominantly containing, or providing links to, content related to the sale of guns, weapons, ammunition or poisonous substances• Displaying or detailing the use of guns, weapons, ammunition or poisonous substances• Clubs which offer training on machine guns, automatics and other assault weapons and/or sniper training <p>NOTE: Weapons are defined as something (as a club, knife, or gun) used to injure, defeat, or destroy.</p>
Web-based E-mail	<ul style="list-style-type: none">• Web-based e-mail accounts• Messaging sites

APPENDIX D

Common Criteria EAL2 Compliance

This appendix describes actions required for a security administrator to properly secure the Network and Security Manager (NSM) system and NSM User Interface to be in compliance with the Common Criteria EAL2 security target for Juniper Networks IDP 4.0 functionality.

The NSM system consists of the Device Server and the GUI Server; the NSM User Interface is a client application used to access information stored in the NSM system.

- [Guidance for Intended Usage on page 917](#)

Guidance for Intended Usage

- The NSM system must be installed on dedicated systems. These dedicated systems must not contain user processes that are not required to operate the NSM software.

Guidance for Personnel

- There must be one or more competent individuals assigned to manage the NSM system and User Interface (UI), and the security of the information that they contain.
- The authorized administrators must not be careless, willfully negligent, or hostile and must follow and abide by the instructions provided by the NSM documentation.
- The NSM system and UI must be accessed only by authorized users.

Guidance for Physical Protection

- The processing resources of the NSM system and UI must be located within facilities with controlled access that prevents unauthorized physical access.

APPENDIX E

Log Entries

This appendix lists the log entry subcategories for the following log entry categories:

- [Screen Alarm Log Entries on page 919](#)
- [Alarm Log Entries on page 921](#)
- [Deep Inspection Alarm Log Entries on page 922](#)
- [Configuration Log Entries on page 997](#)
- [Information Log Entries on page 999](#)
- [Self Log Entries on page 1001](#)
- [Traffic Log Entries on page 1001](#)
- [GTP Log Entries on page 1002](#)

Screen Alarm Log Entries

The Screen category contains the subcategories shown in [Table 124 on page 919](#):

Table 124: Screen Alarm Log Entries

Attack	ScreenOS Message ID
Address Sweep Attack	Attacks > Alert > 00017
Block ActiveX component	Attacks > Critical > 00434
Block EXE component	Attacks > Critical > 00433
Block IP fragment traffic	Attacks > Critical > 00429
Block JAVA component	Attacks > Critical > 00432
Block ZIP component	Attacks > Critical > 00431
Destination IP session limit	Attacks > Critical > 00430
ICMP Flood Attack	Attacks > Alert > 00011
IDS ICMP Fragment	Attacks > Critical > 00422

Table 124: Screen Alarm Log Entries (continued)

Attack	ScreenOS Message ID
IDS ICMP too large	Attacks > Critical > 00436
IDS IP Bad Options	Attacks > Critical > 00415
IDS IP unknown port	Attacks > Critical > 00414
IDS SYN Fragment	Attacks > Critical > 00412
IDS TCP FIN No ACK	Attacks > Critical > 00438
IDS TCP SYN FIN	Attacks > Critical > 00437
IDS TCP No Flag	Attacks > Critical > 00413
IP Source Route Attack	Attacks > Alert > 00009
IP Spoof Attack	Attacks > Alert > 00008
Land Attack	Attacks > Alert > 00010
Malicious URL Protection	Attacks > Critical > 00032
Multiple Authentications Failed	Auth > Alert > 00003
Ping of Death Attack	Attacks > Emergency > 00007
Policy Denied	Policies > Alert > 00018
Port Scan Attack	Attacks > Alert > 00016
SYN Attack	Attacks > Emergency > 00005
SYN Flood	
SYN ACK	
SYN MAC	
SYN-ACK-ACK proxy DoS	Attacks > Critical > 00439
Source IP session limit	Attacks > Critical > 00033
Tear Drop Attack	Attacks > Emergency > 00006
UDP Flood Attack	Attacks > Alert > 00012
VPN Replay Detected	IKE > Critical > 00042

Table 124: Screen Alarm Log Entries (continued)

Attack	ScreenOS Message ID
Winnuke Attack	Attack > Alert > 00004

Alarm Log Entries

The Alarm category contains the subcategories shown in [Table 125 on page 921](#):

Table 125: Alarm Log Entries

Alarm Log Entry SubCategories	ScreenOS Message ID
Admin	Admin > Alert > 00027
Anti Virus - CSP	AntiVirus Scanning (External) > Error > 52
BGP Alarm	BGP > Alert > 00206
CPU Usage High	Logging > Critical > 00030
DHCP	DHCP > Alert > 00029 DHCP > Critical > 00029
DNS Host	DNS > Critical > 00021
Interface Failover	Interface > Critical > 00090
Hardware	Device > Critical > 00022
IP Conflict	ARP > Critical > 00031
Log Overflow	Logging > Critical > 00024
Memory Low	Device > Critical > 00020 Logging > Critical > 00020
NSRP Inconsistent Config	High Availability > 00015
NSRP IP DUP Master	High Availability > 00015
NSRP RTO DOWN	High Availability > 00015
NSRP RTO Duplicate	High Availability > 00015
NSRP RTO UP	High Availability > 00015
NSRP Status	High Availability > Critical > 00015

Table 125: Alarm Log Entries (continued)

Alarm Log Entry SubCategories	ScreenOS Message ID
NSRP TRACKIP Failed	High Availability > 00062
NSRP TRACKIP Failover	High Availability > 00062
NSRP VSD 2nd Path Reply	High Availability > Critical > 00077
NSRP VSD 2nd Path REQ	High Availability > Critical > 00076
NSRP VSD Backup	High Availability > Critical > 00073
NSRP VSD Ineligible	High Availability > Critical > 00074
NSRP VSD Init	High Availability > Critical > 00070
NSRP VSD Inoperable	High Availability > Critical > 00075
NSRP VSD Master	High Availability > Critical > 00071
NSRP VSD Pbackup	High Availability > Critical > 00072
OSPF Packet Flood	OSPF > Critical > 00206
RIP Packet Flood	RIP > Critical > 207
Route add/delete Error	OSPF > Critical > 200
Route RIP Updated Flood	RIP > Critical > 00207
Exceeded Route Entry (Sys)	Route > Critical > 00200
Secure Shell	SSH > Critical > 00034
URL Blk	WEB Filtering > Alert > 00014
VIP Svr Down	VIP > Critical > 00023
VPN	IKE > Alert > 00026
VPN Down	VPN > Critical > 00041
VPN Up	VPN > Critical > 00040

Deep Inspection Alarm Log Entries

The Deep Inspection Alarm category contains the subcategories shown in [Table 126 on page 923](#):

Table 126: Deep Inspection Alarm Log Entries

Attack Name	Attack Description	Severity	Versions
APP:CURL-OF-BANNER	This signature detects buffer overflow attempts against the cURL file retrieval client. cURL 6.1 to 7.4 versions are vulnerable. Attackers may use a malicious server to connect to the cURL client and execute arbitrary code with the permissions of the cURL user.	high	sos5.1.0
CHAT:AIM:MESSAGE-SEND	This signature detects messages sent from AIM clients to other AIM clients.	info	sos5.1.0
CHAT:AUDIT:AIM:INVALID-TLV	This protocol anomaly is a AIM message with an invalid TLV; the TLV data specified in the FLAP header is less than the actual data in the TLV header.	info	sos5.1.0
CHAT:AUDIT:AIM:INV-TLV-LEN	This protocol anomaly is a AIM message with an invalid TLV; the TLV length is less than expected, or the TLV length is greater than the data specified in the FLAP header.	info	sos5.1.0
CHAT:AUDIT:MSN:GROUP-NAME	This protocol anomaly is an MSN message with a group name length that exceeds the user-defined maximum. The default group name maximum is 64.	info	sos5.1.0
CHAT:AUDIT:YMSG:FILE-SEND	This signature detects a Yahoo Messenger client sending a file to another user.	info	sos5.1.0
CHAT:AUDIT:YMSG:MAIL-ADDR	This protocol anomaly is a Yahoo! Messenger e-mail address that exceeds the user-defined maximum. A Yahoo! Messenger server sends an e-mail address as part of a new e-mail alert message. The default number of bytes in an Yahoo! Messenger e-mail address is 84.	info	sos5.1.0
CHAT:AUDIT:YMSG:MSG-TOO-BIG	This protocol anomaly is a Yahoo! Messenger message that exceeds the user-defined maximum. The default number of bytes in an Yahoo! Messenger message is 8192.	info	sos5.1.0
CHAT:AUDIT:YMSG:OFLOW-GRP-NAME	This protocol anomaly is a Yahoo! Messenger group name that exceeds the user-defined maximum. Yahoo! Messenger clients use groups to separate their friends into categories. The default number of bytes in an Yahoo! Messenger group name is 84.	info	sos5.1.0
CHAT:AUDIT:YMSG:OFLOW-PASSWD	This protocol anomaly is a Yahoo! Messenger encrypted password that exceeds the user-defined maximum. The Yahoo! Messenger client sends an encrypted password to the server as part of the authentication process. The default number of bytes in an Yahoo! Messenger encrypted password is 1024.	info	sos5.1.0
CHAT:MSN:ACCESS	This signature detects MSN Messenger chat using the specified content type "text/plain" on port 1863 (default port of MSN Messenger).	info	sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
CHAT:MSN:LOGIN-ATTEMPT	This signature detects attempts to login to the MSN network using an MSN Messenger client.	info	sos5.1.0
DB:MS-SQL:SQLXML-ISAPI-OF	This signature detects buffer overflow attempts against the SQLXML-ASAPI Extension in Microsoft SQL Server 2000. The SQLXML-ASAPI extension handles data queries over HTTP (SQLXML HTTP); attackers may connect to the target host and submit maliciously crafted data to create a buffer overflow.	high	sos5.1.0
DNS:AUDIT:CLASS-NON-IN	This protocol anomaly is a DNS request/reply in which the question/resource address class is not IN (Internet Address). Although allowed by the RFC, this should happen only in rare circumstances and may indicate an exploit attempt.	info	sos5.1.0
DNS:AUDIT:QCLASS-UNEXP	This protocol anomaly is a DNS reply with a resource specifying a CLASS ID reserved for queries only (QCLASS). This may indicate an exploit attempt.	info	sos5.1.0
DNS:AUDIT:REP-QTYPE-UNEXPECTED	This protocol anomaly is a DNS reply with a resource specifying a TYPE ID reserved for queries only (QTYPE). This may indicate an exploit attempt.	info	sos5.1.0
DNS:AUDIT:REP-S2C-QUERY	This protocol anomaly is a DNS reply with a query/reply bit (QR) that is unset (indicating a query). This may indicate an exploit attempt.	info	sos5.1.0
DNS:AUDIT:REQ-C2S-RESPONSE	This protocol anomaly is a DNS request with a query/reply bit (QR) set (indicating a reply). This may indicate an exploit attempt.	info	sos5.1.0
DNS:AUDIT:REQ-INVALID-HDR-RA	This protocol anomaly is a client-to-server DNS message with the recursion-available bit (RA) set. This may indicate an exploit attempt.	info	sos5.1.0
DNS:AUDIT:TYPE-ANY	This protocol anomaly is a DNS request with request type set to "ANY".	info	sos5.1.0
DNS:EXPLOIT:EMPTY-UDP-MSG	This protocol anomaly is an empty DNS UDP message. This may indicate an exploit attempt.	high	sos5.0.0, sos5.1.0
DNS:EXPLOIT:EXPLOIT-BIND9-RT	This protocol anomaly is an rdataset parameter to the dns_message_findtype() function in message.c that is not NULL. In BIND 9 (up to 9.2.0), attackers may cause a shutdown on an assertion failure. Note: Common queries in routine operations (such as SMTP queries) may trigger this anomaly.	high	sos5.0.0, sos5.1.0
DNS:EXPLOIT:POINTER-LOOP	This protocol anomaly is a DNS message with a set of DNS pointers that form a loop. This may indicate a denial-of-service (DoS) attempt.	high	sos5.0.0, sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
DNS:EXPLOIT:REQUEST-SHORT-MSG	This protocol anomaly is a DNS message that ended prematurely. This may indicate an exploit attempt.	high	sos5.1.0
DNS:EXPLOIT:TYPE-AXFR	This protocol anomaly is a zone transfer attempt. This may indicate an attempt to obtain information about an entire domain.	medium	sos5.0.0, sos5.1.0
DNS:HEADERERROR:INVALID-OPCODE	This protocol anomaly is a DNS request/reply with an invalid value in the header OPCODE field. This may indicate an exploit attempt.	medium	sos5.0.0, sos5.1.0
DNS:OVERFLOW:FF-FF-BIN	This signature detects attempts to create buffer overflows. Attackers may send maliciously crafted packets to DNS servers to overflow the buffer and gain root access.	critical	sos5.0.0, sos5.1.0
DNS:OVERFLOW:INVALID-LABEL-LEN	This protocol anomaly is a DNS request/reply with a label that exceeds the maximum length (63) specified in the RFC. This may indicate a buffer overflow attempt.	critical	sos5.0.0, sos5.1.0
DNS:OVERFLOW:INVALID-POINTER	This protocol anomaly is a DNS request/reply with a pointer that points beyond the end of the data. This may indicate a buffer overflow or denial-of-service (DoS) attempt.	critical	sos5.0.0, sos5.1.0
DNS:OVERFLOW:NAME-TOO-LONG	This protocol anomaly is a DNS name that exceeds 255 characters. This may cause problems for some DNS servers.	critical	sos5.0.0, sos5.1.0
DNS:OVERFLOW:NXT-OVERFLOW	This protocol anomaly is a suspiciously large NXT resource record in a DNS transaction. BIND versions 8.2 through 8.2.1 are vulnerable to a buffer overflow in the processing of NXT resource records.	critical	sos5.1.0
DNS:OVERFLOW:OPT-DOS	This protocol anomaly is a suspiciously long OPT resource record. All versions of BIND up to version 8.3.3 are vulnerable to a denial of service attack. An attacker can crash the server by requesting a subdomain that does not exist with an OPT resource record that has a very large UDP payload size.	critical	sos5.1.0
DNS:OVERFLOW:OVERSIZED-UDP-MSG	This protocol anomaly is a DNS UDP-based request/reply that exceeds the maximum length (512) specified in RFC. This may indicate a buffer overflow attempt.	high	sos5.1.0
DNS:OVERFLOW:SIG-OVERFLOW	This protocol anomaly is a TCP-based DNS transaction with a suspiciously small SIG resource record. Bind versions 8 to 8.3.3 are vulnerable to a heap overflow in the code that handles SIG resource records. Attackers may execute arbitrary code on the server.	critical	sos5.1.0
DNS:OVERFLOW:TOO-LONG-TCP-MSG	This protocol anomaly is a DNS TCP-based request/reply that exceeds the maximum length specified in the message header. This may indicate a buffer overflow or an exploit attempt.	high	sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
DNS:QUERY:NULL-QUERY	This protocol anomaly is a DNS request with the question, answer, additional, and name server counts are zero. This can indicate a malicious user trying to crash the DNS server.	high	sos5.1.0
DNS:QUERY:VERSION-QUERY	This protocol anomaly is a DNS query for version.bind with the type set to TXT and the class set to CHAOS. BIND servers support the ability to be remotely queried for their versions. This can indicate a reconnaissance attempt; when attackers know the BIND version, they can then attempt to exploit vulnerabilities on the server.	medium	sos5.1.0
DOS:NETDEV:CISCO-HTTPD-DOS	This signature detects attempts to exploit a vulnerability in Cisco IOS. Versions prior to 11.0, 11.2.8SA1, 12.1(1a)T1, and 12.1(1.3)T are susceptible. Attackers may remotely request URLs containing the %% string from the IP HTTP server, causing the router to crash/reboot/power cycle.	high	sos5.0.0, sos5.1.0
DOS:NETDEV:CISCO-RTR-DOS	This signature detects denial-of-service (DoS) attempts against Cisco (routers). Cisco has identified multiple affected versions of IOS and customers are advised to check with their vendor or on Cisco's Web site for information. Attackers may send invalid HTTP traffic to a Cisco IOS device to cause a DoS on the device.	medium	sos5.1.0
DOS:NETDEV:LINKSYS-GOZILA-DOS2	This signature detects attempts to exploit a vulnerability in a LinkSys Cable/DSL router. Attackers may submit an overly long sysPasswd parameter within a malicious HTTP request to crash a LinkSys Cable/DSL router.	medium	sos5.1.0
DOS:NETDEV:LINKSYS-GOZILA-DOS3	This signature detects attempts to exploit a vulnerability in a LinkSys Cable/DSL router. Attackers may submit an overly long DomainName parameter within a malicious HTTP request to crash a LinkSys Cable/DSL router.	medium	sos5.1.0
DOS:NETDEV:NETWORK-3COM-DOS	This signature detects attempts to exploit a firmware vulnerability in the 3COM OfficeConnect 812 and 840 DSL/ADSL routers. OCR812 versions 1.1.9 and earlier are susceptible. Attackers may remotely request long strings from the HTTP daemon, making the router reboot/power cycle and creating a denial-of-service (DoS).	high	sos5.0.0, sos5.1.0
DOS:NETDEV:WEBJET-FRAMEWORK	This signature detects attempts to exploit a vulnerability in HP Web JetAdmin service. Web JetAdmin version 6.5 is vulnerable. Attackers may access sensitive configuration information. If you run an HP Web JetAdmin server on your network, configure DI to monitor the server port that is configured to listen; by default, the listening port is TCP/8000.	medium	sos5.1.0
DOS:NETDEV:WEBJET-FW-INFOLEAK	This signature detects attempts to exploit a vulnerability in HP Web JetAdmin service. Web JetAdmin version 6.5 is vulnerable. Attackers may access sensitive configuration information.	medium	sos5.0.0, sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
DOS:NETDEV:WEBJET-TRAVERSAL	This signature detects directory traversal attempts against HP Web JetAdmin service. HP Web JetAdmin version 7.5.2546 and earlier are vulnerable. Because JetAdmin does not properly verify input to the setinclude parameter in <code>/plugins/hpjdwm/script/test/setinfo.hts</code> , attackers may use a directory traversal to read and execute arbitrary HTS files.	high	sos5.1.0
DOS:NETDEV:WEBJET-WRITETOFILE	This signature detects attempts to exploit a vulnerability in HP Web JetAdmin service. Web JetAdmin versions 7.x are vulnerable. Attackers may send a maliciously formatted request to a Web JetAdmin script to execute arbitrary commands on the server.	critical	sos5.1.0
FTP:COMMAND:PLATFTP-CD-DOS	This signature detects attempts to exploit a vulnerability in PlatinumFTP. Attackers may submit a maliciously crafted pathname in a CD request to crash the FTP daemon. PlatinumFTP 1.0.6 and earlier versions are vulnerable.	medium	sos5.0.0, sos5.1.0
FTP:COMMAND:SITE-EXEC	This signature detects attempts to exploit a configuration vulnerability in wuFTPd. Version 2.4.1 is susceptible. <code>pathnames.h</code> sets <code>_PATH_EXECPATH</code> to <code>/bin</code> , which is relative to <code>~ftp</code> for anonymous users, but relative to <code>/</code> for users with accounts (specifying the actual <code>/bin</code> rather than <code>~ftp/bin</code>). Attackers may establish an FTP account on the system and run the site exec command to gain access to the <code>/bin</code> directory.	medium	sos5.0.0, sos5.1.0
FTP:DIRECTORY:DOT-DOT	This signature detects <code>'../.'</code> FTP commands sent to FTP/21. Attackers may change the directory to the root directory of the FTP service, and gain access to the system.	medium	sos5.1.0
FTP:DIRECTORY:MSIE-FTP-DIRTRAV	This signature detects a Microsoft Internet Explorer client attempting to download a file from a malicious server. The server may embed a directory traversal attack in the filename to specify the exact file download location on the client machine.	medium	sos5.1.0
FTP:EXPLOIT:BOUNCE-ATTACK	This protocol anomaly is an FTP bounce attack. There are two possibilities: a PORT command specified an IP address different from the client address, or a PASV command resulted in a 227 message with an IP address different than the server.	high	sos5.0.0, sos5.1.0
FTP:EXPLOIT:FTPBIN-WRITEABLE	This signature detects an attempt by a malicious attacker to upload files with the names of common binaries to the FTP server's <code>/bin</code> directory. Successful exploitation of this vulnerability may result in the attacker being able to execute arbitrary code on the victim ftp server, including the reading of sensitive files outside of the ftp server's path.	medium	sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
FTP:EXPLOIT:ILLEGAL-PORT	This protocol anomaly is an FTP PORT command/response to a PASV command ("227...") that specifies a reserved port number. This may indicate an attempt to make the firewall open reserved ports.	high	sos5.0.0, sos5.1.0
FTP:EXPLOIT:OPENFTPD-MSG-FS	This signature detects attempts to exploit a format string vulnerability in the OpenFTP daemon.	critical	sos5.1.0
FTP:EXPLOIT:SYNTAX-ERROR	This protocol anomaly is a syntax error in an FTP command/response, such as a malformed PORT command or 227 response. This may indicate an exploit attempt.	high	sos5.0.0, sos5.1.0
FTP:EXPLOIT:TYPESOFT-DOS	This signature detects denial-of-service (DoS) attempts against TypSoft FTP Server. TypSoft FTP Server 1.10 and earlier versions are vulnerable. Attackers may send known malicious FTP path strings to exhaust all system resources and crash a TypSoft FTP Server.	high	sos5.0.0, sos5.1.0
FTP:EXPLOIT:WIN32-WFTPD-BOF	This signature detects invalid LIST, NLST, and STAT commands. WS-FTPD for Windows (trial versions 3.20 and 3.21, Pro and Standard) contains a vulnerability in the command parser that may allow malicious users to crash the service or execute arbitrary code.	high	sos5.1.0
FTP:FILE:FTP-PUT-AUTOEXECBAT	This signature detects an attempt by an attacker to exploit a directory traversal vulnerability in the SunFTP daemon. Successful exploitation of this vulnerability may allow an attacker to read and write to files outside of the daemon's directory structure. This vulnerability is present in SunFTP build 9.	medium	sos5.1.0
FTP:MS-FTP:ASTERISK	This signature detects denial-of-service (DoS) attempts against Microsoft FTP Service in Microsoft IIS 4.0 and 5.0. Attackers who have previously established an FTP session may send glob characters within a maliciously crafted NLST request to crash the server.	medium	sos5.0.0, sos5.1.0
FTP:MS-FTP:STAT-GLOB	This signature detects denial-of-service (DoS) attempts against Microsoft FTP Service in Microsoft IIS 4.0 and 5.0. Attackers who have previously established an FTP session may send glob characters within a maliciously crafted status request to crash the server.	medium	sos5.0.0, sos5.1.0
FTP:OVERFLOW:BSD-FTPD-MKD-OF	This signature detects buffer overflow attempts against the FTPD that ships with early versions of FreeBSD 4.x and OpenBSD 2.8. FTPD 6.00LS and 6.5/OpenBSD versions are vulnerable. Attackers may gain local host access and root permissions.	critical	sos5.0.0, sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
FTP:OVERFLOW:FREEBSD-FTPD-GLOB	This signature detects buffer overflow attempts against the FreeBSD FTP daemon. FreeBSD-4.2 is vulnerable. Attackers may submit a malicious STAT request that contains file globbing characters to execute arbitrary code on the target host with administrator privileges.	critical	sos5.0.0, sos5.1.0
FTP:OVERFLOW:LINE_TOO_LONG	This protocol anomaly is an incoming FTP line that is too long. This may indicate an attempt to overflow the server.	high	sos5.0.0, sos5.1.0
FTP:OVERFLOW:OPENBSD-X86	This signature detects buffer overflow attempts against ftpd in OpenBSD. OpenBSD versions 2.7 and 2.8, FTP code revisions 1.49 to 1.79 are vulnerable. Attackers with write access may exploit the replydirname() function in BSD-based ftpd daemons to gain root access.	critical	sos5.0.0, sos5.1.0
FTP:OVERFLOW:PASS_TOO_LONG	This protocol anomaly is an FTP client password that exceeds the length threshold. This may indicate a malicious FTP client attempting to overflow the server.	high	sos5.0.0, sos5.1.0
FTP:OVERFLOW:PATH-LINUX-X86-1	This signature detects attempts to exploit a realpath vulnerability in ProFTPD and wuFTPd running on LINUX. Versions ProFTPD 1.2pre1 and earlier and wuFTPd 2.4.2 (beta 18) VR9 and earlier are susceptible. Attackers may gain write access, remotely create long pathnames, and overflow the buffer to gain root access.	critical	sos5.0.0, sos5.1.0
FTP:OVERFLOW:PATH-LINUX-X86-2	This signature detects attempts to exploit a realpath vulnerability in ProFTPD and wuFTPd running on LINUX. Versions ProFTPD 1.2pre1 and earlier and wuFTPd 2.4.2 (beta 18) VR9 and earlier are susceptible. Attackers may gain write access, remotely create long pathnames, and overflow the buffer to gain root access.	critical	sos5.0.0, sos5.1.0
FTP:OVERFLOW:PATH-LINUX-X86-3	This signature detects attempts to exploit a realpath vulnerability in ProFTPD and wuFTPd running on LINUX. Versions ProFTPD 1.2pre1 and earlier and wuFTPd 2.4.2 (beta 18) VR9 and earlier are susceptible. Attackers may gain write access, remotely create long pathnames, and overflow the buffer to gain root access.	critical	sos5.0.0, sos5.1.0
FTP:OVERFLOW:PATH-TOO-LONG	This protocol anomaly is a pathname in an FTP command (RETR, STOR, APPE, SMNT, RNFR, RNTD, DELE, RMD, MKD, STAT, CWD, LIST, NLST) that exceeds the length threshold. This may be an attempt to overflow the server.	high	sos5.0.0, sos5.1.0
FTP:OVERFLOW:SITESTRING-2-LONG	This protocol anomaly is an argument in the FTP SITE command that exceeds the length threshold. This may be an attempt to overflow the server.	high	sos5.0.0, sos5.1.0
FTP:OVERFLOW:USERNAME-2-LONG	This protocol anomaly is a username in an FTP connection that exceeds the length threshold. This may be an attempt to overflow the server.	high	sos5.0.0, sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
FTP:OVERFLOW:WFTPD-MKD-OVERFLOW	This signature detects buffer overflow attempts against the MKD command in Wftpd server 2.34. Attackers may use MKD and CWD commands to create nested directories and execute arbitrary commands with system privileges.	critical	sos5.0.0, sos5.1.0
FTP:OVERFLOW:WUBSD-SE-RACE	This signature detects buffer overflow attempts against the PASS command in Wu-ftp 2.6.0 and BSDi-ftp. Attackers may send a maliciously crafted PASS request to an FTP server to execute arbitrary commands as root.	critical	sos5.0.0, sos5.1.0
FTP:PABLO-FTP:FORMAT-STRING	This signature detects denial-of-service attempts against the Pablo FTP Server. Versions 1.2, 1.3, and 1.5 running on Windows 2000 are vulnerable. Because the FTP server improperly parses format string characters, attackers may supply a maliciously crafted username to execute arbitrary code and crash the server.	high	sos5.0.0, sos5.1.0
FTP:PASSWORD:BRUTE-FORCE	This protocol anomaly is multiple login failures within a short period of time between a unique pair of hosts.	high	sos5.1.0
FTP:PASSWORD:COMMON-PASSWD	This signature detects common passwords used in FTP sessions. Attackers may attempt to log into known accounts using easily guessed passwords.	info	sos5.0.0, sos5.1.0
FTP:PASSWORD:HOTBOX	This signature detects attempts to use the default rootkit password 'h0tb0x' to access a FreeBSD rootkit account. Attackers may gain root access.	high	sos5.0.0, sos5.1.0
FTP:PASSWORD:LRKROX	This signature detects attempts to install the Rootkit hacker utility on a LINUX system. The default password is lrkr0x.	high	sos5.0.0, sos5.1.0
FTP:PASSWORD:SATORI	This signature detects attempts to install the Rootkit lrk4 hacker utility on a system. The default password is satori.	high	sos5.0.0, sos5.1.0
FTP:PASSWORD:WH00T	This signature detects attempts to install the Rootkit hacker utility on a LINUX system. The default password is wh00t.	high	sos5.0.0, sos5.1.0
FTP:PROFTP:LOGXFR-OF1	This signature detects buffer overflow attempts against the log_xfer() function in ProFTPD. This vulnerability affects ProFTPD versions 1.2.0pre1, pre2, and pre3.	critical	sos5.0.0, sos5.1.0
FTP:PROFTP:MKD-OVERFLOW	This signature detects buffer overflow attempts against ProFTPD. Versions 1.2pre3 and earlier are vulnerable. Attackers may send a pathname to the 'MKD' command to gain remote root access.	critical	sos5.0.0, sos5.1.0
FTP:PROFTP:PPC-FS1	This signature detects attempts to exploit a format string vulnerability in ProFTPD. Versions 1.2pre6 and earlier are vulnerable. Attackers may overflow the PWD command.	critical	sos5.0.0, sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
FTP:PROFTP:PPC-FS2	This signature detects attempts to exploit a format string vulnerability in ProFTPD. Versions 1.2pre6 and earlier are vulnerable.	critical	sos5.0.0, sos5.1.0
FTP:PROFTP:PROFTPD-GEN-GLOB-DOS	This signature detects denial-of-service (DoS) attempts against ProFTPD. Because ProFTPD uses inadequate globbing algorithms, attackers may send wildcards in the argument of a maliciously crafted command to DoS the server.	medium	sos5.0.0, sos5.1.0
FTP:PROFTP:SIZE-DOS2	This signature detects attempts to exploit a vulnerability in ProFTPD. Version 1.2.0pre* is vulnerable. Attackers may send multiple SIZE requests with a static pathname to create a denial-of-service (DoS).	high	sos5.0.0, sos5.1.0
FTP:PROFTP:USER-DOS	This signature detects attempts to exploit a vulnerability in ProFTPD. Versions 1.2.0rc* and 1.2.0pre* are vulnerable. Attackers may send a maliciously crafted USER command to create a denial-of-service (DoS).	high	sos5.0.0, sos5.1.0
FTP:REQERR:GNULS-WIDTH-DOS	This signature detects denial-of-service (DoS) attempts against GNU ls. If the FTP daemon uses a vulnerable version of GNU ls, attackers may send an oversized width parameter to GNU ls to cause the server CPU utilization to temporarily reach 100% and exhaust system memory. This condition can persist for several minutes depending on the width specified.	medium	sos5.0.0, sos5.1.0
FTP:REQERR:REQ-MISSING-ARGS	This protocol anomaly is an FTP command with an incomplete argument list, such as a USER command with no user name, a RETR command with no file name, etc. This may indicate command line access to the FTP server or an exploit attempt.	medium	sos5.0.0, sos5.1.0
FTP:SERVU:CHMOD-OVERFLOW	This signature detects attempts to exploit a vulnerability in the ServU FTP server CHMOD command. The CHMOD command is typically used to change the permissions of a file on the server. Attackers may send an overly long filename argument to the CHMOD command to execute arbitrary code with system privileges.	critical	sos5.1.0
FTP:USER:ROOT	This signature detects attempts to login to an FTP server using the "root" account. This may indicate an attacker trying to gain root-level access, or it may indicate poor security practices. FTP typically uses plain-text passwords, and using the root account to FTP could expose sensitive data over the network.	medium	sos5.0.0, sos5.1.0

Table 126: Deep Inspection Alarm Log Entries (continued)

Attack Name	Attack Description	Severity	Versions
FTP:WS-FTP:CPWD	This signature detects buffer overflow attempts against WS FTP Server. The code that handles arguments to the SITE CPWD command, which allows users to change their password, contains an unchecked string copy. Attackers may send a maliciously crafted argument in the SITE CPWD command to overflow the buffer and overwrite the return address.	medium	sos5.0.0, sos5.1.0
FTP:WU-FTP:DELE-OF	This signature detects buffer overflow attempts against the DELE command in a WU-ftp server. WU-ftp versions 2.4 and prior (Academ beta12-18 included) are vulnerable. This may be a variation on the ADM exploit; attackers may log in anonymously using a hardcoded e-mail address as the password.	high	sos5.0.0, sos5.1.0
FTP:WU-FTP:FTPD-BSD-X86	This signature detects attempts to exploit an input validation vulnerability in wuFTPd running on FreeBSD. FreeBSD versions 4.3 and 4.4 are vulnerable. Because user input goes directly into a format string for a *printf function, attackers may overwrite data on a stack (i.e. a return address), access the shellcode pointed to by the overwritten eip, and execute arbitrary commands.	high	sos5.0.0, sos5.1.0
FTP:WU-FTP:GLOBARG	This signature detects attempts to exploit a vulnerability in WU-ftp, a software package that provides File Transfer Protocol (FTP) services for UNIX and Linux systems. WU-ftp versions 2.6.1 to 2.6.18 are vulnerable. Attackers may send a maliciously crafted pathname in a CWD or LIST command to the FTP server to execute arbitrary commands as root.	high	sos5.0.0, sos5.1.0
FTP:WU-FTP:IREPLY-FS	This signature detects attempts to exploit a format string vulnerability in WU-ftp 2.4 running on Solaris 2.8. Attackers may inject malicious code into the WU-ftp daemon memory space; later in the same session, the attacker may exploit a format string vulnerability in the Ireply() function to access that code and execute arbitrary commands as root.	critical	sos5.0.0, sos5.1.0
FTP:WU-FTP:LINUX-OF	This signature detects attempts to exploit an input validation vulnerability in wuFTPd running on LINUX. All versions are susceptible. Because user input goes directly into a format string for a *printf function, attackers may overwrite data on a stack, i.e. a return address, access the shellcode pointed to by the overwritten eip, and execute arbitrary commands. This same attack may be successful seen against ProFTPD servers.	high	sos5.0.0, sos5.1.0
FTP:WU-FTP:REALPATH-OF	This signature detects buffer overflow attempts against the realpath() function in WU-ftp, a software package that provides File Transfer Protocol (FTP) services for UNIX and Linux systems. WU-ftp version 2.5.0 and earlier are vulnerable. Attackers may send a maliciously crafted FTP pathname to overflow a buffer in realpath() and execute arbitrary commands with administrator privileges.	critical	sos5.0.0, sos5.1.0

FTP:WU-FTP:REALPATH-OF2	This signature detects buffer overflow attempts against the <code>realpath()</code> function in <code>Wu-ftpd</code> , a software package that provides File Transfer Protocol (FTP) services for UNIX and Linux systems. <code>Wu-ftpd</code> version 2.5.0 and earlier are vulnerable. Attackers may send a maliciously crafted FTP pathname to overflow a buffer in <code>realpath()</code> and execute arbitrary commands with administrator privileges.	critical	sos5.0.0, sos5.1.0
HTTP:3COM:3COM-PASS-LEAK	This signature detects attempts to access a 3COM wireless router web page that contains sensitive administrative information. No authentication is required to access this page.	high	sos5.0.0, sos5.1.0
HTTP:3COM:ADMIN-LOGOUT	This signature detects direct requests to the logout web service on a 3Com 3crwe754g72-a based device. Attackers that are spoofing a 3Com administrator's IP address may call the logout application to force the administrator to logout.	info	sos5.0.0, sos5.1.0
HTTP:3COM:CONF-DOWNLOAD	This signature detects attempts to download the configuration file from a 3Com 3crwe754g72-a based device. Attackers may use the sensitive information obtained from the configuration file to gain full control over the device.	high	sos5.0.0, sos5.1.0
HTTP:3COM:LOG-CLEAN	This signature detects attempts to cause a 3Com 3crwe754g72-a based device to clear its logs. Attackers may use spoofed IP address to send a log clear request without authenticating.	medium	sos5.1.0
HTTP:APACHE:APACHE-BADIPV6	This signature detects attempts to exploit a vulnerability in Apache Web server. All Apache servers on all platforms running version Apache 2.0.50 and earlier are vulnerable. Using <code>apr-util</code> , attackers may include a crafted IPv6 literal address within an HTTP request to an Apache v2 server to cause the Apache child process to quit. On BSD systems, attackers may also be able to execute arbitrary code.	high	sos5.0.0, sos5.1.0
HTTP:APACHE:APACHE-BADIPV6-2	This signature detects attempts to exploit a vulnerability in Apache Web server. All Apache servers on all platforms running version Apache 2.0.50 and earlier are vulnerable. Using <code>apr-util</code> , attackers may include a crafted IPv6 literal address within an HTTP request to an Apache v2 server to cause the Apache child process to quit. On BSD systems, attackers may also be able to execute arbitrary code.	high	sos5.0.0, sos5.1.0
HTTP:APACHE:CHUNKED-WORM	This signature detects attempts to infect Apache Web servers with the Apache Worm. Apache versions 1.3.26, 2.0.38 and prior are vulnerable. Apache improperly calculates required buffer sizes for chunked encoded requests due to a signed interpretation of an unsigned integer value. The worm sends POST requests containing malicious chunked encoded data to exploit the Apache daemon.	critical	sos5.1.0

HTTP:APACHE:MOD-NTLM-BOFI	This signature detects buffer overflow attempts against Apache Web server. An Apache Web server uses mod_ntlm (an Apache 1.x and 2.x module) to authenticate users against a Microsoft Windows Domain Controller. Attackers may send long or malformed strings to mod_ntlm using the Authorization HTTP header, overflow the buffer, then execute arbitrary code on the Web server.	high	sos5.0.0, sos5.1.0
HTTP:APACHE:MODPHP-UPLOAD-HOF	This signature detects heap overflow attempts against mod_php in Apache. Attackers may send a maliciously crafted HTTP POST request to execute arbitrary code on the server.	critical	sos5.1.0
HTTP:APACHE:NOSEJOB	This signature detects attempts to exploit a vulnerability in Apache Web servers. Apache improperly calculates required buffer sizes for chunked encoded requests due to a signed interpretation of an unsigned integer value. Attackers may send chunked encoded requests with the unique Host header value "Apache-nosejob.c." in the GET request to create a buffer overflow and execute arbitrary code.	critical	sos5.0.0, sos5.1.0
HTTP:APACHE:PHP-INVALID-HDR	This signature detects denial-of-service attempts against the Apache HTTP daemon. PHP versions 4.2.0 and 4.2.1 running on Apache 1.3.26 are vulnerable. Attackers may use invalid headers in an HTTP request to crash the Apache HTTP daemon; the daemon may require a manual restart.	high	sos5.0.0, sos5.1.0
HTTP:APACHE:REDHAT-DIRLIST	By submitting a malformed HTTP GET request to an Apache server using the default configuration supplied with several versions of RedHat Linux an attacker can cause the web server to return a listing of the contents of that directory, even if an index page is present.	low	sos5.1.0
HTTP:APACHE:RESIN-WEB-INF	This signature detects attempts to exploit a flaw in Resin 2.1.12, a Java Scriptlet server. Attackers can send malformed URL requests to a server to allow access to a normally protected subdirectory, the WEB-INF directory.	medium	sos5.0.0, sos5.1.0
HTTP:APACHE:SCALP	This signature detects attempts to exploit a vulnerability in Apache Web servers. Apache improperly calculates required buffer sizes for chunked encoded requests due to a signed interpretation of an unsigned integer value. Attackers may send chunked encoded requests with the unique Host header value "apache-scalp.c." in the GET request to create a buffer overflow and execute arbitrary code.	critical	sos5.0.0, sos5.1.0
HTTP:AUDIT:MSNG-HTTP-VER	This protocol anomaly is an HTTP request with no version number after the 'HTTP/...'. This may indicate command line access to an HTTP server.	info	sos5.1.0
HTTP:AUDIT:UNKNWN-REQ	This protocol anomaly is an unknown HTTP request. Known requests are OPTION, GET, HEAD, POST, PUT, DELETE, TRACE, and CONNECT.	info	sos5.1.0

HTTP:BADBLUE:INVALID-GET-DOS	This signature detects denial-of-service attempts against Working Resources BadBlue Web server. Attackers may send a maliciously crafted HTTP GET request to the Web server to disable the daemon and render it unusable until restarted.	medium	sos5.1.0
HTTP:BADBLUE:PROXY-RELAY	This signature detects attempts to relay a web request through a BadBlue web server. When BadBlue is using its default configuration, attackers may use the web server as a proxy server to attack internal targets or mask attack activity.	medium	sos5.1.0, sos5.0.0
HTTP:BIGBROTHER:DIR-TRAVERSAL	This signature detects attempts to view files on the Web server using the BigBrother bb-hist.sh history browser script. Attackers may view any files on the Web server that are accessible to the user the history browser script is running under.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:ALTAVISTA-TRAVERSAL	This signature detects attempts to exploit a vulnerability in the AltaVista Search engine. The search engine sets up a Web server at port 9000 that listens for search queries. The search function accepts a single '../' string in the query, providing access to the parent, or 'http' directory. This directory typically contains administrative documents that may include the password for the remote administration utility, which is base-64 encoded. Attackers may send multiple '../' strings in hex code (ie. '%2e%2e%2f') in a query to access the remote administration utility password and gain full remote administration abilities.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:ANYFORM-SEMICOLON	This signature detects attempts to exploit the AnyForm CGI script, a popular CGI form designed to support simple forms that deliver responses via e-mail. Some versions of AnyForm did not perform user supplied data sanity checking, and may allow remote execution of arbitrary commands on the server.	high	sos5.1.0, sos5.1.0
HTTP:CGI:APPLE-QT-FILEDISC1	This signature detects attempts to exploit a vulnerability in Apple QuickTime Streaming Server. QuickTime Streaming Server v4.1.1 and earlier versions are vulnerable. Attackers may send a maliciously crafted URL to parse_xml.cgi to view files that are not usually accessible through HTTP.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:AXIS-ACCOUNT	This signature detects a request to an Axis Video Server containing parameters designed to create an Administrator account on the server.	critical	sos5.1.0
HTTP:CGI:BNB-SURVEY-REMOTE-EXEC	This signature detects attempts to access the BNBSurvey survey.cgi program. Attackers may remotely execute commands via shell metacharacters.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:BUGZILLA-SEMICOLON	This signature detects shell access attempts to exploit the process_bug.cgi script vulnerability in Bugzilla. Attackers may send a semicolon as an argument to the script, followed by arbitrary shell commands.	high	sos5.0.0, sos5.1.0

HTTP:CGI:DCFORUM-AZ-EXEC	This signature detects shell attempts to exploit the dcforum.cgi script in DCScripts DC Forum (all versions), which is used to manage web-based discussion boards. Attackers may use maliciously crafted URL requests with the pipe and newline characters to execute arbitrary scripts on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:FORMMAIL-ENV-VAR	This signature detects access to the FormMail CGI program. Attackers may use this program to remotely execute commands.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:HASSAN-DIR-TRAVERSAL	This signature detects attempts to exploit a vulnerability in the Hassan shopping cart script shop.cgi. Attackers may access arbitrary system files.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:HTDIG-INCLUSION	This signature detects attempts to exploit a vulnerability in ht://dig, a Web content search engine for UNIX. Because ht://dig improperly validates form input, attackers may pass a maliciously crafted variable to the htsearch CGI script to read files accessible to the program user.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:HYPERSEEK-DIR-TRAVERSL	This signature detects attempts to exploit a vulnerability in hsx.cgi, which ships as part of iWeb Hyperseek 2000. Attackers may view arbitrary files and directories.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:IKONBOARD-BADCOOKIE	This signature detects attempts to exploit a vulnerability in IkonBoard, a popular Web-based discussion board. Attackers may send a maliciously crafted cookie that contains illegal characters to IkonBoard to execute arbitrary code with IkonBoard privileges (typically user level).	medium	sos5.1.0
HTTP:CGI:INFO2WWW-EXEC	This signature detects attempts to exploit a vulnerability in the info2www CGI script. Attackers may execute arbitrary binaries on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:INFOSRCH-REMOTE-EXEC	This signature detects attempts to exploit a vulnerability in the infosrch.cgi script. Attackers may execute commands on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:LIBCGI-RFP-OVERWRITE	This signature detects attempts to exploit a vulnerability in LIB CGI. Attackers may inject maliciously crafted C code into LIB CGI applications to overwrite the Frame Pointer and execute arbitrary code on the host.	medium	sos5.1.0
HTTP:CGI:MOREOVER-CACHE-FEED	This signature detects attempts to exploit a vulnerability in the cached_feed.cgi script provided by moreover.com. Attackers may view arbitrary system files that are readable by the HTTPd process.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:TECHNOTE-MAIN-DCLSR	This signature detects directory traversal attempts that exploit the main.cgi script in TECH-NOTE 2000. Because the script validates input incorrectly, attackers may remotely access arbitrary files from the server.	medium	sos5.0.0, sos5.1.0

HTTP:CGI:TECHNOTE-PRINT-DSCLSR	This signature detects directory traversal attempts that exploit the print.cgi script in TECH-NOTE 2000. Because the script validates input incorrectly, attackers may remotely access arbitrary files from the server.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:W3-MSQL-CGI-OF	This signature detects attempts to exploit a vulnerability in W3-msql, a CGI program that acts as a Web interface for Mini SQL (mSQL). W3-msql version 2.0.11 is vulnerable. Attackers may remotely send a maliciously crafted scanf call to overflow the content-length field and execute arbitrary code with Web server privileges.	medium	sos5.0.0, sos5.1.0
HTTP:CGI:W3-MSQL-FILE-DISCLSR	This signature detects buffer overflow attempts that exploit the w3-msql CGI script in mini-SQL. Attackers may execute arbitrary commands on the server.	medium	sos5.1.0
HTTP:CGI:WEBPALS-EXEC	This signature detects attempts to exploit a vulnerability in the WebPALS CGI script. Attackers may remotely execute arbitrary code with root permissions.	critical	sos5.0.0, sos5.1.0
HTTP:CGI:WEBSPEED-WSMADMIN	This signature detects attempts to gain administrative access to the WebSpeed server without normal authentication.	critical	sos5.0.0, sos5.1.0
HTTP:CGI:WEBSPIRS-FILE-DISCLSR	This signature detects attempts to exploit a vulnerability in the SilverPlatter WebSPIRS webspiers.cgi file. Attackers may access arbitrary system files	medium	sos5.0.0, sos5.1.0
HTTP:CGI:YABB-DIR-TRAVERSAL	This signature detects attempts to exploit a vulnerability in the YaBB.pl CGI script. Attackers may view arbitrary files.	medium	sos5.0.0, sos5.1.0
HTTP:CHKP:AUTH-FMT-STR	This signature detects attempts to exploit a vulnerability in some Web servers and Web proxies. Attackers may send user authentication that includes format strings to crash some Web servers, creating a denial-of-service (DoS) or enabling the attackers to take control of the firewall as root.	critical	sos5.1.0
HTTP:CHKP:FW1-FORMAT-STR	This signature detects attempts to exploit a vulnerability in the CheckPoint AI/Smart Defense HTTP proxy engine. Attackers may send a scheme that includes format strings to crash the proxy engine, creating a denial-of-service (DoS) or enabling the attackers to take control of the firewall as root.	critical	sos5.1.0, sos5.0.0
HTTP:CHKP:FW1-PROXY	This signature detects attempts to exploit the web proxy functions of CheckPoint FireWall-1. When the HTTP CONNECT method, used to build generic Transit Layer Security over HTTP, is used by default, the firewall web proxies may be used as open TCP proxies. Attackers may use an HTTP proxy to connect to a server, then use the CONNECT method to access other servers and launch further attacks.	medium	sos5.0.0, sos5.1.0
HTTP:CISCO:IOS-ADMIN-ACCESS	This signature detects attempts to exploit a vulnerability in Cisco IOS. Attackers may remotely gain full administrative access to the router.	critical	sos5.0.0, sos5.1.0

HTTP:CISCO:VOIP:PORT-INFO-DOS	This signature detects attempts to exploit a vulnerability in Cisco VoIP phones. Versions CP-7910 and later are vulnerable. Attackers may send an arbitrarily long (120000+) StreamID to the PortInformation script to cause an error message that displays a memory dump. Attackers may use this information to reconstruct the calling patterns of a particular phone.	medium	sos5.0.0, sos5.1.0
HTTP:CISCO:VOIP:STREAM-ID-DOS	This signature detects denial-of-service (DoS) attempts against Cisco VoIP phones. Versions CP-7910 and later are vulnerable. Attackers may send an arbitrarily long (120000+) StreamID to the StreamingStatistics script to cause the phone to reset, creating a DoS for 30 seconds (or until the phone reboots).	medium	sos5.0.0, sos5.1.0
HTTP:COLDFUSION:EXPRCALC-OPNFIL	This signature detects attempts to exploit a vulnerability in the ColdFusion ExprCalc.cfm script. Attackers may delete files from a Web server.	medium	sos5.0.0, sos5.1.0
HTTP:COLDFUSION:HEADER-LOG-OF	This signature detects attempts to exploit a vulnerability in the JRun component of Macromedia ColdFusion web server. Attackers may send overly long HTTP headers to overflow the logging function, enabling an attacker to crash or take control of the web server.	high	sos5.1.0
HTTP:COLDFUSION:JRUN-SC-PARSE	This signature detects attempts to exploit a vulnerability in the JRun component of Macromedia ColdFusion web server. Attackers may pass a semicolon character to JRun to expose the script source code and other sensitive files.	low	sos5.0.0, sos5.1.0
HTTP:DIR:CRYSTAL-REPORTS	This signature detects attempts to exploit a vulnerability in Microsoft Crystal Reports. Users of Visual Studio .NET 2003, Outlook 2003 with Business Contact Manager, or Microsoft Business Solutions Customer Relationship Management (CRM) 1.2 are affected. Attackers may send a malformed URL to the server to read or write to any file on the server.	high	sos5.1.0
HTTP:DIR:DEEP-PARAM-TRAVERSE	This signature detects directory traversal attempts within HTTP GET or POST form parameters that extend three or more directories. Attackers may exploit a poorly-written CGI program to access or modify private files.	medium	sos5.1.0
HTTP:DIR:PARAM-TRAVERSE	This signature detects directory traversal attempts within HTTP GET or POST form parameters. Attackers may exploit a poorly-written CGI program to access or modify private files.	low	sos5.1.0
HTTP:DIR:TRAVERSE-DIRECTORY	This protocol anomaly is an HTTP directory traversal attempt, i.e. ../../ or ../. This may indicate an attempt to evade an IDS (DI is not vulnerable). Note that some Websites refer to directories in a way that looks like a traversal.	medium	sos5.0.0, sos5.1.0
HTTP:EXPLOIT:AMBIG-CONTENT-LEN	This protocol anomaly is an HTTP request that has a Content-Length and Transfer-Encoding header. RFC-2616#4.4 specifies that only one of these two headers should be used in an HTTP request.	low	sos5.0.0, sos5.1.0

HTTP:EXPLOIT:BLAZIX-JSPVIEW	This signature detects attempts to exploit a vulnerability in the Blazix, a Java-based Web server. Blazix 1.2 and earlier versions are vulnerable. Because Blazix does not strip bad characters (such as '+' and '"') from URL requests, attackers may send a malicious URL to the Web server to view the jsp server side scripts.	medium	sos5.0.0, sos5.1.0
HTTP:EXPLOIT:BRUTE-FORCE	This protocol anomaly is too many authentication failures (Web pages that require authentication) within a short period of time between a unique pair of hosts.	high	sos5.1.0
HTTP:EXPLOIT:BRUTE-SEARCH	"This protocol anomaly is multiple 301 (Moved Permanently), 403 (Forbidden), 404 (Not Found) and 405 (Method Not Allowed) errors between a unique pair of hosts within a short period of time. This could indicate that a search robot or a script is methodically searching a Web site for vulnerable directories or CGI scripts. The default maximum number of 301/403/404/405 errors is 16.		
HTTP:EXPLOIT:IE-ZONE-SPOOF	This signature detects attempts to access potentially malicious Web sites. When using Microsoft Internet Explorer, a user can be tricked into visiting a malicious Web site that they believe is benign. Additional IE vulnerabilities may allow the malicious Web site to run scripts in the Local Computer zone, which bypasses security checks on the user's machine. In your logs for the event, the malicious Web site appears as the destination IP address.	high	sos5.0.0, sos5.1.0
HTTP:EXPLOIT:ILLEGAL-HOST-CHAR	This signature detects illegal characters in a Host header field of an HTTP/1.1 request. Attackers may send an HTTP link, that, when selected by the user, generates an HTTP request to a malicious Web site. In your logs, the destination IP address for the event may be the malicious Web site; however, some foreign Web sites may also trigger this signature, creating a false positive. Per RFC, '_' is not a legal character for a host name.	medium	sos5.0.0, sos5.1.0
HTTP:EXPLOIT:REALPLAYER-SKIN	This signature detects malicious RealPlayer skin files.	high	sos5.1.0
HTTP:EXPLOIT:SHOUTCAST-FMT-STR	This signature detects attempts to exploit a known vulnerability in the Shoutcast streaming audio server. Attackers may gain complete control of the target host.	medium	sos5.1.0
HTTP:EXPLOIT:WIN-MAL-COMP-FILE	This signature detects attempts to exploit a vulnerability in Microsoft Windows native compressed file handler. Attackers may send .zip files with overly long filenames to overflow the file handler and run arbitrary code.	high	sos5.1.0
HTTP:EXT:GRP-EXT-HTTP	This signature detects GRP files sent over HTTP. GRP files can contain Windows Program Group information, and may be exploited by malicious users to deposit instructions or arbitrary code on a target's system. User involvement is required to activate GRP files; typically they are attached or linked to a harmless-appearing e-mail message.	medium	sos5.1.0

HTTP:EXT:JOB	This signature detects an attempt to download a Microsoft Task Scheduler (.job) file. Opening a malicious .job file in Task Scheduler may allow for arbitrary code execution, leading to system compromise. This vulnerability is present in Microsoft Windows 2000 Service Pack 2 and later. It is also present in Microsoft Windows XP Service Pack 1.	medium	sos5.0.0, sos5.1.0
HTTP:FRONTPAGE:ADMIN.PWD-REQ	This signature detects attempts to access the Microsoft FrontPage Extensions for UNIX .pwd file that contains sensitive account information.	medium	sos5.0.0, sos5.1.0
HTTP:FRONTPAGE:DOS-NAME-DOS	This signature detects attempts to exploit a known vulnerability in Microsoft Frontpage. Attackers may send a malformed request with an MS-DOS device name to shtml.exe to crash the server.	medium	sos5.1.0
HTTP:FRONTPAGE:FOURDOTS	This signature detects attempts to exploit the '/.../' directory traversal vulnerability in Microsoft FrontPage PWS.	medium	sos5.0.0, sos5.1.0
HTTP:FRONTPAGE:FP30REG.DLL-OF	This signature detects buffer overflow attempts against Microsoft FrontPage extensions in Windows 2000 and XP. Attackers may execute arbitrary code on the target host.	critical	sos5.1.0
HTTP:FRONTPAGE:SERVICE.PWD-REQ	This signature detects attempts to access the Microsoft FrontPage extensions for UNIX .pwd file which contains sensitive account information.	medium	sos5.0.0, sos5.1.0
HTTP:HOSTCTRL:BROWSE-ASP	This signature detects attempts to exploit a vulnerability in the browse.asp script supplied with Hosting Controller, a tool that allows Microsoft Windows network administrators to centralize administrative tasks into one interface. Attackers may send a maliciously crafted URL request for browse.asp to view arbitrary directories and files on hard drives.	medium	sos5.0.0, sos5.1.0
HTTP:HOTMAIL:EXE-DOWNLOAD	This signature detects attempts by users to download potentially hazardous attachments from MSN Hotmail.	medium	sos5.1.0
HTTP:IIS:AD-SERVER-CONFIG	This signature detects attempts to download the site.csc configuration file for Microsoft Ad Server. Attackers may access sensitive information.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:ASP-CODEBROWSER-EXAIR	This signature detects attempts to exploit the Showcode ASP vulnerability in Microsoft IIS.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:ASP-DOT-NET-BACKSLASH	This signature detects backslash (\) characters in the URL portion of an HTTP request. Attackers may use a backslash as a directory separator instead of the normal forward slash (/) to bypass the Microsoft IIS ASP.Net authentication capabilities and access protected resources. Note: A poorly configured web server may also display a backslash in a non-malicious URL request.	medium	sos5.1.0
HTTP:IIS:BAT-&	This signature detects attempts to execute a command by specifying a .bat or .cmd extension to a Microsoft Windows Web server.	high	sos5.0.0, sos5.1.0

HTTP:IIS:COMMAND-EXEC	This signature detects attempts to exploit Microsoft Windows Web servers. Attackers may send a maliciously crafted url containing the string "cmd.exe" to execute commands on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:COMMAND-EXEC-2	This signature detects attempts to exploit a vulnerability in Microsoft IIS. Attackers may execute arbitrary commands on the Web server.	high	sos5.0.0, sos5.1.0
HTTP:IIS:DATA-DISCLOSURE	This signature detects attempts to obtain the sourcecode of Active Server Pages served by Microsoft's Internet Information Server. In IIS, remote attackers can obtain source code for ASP files by appending ":::\$DATA" to the URL.	high	sos5.0.0, sos5.1.0
HTTP:IIS:HEADER-HOST-DOS	This signature detects denial-of-service (DoS) attempts against Microsoft IIS. Attackers may pass maliciously malformed header values to the host to crash the IIS service.	high	sos5.0.0, sos5.1.0
HTTP:IIS:IIS-NSIISLOG-OF	This signature detects buffer overflow attempts against Microsoft Windows Media Services, included with Microsoft Windows 2000 Server SP4. Attackers may send a maliciously crafted HTTP 'POST' request to overflow the buffer.	critical	sos5.1.0, sos5.1.0
HTTP:IIS:ISAPI-IDA-OVERFLOW	This signature detects buffer overflow attempts against Microsoft ISAPI Indexing Service for IIS. Index Server 2.0 and Indexing Service 2000 in IIS 6.0 beta and earlier versions are vulnerable. Attackers may send a long argument to Internet Data Administration (.ida) and Internet Data Query (.idq) files to overflow the buffer in the ISAPI extension (idq.dll) and execute arbitrary commands.	critical	sos5.1.0
HTTP:IIS:ISAPI-IDQ-OVERFLOW	This signature detects buffer overflow attempts against Microsoft ISAPI Indexing Service for IIS. Index Server 2.0 and Indexing Service 2000 in IIS 6.0 beta and earlier versions are vulnerable. Attackers may send a long argument to Internet Data Administration (.ida) and Internet Data Query (.idq) files to overflow the buffer in the ISAPI extension (idq.dll) and execute arbitrary commands.	critical	sos5.1.0, sos5.0.0
HTTP:IIS:ISAPI-PRINTER-OVERFLOW	This signature detects attempts to execute a buffer overflow in the Microsoft IIS 5.0 .printer ISAPI extension.	critical	sos5.0.0, sos5.1.0
HTTP:IIS:MALFORMED-HTR-REQUEST	This signature detects malformed .htr requests that may cause a denial-of-service (DoS).	high	sos5.0.0, sos5.1.0
HTTP:IIS:MDAC-RDS	This signature detects attempts to exploit the Microsoft Data Access Components (MDAC) Remote Data Services (RDS) component. Attackers may access files and other services.	high	sos5.0.0, sos5.1.0

HTTP:IIS:MDAC-RDS-2	This signature detects attempts to exploit the Remote Data Services (RDS) component included in Microsoft Data Access Components (MDAC) using ActiveDataFactory. Microsoft IIS 3.x and 4.x are vulnerable. Attackers may remotely access exposed unsafe methods to execute arbitrary commands.	high	sos5.0.0, sos5.1.0
HTTP:IIS:MFC-EXT-OF	This signature detects buffer overflow attempts against Microsoft IIS. A maliciously crafted HTTP request can exploit a buffer overflow condition in mfc42.dll by way of ext.dll. Attackers may gain local access to an IIS server.	high	sos5.1.0
HTTP:IIS:NEWDSN-FILE-CREATION	This signature detects attempts to create a file on the Web server by exploiting the newdsn.exe vulnerability in Microsoft IIS 3.0.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:NSIISLOG-CHUNKED-POST	This signature detects chunked POST requests to NSIISLOG.DLL. Attackers may exploit Windows Media Services that have logging enabled, and other vulnerabilities using this method.	high	sos5.1.0
HTTP:IIS:OUTLOOK-WEB-DOS	This signature detects denial-of-service (DoS) attempts against Microsoft Outlook Web. Attacker may send a long string of '%' characters as the user name and/or password.	high	sos5.0.0, sos5.1.0
HTTP:IIS:PROPFIND	This signature detects attempts to exploit a vulnerability in Microsoft IIS 5.0. Attackers may send malicious 'PROPFIND' requests to the server to crash it.	medium	sos5.1.0
HTTP:IIS:SADMIND-WORM-ACCESS	This signature detects the sadmind/IIS worm attempting to infect Microsoft IIS. The sadmind/IIS worm first exploits a vulnerability in a Solaris system, then attacks Microsoft IIS Web servers using the Web server folder directory traversal exploit.	medium	sos5.1.0
HTTP:IIS:SENSEPOST.EXE	This signature detects attempts to locate sensepost.exe on a Microsoft ISS Web Server. Attackers may use a proof-of-concept hacking tool to break into a vulnerable Web server, then copy cmd.exe to the Web server script directory and rename it sensepost.exe to avoid detection by log viewers. To identify this event, check your Web server logs for details--if the server returned a '200' to the request, your Web server may be compromised.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:SITE-SERVER-FILE-UPLD	This signature detects attempts to exploit a vulnerability in MS Site Server 2.0 with IIS 4. Attackers may upload content (including ASP) to the target web site and remotely execute commands.	high	sos5.1.0
HTTP:IIS:WEBDAV:LOCK-OF	This signature detects buffer overflow attempts against Microsoft IIS WebDAV. Attackers may send a maliciously crafted WebDAV URL request that contains 65535 or 65536 bytes to the Web server to execute arbitrary code as the system account.	critical	sos5.1.0

HTTP:IIS:WEBDAV:MALFORMED-REQ1	This signature detects denial-of-service (DoS) attempts against Microsoft IIS 5.0 servers with WebDAV extensions enabled. Attackers may send a maliciously crafted WebDAV SEARCH request in an HTTP request to DoS the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:WEBDAV:MALFORMED-REQ2	This signature detects denial-of-service (DoS) attempts against Microsoft IIS 5.0 servers with WebDAV extensions enabled. Attackers may send a maliciously crafted WebDAV SEARCH request in an HTTP request to DoS the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:IIS:WEBDAV:SEARCH-OF	This signature detects buffer overflow attempts against Microsoft IIS WebDAV. Attackers may send a maliciously crafted WebDAV URL request that contains 65535 or 65536 bytes to the Web server to execute arbitrary code as the system account.	critical	sos5.1.0
HTTP:IIS:WEBDAV:XML-HANDLER-DOS	This signature detects denial-of-service (DoS) attempts against the WebDAV XML Message Handler in Microsoft IIS. Attackers may send a malicious HTTP request to a WebDAV enabled IIS server to cause it to consume all system resources. A machine reboot is required to resume service.	medium	sos5.1.0
HTTP:INFO:HTTPPOST-GETSTYLE	This signature detects HTTP POST requests with GET parameters. POST requests should not have parameters on the same line as the request method. This may indicate a poorly-written Web application or HTTP tunneling.	info	sos5.1.0
HTTP:INFO-LEAK:GOAHEAD-PERM	"This signature detects attempts to bypass directory permissions set on the /cgi-bin directory of a GoAhead web server. GoAhead WebServer versions 2.1.8 and earlier are vulnerable. Attackers may supply an invalid URL to the server to reveal the contents of certain private directories on the server.		
HTTP:INFO-LEAK:HTACCESS	This signature detects probes for the .htaccess file, used by the Apache Web server for configuration directives. Attackers may be attempting to gain access to the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:INFO-LEAK:HTPASSWD-REQUEST	This signature detects attempts to access the .htpasswd file on a Web server.	medium	sos5.0.0, sos5.1.0
HTTP:INFO-LEAK:VIGNETTE-DIAG	This signature detects attempts to access the diagnostic utility supplied with the Vignette Application server. Because the utility does not use access controls, attackers (or any client) may connect to the utility and access sensitive configuration information.	low	sos5.0.0, sos5.1.0
HTTP:INFO-LEAK:VIGNETTE-LEAK	"This signature detects attempts to exploit a vulnerability in Vignette Story Server. Vignette Story Server versions 4.1 and 6 are vulnerable. Attackers may expose information about user sessions, server side code, and other sensitive information.		

HTTP:INFO-LEAK:WEB-INF-DOT	This signature detects attempts to exploit a vulnerability in Windows Web servers with J2EE. Attackers may append a '.' character to a request for the WEB-INF directory (where J2EE class files are typically stored) to bypass directory security and gain access to normally protected files.	medium	sos5.0.0, sos5.1.0
HTTP:INFO-LEAK:WR850-CONF-DL	This signature detects attempts to download the configuration file from a Motorola WR850G Wireless Router.	medium	sos5.0.0, sos5.1.0
HTTP:INVALID:INVLD-AUTH-CHAR	This protocol anomaly is an HTTP header with an authorization string that contains an invalid character. The authorization line is decoded using base64.	high	sos5.1.0
HTTP:INVALID:INVLD-AUTH-LEN	This protocol anomaly is an HTTP header with an authorization string that has an invalid length (a length that is not a multiple of 4). Because the authorization line is encoded/decoded using base64, the length must be a multiple of 4.	high	sos5.1.0
HTTP:INVALID:MISSING-REQ	This protocol anomaly is an HTTP header that has no request line or request uniform resource identifier (URI).	medium	sos5.1.0
HTTP:IRIX:CGI-BIN-WRAP	This signature detects attempts to exploit the wrap CGI script in SGI IRIX. Attackers may list the contents of Web server directories.	medium	sos5.0.0, sos5.1.0
HTTP:MISC:EMULIVE-ADMIN	This signature detects an attempt to gain unauthorized administrative access to an EmuLive Server4 daemon.	medium	sos5.1.0
HTTP:MISC:HP-PROCURVE-RESET	This signature detects denial-of-service (DoS) attempts against the HP Procurve 4000M switch. Configuration changes for the switch are made via an HTTP-based interface; however, the script that resets the switch after a configuration change does not properly authenticate the IP address that calls the script. Attackers may call the script repeatedly to perform a DoS.	medium	sos5.0.0, sos5.1.0
HTTP:MISC:MOBY-LENGTH-DOS	This signature detects denial-of-service (DoS) attacks against the Moby NetSuite. Attackers may send a maliciously crafted HTTP POST request that contains an invalid Content-Length field to the host to crash the Web server.	medium	sos5.1.0
HTTP:MISC:MOODLOGIC-CLIENT	This signature detects use of the Mood Logic client. Mood Logic is an MP3 catalogue system that helps users identify and classify MP3s. If your organization prohibits the use of MP3s, use this signature to detect Mood Logic clients.	info	sos5.1.0
HTTP:MISC:NG-WG602-BACKDOOR	This signature detects attempts to administer a Netgear WG602 using an undocumented administrator username/password that cannot be changed or disabled. Attackers can modify any setting on the WG602 to perform a denial-of-service (DoS) on the Netgear device or circumvent other access control protocols.	high	sos5.1.0

HTTP:MISC:NOOP-SLIDE-HEAD-OF	This signature detects buffer overflow attempts against Web servers on Intel x86 platforms. Attackers may use the "No-Op Slide" attack to pad the stack with "No Operation" x86 CPU instructions and overwrite the return address.	critical	sos5.0.0, sos5.1.0
HTTP:MISC:NOOP-SLIDE-REQ-OF	This signature detects buffer overflow attempts against Web servers on Intel x86 platforms. Attackers may use the "No-Op Slide" attack to pad the stack with "No Operation" x86 CPU instructions and overwrite the return address.	critical	sos5.1.0
HTTP:MISC:SHAMBALA-DOS1	This signature detects denial-of-service (DoS) attempts against Evolvable Shambala Server, an FTP, Web, and Chat server. Version 4.5 is vulnerable. Attackers may send a maliciously crafted request to the Web server to cause a DoS.	medium	sos5.1.0
HTTP:MISC:VISNETIC-DOS	This signature detects attempts to exploit a vulnerability in VisNetic WebSite. Versions 3.5.13.1 and earlier are vulnerable. Attackers may send a malicious OPTIONS request to crash the server.	medium	sos5.1.0
HTTP:MISC:WR850-WEBSHELL	This signature detects attempts to access a debug mode web shell supplied with the Motorola WR850 Wireless Router. Attackers may use this access exploit in conjunction with an authentication bypass exploit to gain full control over the router.	high	sos5.0.0, sos5.1.0
HTTP:NETSCAPE:ENTERPRISE-DOS	This signature detects denial-of-service (DoS) attempts that exploit the Web Publishing REVLOG command in Netscape Enterprise Server 3.x.	high	sos5.1.0
HTTP:NOVELL:NETWARE-CONVERT.BAS	This signature detects directory traversal attempts on Novell NetWare Web Server 2.x. The convert.bas CGI script allows file retrieval outside of normal Web server context. Attackers may submit the filename and path as a parameter to the script using relative paths (../..) to traverse directories.	medium	sos5.0.0, sos5.1.0
HTTP:OREILLY:WIN-C-SMPLE-OVFLOW	This signature detects buffer overflow attempts that exploit the win-c-sample.exe sample script vulnerability in O'Reilly Website Pro 2.0 Web server. The script is placed in the /cgi-shl directory off of the Web root by default.	medium	sos5.0.0, sos5.1.0
HTTP:OVERFLOW:ACCEPT	DI has detected a suspiciously long Accept header.	medium	sos5.1.0
HTTP:OVERFLOW:ACCEPT-ENCODING	DI has detected a suspiciously long Accept-Encoding header.	medium	sos5.1.0
HTTP:OVERFLOW:ACCEPT-LANGUAGE	DI has detected a suspiciously long Accept-Language header.	medium	sos5.1.0
HTTP:OVERFLOW:ATP-HTTPD-OF	This signature detects buffer overflow attempts against ATPhttp versions 0.4b and earlier. Attackers may send an overly long GET request to the Web server daemon to overflow the buffer.	critical	sos5.1.0
HTTP:OVERFLOW:AUTHORIZATION	This protocol anomaly is an HTTP authorization header that exceeds the user-defined maximum. The default length is 128.	medium	sos5.1.0

HTTP:OVERFLOW:AUTH-OVFLW	This protocol anomaly is an HTTP header with an authorization line that exceeds the user-defined maximum. The default authorization line length is 128.	high	sos5.1.0
HTTP:OVERFLOW:CHUNK-LEN-OFLOW	This protocol anomaly is an HTTP message that has a chunk length in a Transfer-Encoding: chunk request that is greater than 0x7fffffff. Apache servers 1.3 to 1.3.24 and 2.0 to 2.0.36 are vulnerable. Attackers may cause a denial-of-service (DoS) or execute arbitrary code on the server.	critical	sos5.0.0, sos5.1.0
HTTP:OVERFLOW:CHUNK-OVERFLOW	This protocol anomaly is an invalid data chunk length in an HTTP request that uses chunked encoding. The chunked encoding transfer method sends data length requests followed by data chunks that match the negotiated data lengths. Attackers may cause a stack overflow and execute arbitrary code on the server.	critical	sos5.1.0
HTTP:OVERFLOW:CONNECTION	DI has detected a suspiciously long Connection header.	medium	sos5.1.0
HTTP:OVERFLOW:CONTENT-ENCODING	DI has detected a suspiciously long Content-Encoding header.	medium	sos5.1.0
HTTP:OVERFLOW:CONTENT-LANGUAGE	DI has detected a suspiciously long Content-Language header.	medium	sos5.1.0
HTTP:OVERFLOW:CONTENT-LENGTH	DI has detected a suspiciously long Content-Length header.	medium	sos5.1.0
HTTP:OVERFLOW:CONTENT-LOCATION	DI has detected a suspiciously long Content-Location header.	medium	sos5.1.0
HTTP:OVERFLOW:CONTENT-MD5	DI has detected a suspiciously long Content-MD5 header.	medium	sos5.1.0
HTTP:OVERFLOW:CONTENT-OVERFLOW	This protocol anomaly is a missing line break after a specified data length in an HTTP request using content length transfer. The content length transfer method sends the specified data length in the BODY of the request followed by a line break.	critical	sos5.1.0
HTTP:OVERFLOW:CONTENT-TYPE	This protocol anomaly is a Content-Type header length that exceeds the user-defined maximum. The default length is 64.	medium	sos5.1.0
HTTP:OVERFLOW:COOKIE	This protocol anomaly is an HTTP Cookie header length that exceeds the user-defined maximum. The default length is 8192.	medium	sos5.1.0
HTTP:OVERFLOW:HEADER	This protocol anomaly is an HTTP header field that is too long, and may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
HTTP:OVERFLOW:HOST	This protocol anomaly is an HTTP Host header length that exceeds the user-defined maximum. The default length is 64.	medium	sos5.1.0
HTTP:OVERFLOW:HTTPA-OFF	This signature detects buffer overflow attacks against the HTTPa daemon. Attackers may send a maliciously crafted HTTP GET request to the host to overflow the buffer.	high	sos5.1.0

HTTP:OVERFLOW:INV-CHUNK-LEN	This protocol anomaly is an invalid chunk length specification in a chunked transfer encoded HTTP request. RFC-2616#3.6.1 specifies that the size of a chunk should be represented using hexadecimal notation.	high	sos5.0.0, sos5.1.0
HTTP:OVERFLOW:JANASRV-VER-OF	This signature detects buffer overflow attempts against JanaServer HTTP Server, an Internet gateway for Windows. JanaServer 2.21 and prior are vulnerable. Attackers may send a maliciously crafted HTTP GET request to overflow the buffer.	medium	sos5.1.0
HTTP:OVERFLOW:LIBHTTPD-GET-OF	This signature detects buffer overflow attempts against LibHTTPd. LibHTTPd 1.2 and earlier are vulnerable. Attackers may send a maliciously crafted GET request to execute arbitrary code on the host.	high	sos5.1.0
HTTP:OVERFLOW:METHOD-GENRC-OF	This signature detects buffer overflow attempts against HTTP request methods. Attackers may send an invalid or long HTTP request to overflow vulnerable buffers on the target Web server.	high	sos5.1.0
HTTP:OVERFLOW:NULLHTTPD-ROOT-OF	This signature detects buffer overflow attempts against Null HTTPD. Attackers may remotely send shellcode in a maliciously crafted POST command to gain local access.	critical	sos5.1.0
HTTP:OVERFLOW:PI3WEB-SLASH-OF	This signature detects denial-of-service (DoS) attempts against Pi3Web Server. Attackers may send a URL with more than 354 Slashes (/) to crash the server.	medium	sos5.0.0, sos5.1.0
HTTP:OVERFLOW:REFERER	This protocol anomaly is an HTTP Referrer header length that exceeds the user-defined maximum. The default length is 8192.	medium	sos5.1.0
HTTP:OVERFLOW:SAMBAR-SEARCH	This signature detects buffer overflow attempts against Sambar Server, a free Web server. Attackers may include an oversized HTTP header within a maliciously crafted request to the server to execute arbitrary code.	critical	sos5.1.0
HTTP:OVERFLOW:SERVER	DI has detected a suspiciously long Server header.	medium	sos5.1.0
HTTP:OVERFLOW:SET-COOKIE	DI has detected a suspiciously long Set-Cookie header.	medium	sos5.1.0
HTTP:OVERFLOW:TRANSFER-ENCODING	DI has detected a suspiciously long Transfer-Encoding header.	medium	sos5.1.0
HTTP:OVERFLOW:USER-AGENT	This protocol anomaly is an HTTP User-Agent header length that exceeds the user-defined maximum. The default length is 258.	medium	sos5.1.0
HTTP:PHP:ALEXPHP-INCLUDE	This signature detects attempts to exploit a remote file inclusion vulnerability in AlexPHP. Attackers may send a maliciously crafted HTTP request to execute PHP code from a remote server on the host running AlexPHP.	high	sos5.1.0

HTTP:PHP:BLACKBOARD-INC	This signature detects attempts to exploit a vulnerability in the admin.inc.php script that shipped as part of the BlackBoard suite. Attackers may force the admin.inc.php script to include and execute PHP code from a remote source.	high	sos5.1.0
HTTP:PHP:COOLPHP-DIRTRAV	This signature detects directory traversal attempts against CoolPHP. Attackers may use this exploit to execute arbitrary scripts on the PHP server.	medium	sos5.1.0
HTTP:PHP:DFORUM-PHP-INC	This signature detects attempts to exploit a vulnerability in D-Forum. D-Forum versions 1.0 through 1.11 are vulnerable. Attackers may exploit header.php3 and footer.php3 to include PHP code from a remote host and execute arbitrary commands.	high	sos5.0.0, sos5.1.0
HTTP:PHP:FI-DIR-TRAVERSAL	This signature detects attempts to exploit a design vulnerability in PHP/FI. Attackers may remotely access files and directories that are readable by the Web server UID to gather information on the local host and retrieve encrypted user passwords on the system.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:GALLERY:EMBED-AUTH	This signature detects attempts to exploit a vulnerability in Gallery, a Web-based photo album application written in php. Attackers may bypass user authorization to gain administrative privileges.	high	sos5.0.0, sos5.1.0
HTTP:PHP:GALLERY:HTTP-VARS	This signature detects attempts to exploit a vulnerability in Gallery, a Web-based photo management application. Gallery uses the variables HTTP_POST_VARS, HTTP_GET_VARS, HTTP_COOKIE_VARS, and HTTP_POST_FILES to transfer data between pages, including the GALLERY_BASEDIR variable. Attackers may manually control these variables to include a malicious setting for GALLERY_BASEDIR, enabling them to execute arbitrary PHP code on the Gallery server with the permissions of the HTTP server.	high	sos5.1.0
HTTP:PHP:GALLERY:MAL-INCLUDE	This signature detects attempts to exploit a vulnerability in Gallery online photo gallery software. Attackers may inject malicious PHP code into the software to execute operations on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:MANTIS-ARB-EXEC1	This signature detects attempts to exploit a vulnerability in Mantis, an open source Web-based bug tracking system. Mantis 0.17.3 and earlier versions are vulnerable. Attackers may send a maliciously crafted URL to cause the Web server to download PHP code from a remote server, allowing the attacker to execute arbitrary code with the permissions of the user that is running the Web server daemon.	medium	sos5.0.0

HTTP:PHP:MANTIS-ARB-EXEC2	This signature detects attempts to exploit a vulnerability in Mantis, an open source Web-based bug tracking system. Mantis 0.17.3 and earlier versions are vulnerable. Attackers may send a maliciously crafted URL to cause the Web server to download PHP code from a remote server, allowing the attacker to execute arbitrary code with the permissions of the user that is running the Web server daemon.	medium	sos5.0.0
HTTP:PHP:MLOG-SCREEN	This signature detects attempts to exploit the vulnerable mlog.phtml script. Attackers may remotely access arbitrary files on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:NULL-CHAR-IN-TAG	This signature detects attempts to exploit a known vulnerability in the PHP Hytertext Processor (PHP) scripting language used on many Unix/POSIX-based web servers. PHP does not properly check for an encoded NULL character (%00) within parameters passed to it. Because PHP does not properly filter the HTML for malicious content, attackers may post HTML that contains malicious code to a PHP-enabled web site. When other users visit the web site, the malicious code runs on their web browser with credentials allowed for the site by that user.	medium	sos5.1.0
HTTP:PHP:PHORUM:ADMIN-PW-CHG	This signature detects attempts to exploit the vulnerable admin.php3 script in Phorum. Attackers may remotely send a maliciously crafted string to the script, change the administrative password of the board without user verification, and access restricted files on the local system.	high	sos5.0.0, sos5.1.0
HTTP:PHP:PHORUM:READ-ACCESS	This signature detects access to the vulnerable read.php3 script installed with Phorum. Because the script does not validate input, attackers may execute arbitrary SQL statements to modify the database contents, insert new entries, create and drop tables, etc.	high	sos5.0.0, sos5.1.0
HTTP:PHP:PHORUM:REMOTE-EXEC	This signature detects attempts to exploit a vulnerability in the PHP Phorum bulletin board system. Attackers may remotely execute arbitrary commands with the privileges of the HTTP server.	high	sos5.0.0, sos5.1.0
HTTP:PHP:PHPBB:HIGHLIGHT-EXEC	This signature detects attempts to exploit a vulnerability in phpBB. Attackers may send a malformed HTTP request to phpBB to force phpBB to execute arbitrary perl commands on the server with Web server permissions.	high	sos5.1.0
HTTP:PHP:PHPBB:HIGHLIGHT-EXEC2	This signature detects attempts to exploit a vulnerability in phpBB. Attackers may send a malformed HTTP request to phpBB to force phpBB to execute arbitrary perl commands on the server with Web server permissions.	high	sos5.1.0
HTTP:PHP:PHPBB:PM_SQL_USR	This signature detects attempts to inject SQL code into a request to phpBB, a popular open-source bulletin board application written in php. Attackers may send a maliciously crafted request that supplies SQL commands to the pm_sql_user parameter, changing database values and escalating client privileges.	low	sos5.1.0

HTTP:PHP:PHPBB:SEARCH-INJECT	This signature detects attempts to exploit a vulnerability in phpBB, an open-source bulletin board package. The search_id parameter in phpBB is vulnerable to SQL injection. Attackers may query private data (such as hashed passwords) then embed the password in a cookie to gain administrative access to the Web site.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:PHPDIG-FILE-INC	This signature detects attempts to exploit a vulnerability in PhpDig 1.6. Attackers may include a malicious 'relative_script_path' parameter in a direct request to the config.php script; this request causes the server to download php code from remote location and execute it. Attackers may execute arbitrary code on the server with permissions of the web server.	high	sos5.1.0
HTTP:PHP:PHPLIB-REMOTE-EXEC	This signature detects attempts to exploit a vulnerability in PHPLILB, a code library that provides support for managing sessions in Web applications. Attackers may remotely submit maliciously crafted Web requests to cause the application to fetch and execute scripts from another host, allowing local access to the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:PHPMYADMIN:SVR-PARAM	This signature detects attempts to exploit a vulnerability in PHPMyAdmin. Attackers may use HTTP form parameters to remotely provide mysql server configuration data. This attack is typically one stage in a multi-stage exploit attempt.	medium	sos5.1.0
HTTP:PHP:PHPNUKE:CID-SQL-INJECT	This signature detects attempts to exploit a vulnerability in PHP-Nuke. Attackers may execute arbitrary SQL commands on a Web server.	medium	sos5.1.0
HTTP:PHP:PHPNUKE:MODULES-DOS	This signature detects attempts to exploit a SQL injection vulnerability in the modules.php script that ships with PHPNuke. PHPNuke 6.0 and earlier are vulnerable. Attackers may produce a process that increases system load on the server, making it unusable until the process is killed.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:PHPROJEKT-INC	This signature detects attempts to exploit a vulnerability in the authform.inc.php script included in the PHProjekt package. Attackers may supply a remote location in the 'path_pre' input parameter to force the target to download and execute arbitrary PHP code from the remote location.	medium	sos5.1.0
HTTP:PHP:PHPWEB-REMOTE-FILE	This signature detects attempts to exploit a vulnerability in phpWebsite. Version 0.8.2 and earlier are vulnerable. Attackers may specify a remote file location for file inclusion to cause phpWebsite to execute arbitrary PHP code; attackers may execute commands with HTTP daemon user permissions.	high	sos5.0.0, sos5.1.0
HTTP:PHP:PMACHINE-INCLUDE	This signature detects attempts to exploit a vulnerability in pMachine, an online publishing application. pMachine version 2.2.1 and other versions are vulnerable. Attackers may send a malicious HTTP request to force the pMachine Web server to execute PHP code from a remote server; commands are executed with web server privileges.	high	sos5.0.0, sos5.1.0

HTTP:PHP:POPPER-OPEN-ADMIN	This signature detects attempts to exploit a vulnerability in popper_mod 1.2.1, a Web-based PHP POP3 e-mail client based on Qpopper. Popper_mod relies on htaccess authentication to authenticate administrators; if htaccess is not used to protect administrator access, popper_mod does not authenticate administrators. Attackers may browse to the /mail/administrator directory to access the administration PHP script and view a complete list of user accounts and passwords, delete accounts, modify accounts, and edit settings.	high	sos5.0.0, sos5.1.0
HTTP:PHP:REDHAT-PIRANHA-PASSWD	This signature detects attempts to exploit the vulnerable passwd.php3 cgi-bin script in the Piranha virtual server package (RedHat Linux 6.2). Because the script does not validate input properly, attackers may authenticate to the Piranha package with the effective ID of the Web server and execute arbitrary commands.	high	sos5.0.0, sos5.1.0
HTTP:PHP:SILENT-STORM-ADMIN	This signature detects attempts to raise the privileges on an account for the Silent Storm PHP Portal.	low	sos5.1.0
HTTP:PHP:UPLOAD-LOCATION	This signature detects a maliciously crafted HTTP POST request. Attackers may use a directory traversal attack within the Content-Disposition field of a POST request to force PHP to execute arbitrary code.	high	sos5.0.0, sos5.1.0
HTTP:PHP:VBULL-CAL-EXEC	This signature detects attempts to exploit a vulnerability in the calender.php script that is included with the VBulletin package. Attackers may run the vbull.c exploit to execute arbitrary commands with Web Server user permissions.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:WOLTAB-SQL-INJ	Any user on the bulletin board can compromise any other user's account by exploiting a vulnerability in board.php. Board.php does not perform proper input validation, and therefore is subject to executing user-supplied SQL statements. This is known to affect Woltlab Burning Board 2.0 RC 1 and earlier versions.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:YABBSE-PKG-EXEC	This signature detects attempts to exploit a vulnerability in Packages.php in YabbSE. YabbSE 1.5.0 and earlier are vulnerable. Attackers may include remote malicious code in Packages.php to include remote malicious code to execute arbitrary commands with Web server privileges.	high	sos5.0.0, sos5.1.0
HTTP:PHP:YABBSE-SSI-INCLUDE	This signature detects attempts to exploit a vulnerability in YabbSE, a PHP/MySQL port of the forum software YaBB (yet another bulletin board). YabbSE versions 1.5.2 and earlier are vulnerable. Attackers may include PHP code in a maliciously crafted URL request; when YabbSE receives the request it runs the PHP code, enabling the attacker to execute arbitrary commands on the server.	medium	sos5.0.0, sos5.1.0
HTTP:PHP:ZENTRACK-CMD-EXEC	This signature detects attacks against the PHP-based zenTrack CRM system. A vulnerability exists in the header.php that holds zenTrack configuration settings. It allows remote command execution as the webserver process privilege. This applies to zenTrack 2.4.1 and below.	high	sos5.0.0, sos5.1.0

HTTP:PKG:ALLAIRE-JRUN-DOS	This signature detects an attempt to launch a denial-of-service (DoS) in Allaire JRun 3.0/3.1. Attackers may send a long string of '.' characters after the /servlet/ prefix in the URL to cause the server to interpret the URL as a very large tree of nonexistent directories and to consume system resources.	high	sos5.0.0, sos5.1.0
HTTP:PKG:DB4WEB-FILE-ACCESS-LIN	This signature detects attempts to exploit a vulnerability in DB4Web (R) Application Server for Windows. Attackers may use a Web browser to download arbitrary files to the target host and obtain system information such as passwords.	medium	sos5.0.0, sos5.1.0
HTTP:PKG:EWAVE-SERVLET-DOS	This signature detects denial-of-service (DoS) attempts against the eWave Servlet JSP. Attackers may remotely send URL requests to cause the Servlet engine to terminate abruptly.	medium	sos5.0.0, sos5.1.0
HTTP:PKG:MOUNTAIN-ORDR-DSCLSR	This signature detects attempts to exploit a vulnerability in Mountain Network Systems Webcart software. Attackers may remotely execute arbitrary commands on the server.	high	sos5.0.0, sos5.1.0
HTTP:PKG:WEBGAIS-REMOTE-EXEC	This signature detects attempt to exploit the websendmail script in WebGais. Attackers may execute arbitrary commands on the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:PROXY:DOUBLE-AT-AT	This signature detects URLs that contain multiple @ characters. Squid/2.3.STABLE5 is vulnerable. Internet Explorer users may use these malicious URLs to evade web proxies and gain direct access to the internet.	medium	sos5.0.0, sos5.1.0
HTTP:REQERR:HEADER-INJECT	This signature detects attempts to exploit an input validation vulnerability in HTTP. Attackers may use encoded CR/LF (carriage return/line feed) characters in an HTTP response header to split HTTP responses into multiple parts, enabling them to misrepresent web content to the recipient.	medium	sos5.0.0, sos5.1.0
HTTP:REQERR:REQ-INVALID-FORMAT	This protocol anomaly is an invalid HTTP request format, such as a request that begins before a previous one ends.	medium	sos5.0.0, sos5.1.0
HTTP:REQERR:REQ-LONG-UTF8CODE	This protocol anomaly is an HTTP request with an exceedingly long UTF8 codes. This may be an attempt to overflow a portion of the Web server, or that a script is being made available to the Web server.	medium	sos5.0.0, sos5.1.0
HTTP:REQERR:REQ-MALFORMED-URL	This protocol anomaly is a malformed URL, such as a Unicode encoded field with non-hex digits or an encoded NULL byte.	medium	sos5.0.0, sos5.1.0
HTTP:SAVANT:GET-DOT1	This signature detects denial-of-service (DoS) attempts against the Savant HTTP server. Savant HTTP server 3.0 and earlier versions are vulnerable. Attackers may send a maliciously crafted HTTP GET request to the Web server to crash the server and create a DoS.	medium	sos5.1.0
HTTP:SPYWARE:DOWNLOAD-ACCEL	This signature detects the use of Download Accelerator, a spyware application.	info	sos5.1.0

HTTP:SPYWARE:GATOR	This signature detects the use of Gator, a spyware application.	info	sos5.1.0
HTTP:SPYWARE:NEW-DOT-NET	This signature detects the use of New.net, a spyware application.	info	sos5.1.0
HTTP:SQL:INJECTION:CMD-CHAIN-1	This signature detects a SQL command sequence in a URL. Because SQL commands are not normally used in HTTP connections, this may indicate a SQL injection attack. However, it may also be a false positive.	medium	sos5.1.0
HTTP:SQL:INJECTION:CMD-CHAIN-2	This signature detects a long SQL command sequence in a URL. Because SQL commands are not normally used in HTTP connections, this may indicate a SQL injection attack.	high	sos5.1.0
HTTP:SQL:INJECTION:CMD-IN-URL	This signature detects SQL commands within a URL. Because SQL commands are not normally used in HTTP connections, this may indicate a SQL injection attack. However, it may be a false positive.	low	sos5.1.0
HTTP:SQL:INJECTION:FACTO-CMS	This signature detects attempts to exploit a vulnerability in the FactoSystem Content Management System (CMS). Attackers may introduce instructions into a SQL query to create a non-authorized CMS account.	medium	sos5.0.0, sos5.1.0
HTTP:SQL:INJECTION:GENERIC	This signature detects specific characters, typically used in SQL, within an HTTP connection. Because these characters are not normally used in HTTP, this may indicate a SQL injection attack. However, it may be a false positive. Some attempts at Cross Site Scripting attacks will also trigger this signature.	info	sos5.1.0
HTTP:SQL:INJECTION:POSTNUKE	This signature detects directory traversal attempts against the modules.php script included with PostNuke. PostNuke versions 0.723 and earlier are vulnerable. Attackers may send a maliciously crafted request to the modules.php to traverse the directory structure and execute SQL queries to the PostNuke database.	medium	sos5.0.0, sos5.1.0
HTTP:SQL:INJECTION:WS2000	This signature detects SQL injection attempts against a WebStore2000 server. Attackers may inject SQL code into the Item_ID parameter of a maliciously crafted request, enabling them to execute arbitrary SQL commands on the WebStore2000 server.	medium	sos5.1.0, sos5.1.0
HTTP:STC:ACROBAT-EXT-OF	This signature detects buffer overflow attempts against Adobe Acrobat Reader. A malicious HTTP server may host an Adobe Acrobat file with an overly long extension; when a client opens this file in Adobe Acrobat Reader, the file triggers a buffer overflow, enabling the server to execute arbitrary code on the client.	high	sos5.1.0
HTTP:STC:ACROBAT-UUEXEC	This signature detects a maliciously crafted PDF file downloaded via HTTP. Attackers may insert certain shell metacharacters at the beginning of a uuencoded PDF file to force Adobe Acrobat to execute arbitrary commands upon loading the file.	high	sos5.1.0

HTTP:STC:EICAR-DOWNLOAD	This signature detects the EICAR antivirus test file downloaded via HTTP.	info	sos5.1.0
HTTP:STC:EXCEL-CELL-OF	This signature detects a maliciously crafted Microsoft Excel file downloaded via HTTP. Attackers may supply an Excel document that contains an overly long Cell Length field to overflow the buffer and execute arbitrary code on the client.	high	sos5.1.0
HTTP:STC:IE:CONT-LOC-ZON-BYPASS	This signature detects attempts to circumvent a security zone feature that warns when executable files are downloaded. WindowsXP Service Pack 2 and Internet Explorer 6 are vulnerable. Attackers may trick a user into downloading a file that the user did not know was executable. Similarly, viruses and worms may use this method to download themselves onto target computers.	medium	sos5.1.0
HTTP:STC:IE:EXEC-CMD-FILE-SPOOF	This signature detects attempts to exploit a vulnerability in the way that Internet Explorer handles the Javascript execCommand function. Attackers may trick a user into saving a file that the user thinks is HTML, but is actually an executable file.	medium	sos5.1.0
HTTP:STC:WINAMP:CDDA-OF	This signature detects the download of a maliciously crafted WinAmp playlist file. Using WinAmp to open this file may execute arbitrary code.	high	sos5.1.0
HTTP:STC:WINAMP:CDDA-OF2	This signature detects the download of a maliciously crafted WinAmp playlist file. Using WinAmp to open this file may execute arbitrary code.	high	sos5.1.0
HTTP:TOMCAT:JSP-AS-HTML	This signature detects attempts to exploit a vulnerability in Apache Tomcat. Apache Tomcat 3.3.1 and earlier are vulnerable. Attackers may send a maliciously crafted URL to cause the server to parse a .jsp file as HTML code and display the JSP code, allowing attackers to retrieve normally inaccessible files.	medium	sos5.1.0
HTTP:TOMCAT:SERVLET-DEVICE-DOS	This signature detects denial-of-service (DoS) attempts against Apache Group Tomcat Server. Attackers may request a device name from the /examples/servlet directory to render the server inaccessible. This signature also detects attempts to run neuter.c and similar exploits.	medium	sos5.0.0, sos5.1.0
HTTP:TUNNEL:ALTNET-OVER-HTTP	This signature detects attempts to connect to a Altnet server over HTTP. Altnet is a component of Kazaa, a common Peer to Peer file sharing system. Users may be attempting to download files.	info	sos5.1.0
HTTP:TUNNEL:CHAT-AOL-IM	This signature detects AOL Instant Messenger Proxy over HTTP. Users may use proxy connections over the HTTP port to circumvent firewall policies.	info	sos5.1.0
HTTP:TUNNEL:CHAT-MSN-IM	This signature detects MSN Instant Messenger over HTTP. Users may use proxy connections over the HTTP port to circumvent firewall policies.	info	sos5.1.0

HTTP:TUNNEL:CHAT-YIM	This signature detects Yahoo Instant Messenger Proxy over HTTP. Users may use proxy connections over the HTTP port to circumvent firewall policies.	info	sos5.1.0
HTTP:TUNNEL:HTTPTUNNEL-URL	This signature detects traffic from the HTTP Tunnel utility. HTTP Tunnel masquerades a network session in HTTP traffic.	low	sos5.1.0
HTTP:TUNNEL:KAZAA-OVER-HTTP	This signature detects attempts to connect to a Kazaa server over HTTP. Kazaa is a common Peer to Peer file sharing system. Users may be attempting to download files.	info	sos5.1.0
HTTP:TUNNEL:SSH	This signature detects SSH over HTTP. Attackers may send SSH over the HTTP port to circumvent firewall policies.	info	sos5.1.0
HTTP:TUNNEL:TELNET	This signature detects Telnet over HTTP. Attackers may send Telnet over the HTTP port to circumvent firewall policies.	info	sos5.1.0
HTTP:WASD:CONF-ACCESS	This signature detects attempts to exploit a vulnerability in the WASD HTTP Server for OpenVMS. Default installations of 1.0 and earlier are vulnerable. Attackers may download the configuration file for the server and obtain information on the ACL and internal directory structure.	medium	sos5.0.0, sos5.1.0
HTTP:WASD:DIR-TRAV	This signature detects directory traversal attempts against WASD HTTP Server for OpenVMS. WASD version 1.0 and earlier are vulnerable. Attackers may navigate to any directory on the server.	medium	sos5.0.0, sos5.1.0
HTTP:WEBLOGIC:URL-REVEAL-SRC	This signature detects attempts to exploit a vulnerability in Bea Weblogic. Version V6.1 Service Pack 2 on Windows 2000 Server is vulnerable. Attackers may append the string "%00x" to a URL request to read the contents of a .jsp file.	medium	sos5.0.0, sos5.1.0
HTTP:WEBLOGIC:WEBROOT	This signature detects attempts to exploit a vulnerability in Bea Weblogic. Version V6.1 Service Pack 2 on Windows 2000 Server is vulnerable. Attackers may append the string "%00.jsp" to a normal .html request, causing a compiler error that prints the path to the physical web root.	medium	sos5.0.0, sos5.1.0
HTTP:WEBPLUS:DIR-TRAVERSAL	This signature detects attempts to exploit the input validation vulnerability in the main CGI in TalentSoft Web+, an e-commerce storefront provider. Attackers may pass a script variable that specifies a filepath to the webpsvr daemon, and gain access to any file on the system that the UID of the Web server has access to.	medium	sos5.0.0, sos5.1.0
HTTP:WEBSphere:VER-DOS	This signature detects denial-of-service (DoS) attempts against the caching proxy in IBM WebSphere Edge Server. Version 2.0 is vulnerable. Attackers may send a maliciously crafted HTTP GET request that does not have a proper version identifier to crash the proxy service and render the proxy unusable.	medium	sos5.1.0

HTTP:WIN-CMD:WIN-CMD-EXE	This signature detects the Windows command 'cmd.exe' within a URL. This command does not normally appear in a URL, and may indicate an attempt to compromise the system.	medium	sos5.0.0, sos5.1.0
HTTP:WIN-CMD:WIN-RGUEST	This signature detects the Windows command 'rguest.exe' within a URL. This command does not normally appear in a URL, and may indicate an attempt to compromise the system.	medium	sos5.0.0, sos5.1.0
HTTP:WIN-CMD:WIN-WGUEST	This signature detects the Windows command 'wguest.exe' within a URL. This command does not normally appear in a URL, and may indicate an attempt to compromise the system.	medium	sos5.0.0, sos5.1.0
HTTP:XSS:HDR-REFERRER	This signature detects attempts to exploit a cross-site scripting vulnerability. Attackers may embed malicious HTML tags within the HTTP Referrer header; because some web servers and server-side applications parse this data incorrectly, attackers can successfully execute a cross-site scripting attack.	low	sos5.0.0, sos5.1.0
HTTP:XSS:HTML-SCRIPT-IN-URL-PRM	This signature detects attempts at cross site scripting attacks. Attackers may create a malicious Web site that includes HTML embedded in the hyperlinks, which might violate site security settings. Attackers may then view the Web cookies from your computer; Web cookies typically contain sensitive information such as usernames, passwords, credit card numbers, social security numbers, bank accounts, etc.	medium	sos5.1.0
HTTP:XSS:HTML-SCRIPT-IN-URL-PTH	This signature detects cross site scripting attacks. Attackers may create a malicious Web site that includes HTML embedded in the hyperlinks, which might violate site security settings. Attackers may then view the Web cookies from a target computer. Web cookies typically contain sensitive information such as usernames, passwords, credit card numbers, social security numbers, and bank account numbers.	medium	sos5.1.0
HTTP:XSS:URL-IMG-XSS	This signature detects HTML tags in URLs that include Javascript. Because tags should never be present in URLs, the presence of Javascript in such a URL is a clear indication of a Cross-Side Scripting (XSS) attack. XSS attacks are typically Web browser-independent.	high	sos5.1.0
IMAP:FAILURE:BRUTE-FORCE	This protocol anomaly is multiple login failures within a short period of time between a unique pair of hosts.	high	sos5.1.0
IMAP:IPSWITCH:DELE-OF	This signature detects buffer overflow attempts against IPSwitch IMAP server. Attackers may send an overly long delete command (DELE) to overflow the buffer and take complete control of the server.	high	sos5.1.0
IMAP:OVERFLOW:COMMAND	This protocol anomaly is an IMAP command that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0

IMAP:OVERFLOW:FLAG	This protocol anomaly is an IMAP flag that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:IMAP4-LSUB-OF	This signature detects buffer overflow attempts against the IMAP package included with several Linux distributions. Attackers may send a long string to the IMAP package to execute code with daemon-level permissions.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:LINE	This protocol anomaly is an IMAP line (from the client to the server) that is too long. This may indicate a buffer overflow attempt. NOTE: Long lines are parsed, which may generate other IMAP overflow errors.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:LIT_LENGTH_OVERFLOW	This protocol anomaly is an IMAP literal that specifies more octets than the user-defined maximum. A literal is a sequence of zero or more octets. The default maximum number of octets is 65535.	high	sos5.1.0
IMAP:OVERFLOW:MAILBOX	This protocol anomaly is an IMAP mailbox name that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:PASS	This protocol anomaly is an IMAP user password that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:REFERENCE	This protocol anomaly is an IMAP reference field that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:TAG	This protocol anomaly is an IMAP tag field that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
IMAP:OVERFLOW:USER	This protocol anomaly is an IMAP user name that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
IMAP:REQERR:INVALID_LITERAL_LEN	This protocol anomaly is a literal that specifies a number of octets containing a character that is not 0 or 9.	high	sos5.1.0
IMAP:REQERR:REQ-INVALID-TAG	This protocol anomaly is an invalid IMAP tag, i.e., a tag that begins with a white space or contains non-alphanumeric characters. This may indicate a nonstandard IMAP client or command line access to an IMAP server.	medium	sos5.0.0, sos5.1.0
IMAP:REQERR:REQ-UNEXPECTED-ARG	This protocol anomaly is an IMAP command with too many arguments. This may indicate a nonstandard IMAP client or command line access to an IMAP server.	medium	sos5.0.0, sos5.1.0
Key	Description	Severity	Versions
MS-RPC:DCOM:SVRNAME-2LONG	This protocol anomaly is a DCOM servername that is longer than 32 octets in unicode.	critical	sos5.1.0
MS-RPC:EPDUMP-SCAN	This anomaly detects a client enumerating MSRPC endpoints on a windows server. This may indicate a probing scan prior to a more sophisticated attack.	low	sos5.1.0

MS-RPC:ERR:CL-PTYPE-IN-CO-PDU	This protocol anomaly is an MSRPC connection-oriented message with a packet type that is allowed only in Connectionless PDUs.	medium	sos5.1.0
MS-RPC:ERR:CO-PTYPE-IN-CL-PDU	This protocol anomaly is a connectionless MSRPC message with a packet type that is allowed only in connection-oriented PDUs.	medium	sos5.1.0
MS-RPC:ERR:EPM-INV-LHS-LEN	This protocol anomaly is an EPM message with an LHS length that is larger than the rest packet length.	high	sos5.1.0
MS-RPC:ERR:EPM-INV-OP-NUM	This protocol anomaly is an invalid EPM operation number.	medium	sos5.1.0
MS-RPC:ERR:EPM-INV-RHS-LEN	This protocol anomaly is an EPM message with an RHS length that is larger than the rest packet length.	high	sos5.1.0
MS-RPC:ERR:EPM-INV-TOWER-LEN	This protocol anomaly is an EPM message with a tower length that is larger than 8192 bytes, or larger than the rest fragment length.	high	sos5.1.0
MS-RPC:ERR:EPM-WRONG-LHS-LEN	This protocol anomaly is an EPM packet with a UUID LHS length that is not equal to 19.	high	sos5.1.0
MS-RPC:ERR:EPM-WRONG-RHS-LEN	This protocol anomaly is an EPM message with an RHS length that is larger than the rest packet length.	high	sos5.1.0
MS-RPC:ERR:EPM-WRONG-TOWER-LEN	This protocol anomaly is an EPM message with a tower length that is inconsistent with message's LHS and RHS lengths.	high	sos5.1.0
MS-RPC:ERR:FRAG-BIGGER-THEN-NEG	This protocol anomaly is a MSRPC fragment length that is larger than the negotiated maximum.	medium	sos5.1.0
MS-RPC:ERR:FRAG-LEN-TOO-SMALL	This protocol anomaly is an MSRPC fragment length that is less than the common header length.	high	sos5.1.0
MS-RPC:ERR:INV-AUTH-LEN	This protocol anomaly is an MSRPC message with an authentication length that is larger than the entire MS-RPC message payload length.	high	sos5.1.0
MS-RPC:ERR:INV-AUTH-PAD-LEN	This protocol anomaly is an MSRPC message with authentication padding length plus authentication section length that is larger than the entire MSRPC message payload length.	high	sos5.1.0
MS-RPC:ERR:INV-PTYPE	This protocol anomaly is an MSRPC message that contains an invalid packet type value.	medium	sos5.1.0
MS-RPC:ERR:LEN-CONFLICT	This protocol anomaly is an MSRPC connectionless message with a fragment length that conflicts with the common header length and the whole message length.	high	sos5.1.0
MS-RPC:ERR:RESPONSE-NO-REQ	This protocol anomaly is an MSRPC response that precedes the request.	medium	sos5.1.0

MS-RPC:ERR:SHORT-MSG	This protocol anomaly is an incomplete MSRPC message.	high	sos5.1.0
MS-RPC:LOC-SVC-OF	This signature detects attempts to exploit a flaw in the Windows DCE RPC Locator service. This service is turned on by default on all Windows NT 4 and Windows 2000 Domain Controllers, or can be turned on manually on all Windows NT, 2000, and XP systems. Attackers can deny the service of the locator, causing network-wide outages, or take control of the service and run code of choice.	high	sos5.1.0
MS-RPC:LSASS:MAL-OPCODE	This signature detects attempts to exploit a known vulnerability in Microsoft Windows LSASS (Local Security Authority Subsystem Service). Attackers may remotely run arbitrary code on the target system. Note: This vulnerability is exploited by many worms.	critical	sos5.1.0
MS-RPC:LSASS:OVERSIZED-FRAG	This signature detects attempts to remotely attack a known vulnerability in the Microsoft Windows LSASS (Local Security Authority Subsystem Service). A successful attack could run code of an attacker's choice on the target system. By supplying an oversized fragment to the LSASS service, a buffer can be overflowed that can result in remote code execution.	critical	sos5.1.0
MS-RPC:MSRPC-ISYSACTIVATE-RACE	This protocol anomaly is too many DCE/RPC ISystemActivate requests. Excessive requests can cause a denial-of-service (DoS) in the RPCSS module.	high	sos5.1.0
MS-RPC:NOOP-SLIDE-RPC-REQ	This signature detects Unicode NOOP sleds in an RPC request. Because these patterns are usually malicious, they might indicate an attack.	medium	sos5.1.0
MS-RPC:SAMR-ACCESS-DENIED	This signature detects failed attempts to connect to the Security Account Manager Remote (SAMR) service on Windows. Attackers may be probing your server for vulnerabilities, as a successful login to this service provides important information such as administrator account details, default domain names, open users, and active groups. However, because system administrators also use the SAMR service legitimately, this signature may also detect non-malicious activity.	info	sos5.1.0
MS-RPC:SAMR-ACCESS-REQUEST	This signature detects attempts to connect to the Security Account Manager Remote (SAMR) service on Windows. Attackers may be probing your server for vulnerabilities, as a successful login to this service provides important information such as administrator account details, default domain names, open users, and active groups. However, because system administrators also use the SAMR service legitimately, this signature may also detect non-malicious activity.	low	sos5.1.0

MS-RPC:WKST-SVC-OFLOW	This protocol anomaly is a suspiciously long argument for the NetrValidateName, NetrValidateName2, or NetrAddAlternateComputerName functions requested using a named-pipe transaction. An unauthenticated user may exploit this vulnerability on Windows 2000/XP servers to execute arbitrary code with system-level privileges.	critical	sos5.1.0
NETBIOS:ACCESS:ADMIN	This signature detects attempts to exploit a null session vulnerability in NETBIOS SMB protocols. Attackers may initiate SMB sessions with no user name or password, obtain the remote administrator share on the server, and use this information to plan further attacks.	medium	sos5.1.0
NETBIOS:ACCESS:C-DRIVE	This signature detects attempts to exploit a null session vulnerability in NETBIOS SMB protocols. Attackers may initiate SMB sessions with no user name or password, obtain the C Drive share on the server, and use this information to plan further attacks.	medium	sos5.1.0
NETBIOS:ACCESS:D-DRIVE	This signature detects attempts to exploit a null session vulnerability in NETBIOS SMB protocols. Attackers may initiate SMB sessions with no user name or password, obtain the D Drive share on the server, and use this information to plan further attacks.	medium	sos5.1.0
NETBIOS:NBDS:BAD_LABEL_FORMAT	This protocol anomaly is label for the second level encoding of a Netbios name that contains a pointer.	high	sos5.1.0
NETBIOS:NBDS:INVALID:1STLVL_ENC	This protocol anomaly is an invalid first level encoding of a Netbios name.	medium	sos5.1.0
NETBIOS:NBDS:INVALID:DGM_LEN	This protocol anomaly is a Netbios datagram header with a DGM_LENGTH field value that is bigger than the packet length.	high	sos5.1.0
NETBIOS:NBDS:INVALID:HDR_FLGS	This protocol anomaly is a Netbios datagram header with a FLAGS field that contains non-zero values for bits 0-3.	high	sos5.1.0
NETBIOS:NBDS:INVALID:LABEL_LEN	This protocol anomaly is a label for the second level encoding of a netbios name; the label length is larger than 63, or the label is the first label and the length is not 32.	high	sos5.1.0
NETBIOS:NBDS:INVALID:MSG_TYPE	This protocol anomaly is a Netbios datagram header with a MSG_TYPE field value that is invalid.	high	sos5.1.0
NETBIOS:NBDS:INVALID:PROTO	This protocol anomaly is a Netbios message with a USER_DATA section that is less than the size of SMB header, or the protocol field of the SMB header does not start with 0xff 'S' 'M' 'B'.	high	sos5.1.0
NETBIOS:NBDS:OVERFLOW:MSG	This protocol anomaly is a Netbios datagram that is bigger than 1064.	high	sos5.1.0
NETBIOS:NBDS:OVERFLOW:NAME	This protocol anomaly is a Netbios name that is longer than 255.	high	sos5.1.0

NETBIOS:NBNS:C2S_AA_FLAG	This protocol anomaly is query message with an NM_FLAGS field containing an authoritative answer flag that is set.	medium	sos5.1.0
NETBIOS:NBNS:C2S_RESPONSE	This protocol anomaly is query message with an OPCODE field containing an response flag that is set.	medium	sos5.1.0
NETBIOS:NBNS:CLASS-UNKNOWN	This protocol anomaly is an invalid value in the QUESTION_CLASS field of the question section or in the RR_CLASS field of the resource record header.	medium	sos5.1.0
NETBIOS:NBNS:INVALID:FIRST-ENC	This protocol anomaly is an invalid first level encoding of a Netbios name.	high	sos5.1.0
NETBIOS:NBNS:INVALID:HDR-CNT	This protocol anomaly is a 1) a query message with ARCOUNT (answer count) or NSCOUNT (number of records in the authority section of a name service packet) fields of the header that are not zero, or 2) a response message with a QDCOUNT (number of entries in the question section) that is not zero.	medium	sos5.1.0
NETBIOS:NBNS:INVALID:HDR-OPCODE	This protocol anomaly is a header with an OPCODE field value that is not 0, 5, 6, 7, or 8.	high	sos5.1.0
NETBIOS:NBNS:INVALID:HDR-Z	This protocol anomaly is a Netbios name header with a NM_FLAGS field that contains nonzero values for bit 4 or bit 5.	medium	sos5.1.0
NETBIOS:NBNS:INVALID:LABEL-LEN	This protocol anomaly is a label for the second level encoding of a Netbios name that has a label length larger than 63, or the label is the first label and the length is not 32.	high	sos5.1.0
NETBIOS:NBNS:INVALID:NAME-FLGS	This protocol anomaly is a Netbios name header with a NM_FLAGS field that contains nonzero values for bits 3-15.	medium	sos5.1.0
NETBIOS:NBNS:INVALID:PTR	This protocol anomaly is a pointer offset in the second level encoding of a Netbios name that exceeds the message length (the pointer is pointing out of the message).	high	sos5.1.0
NETBIOS:NBNS:INVALID:RRNB-FLG	"This protocol anomaly is a type node status response message		
NETBIOS:NBNS:NAME_TOO_LONG	This protocol anomaly is a Netbios name that is longer than 255.	high	sos5.1.0
NETBIOS:NBNS:POINTER_LOOP	This protocol anomaly is a second level encoding of a Netbios name that contains more nested pointers than the user-defined maximum. Default setting for the sc_nbname_pointer_loop_limit is 8.	medium	sos5.1.0
NETBIOS:NBNS:RESCODE:FORMAT_ERR	This protocol anomaly is Netbios name response with an RCODE that indicates the request has an invalid format.	high	sos5.1.0
NETBIOS:NBNS:S2C_QUERY	This protocol anomaly is a Netbios name response header with an OPCODE field that contains an unset response bit.	medium	sos5.1.0

NETBIOS:NBNS:SHORT_MSG	This protocol anomaly is a Netbios name message that is shorter than expected.	medium	sos5.1.0
NETBIOS:NBNS:TYPE_UNKNOWN	This protocol anomaly is an invalid value in 1) the QUESTION_TYPE field in the question section or 2) the RR_TYPE field in the resource record header.	high	sos5.1.0
P2P:AUDIT:GNUTELLA-BYE-TTL	This protocol anomaly is a Gnutella BYE message that does not contain a TTL of 1 and a HOPS of 0.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-EOL	This protocol anomaly is a Gnutella message that does not use the end-of-line (EOL) terminator characters <CR><LF>.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-HDR-ATRB	This protocol anomaly is a Gnutella message with a header line that does not have a value for an attribute; a blank space exists after the attribute colon.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-HTTP-GET	This protocol anomaly is a Gnutella GET command that does not use the expected syntax. Correct syntax is: GET /get/<File Index>/<File Name> HTTP/1.1<CR><LF>.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-LINE	This protocol anomaly is a Gnutella message with a line length that exceeds the user-defined maximum number of bytes. The default line length is 2048.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-MESSAGE	This protocol anomaly is a Gnutella message with a payload type that is not defined in the Gnutella RFC.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-MSG	This protocol anomaly is a Gnutella message with a payload length that exceeds 4096 bytes.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-OK-RESP	This protocol anomaly is a Gnutella client response that does not use the expected syntax. Correct syntax for Gnutella 0.6 is: GNUTELLA/0.6 200 OK<CR><LF>.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-PING-LEN	This protocol anomaly is a Gnutella 0.4 PING message that has a nonzero payload length.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-PONG-LEN	This protocol anomaly is a Gnutella PONG message that has an invalid payload length. Gnutella 0.4 PONG messages should have exactly 14 bytes; Gnutella 0.6 PONG messages should have a minimum of 14 bytes.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-PUSH-LEN	This protocol anomaly is a Gnutella PUSH message with a payload length that is less than 26 bytes.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-QUERY	This protocol anomaly is a Gnutella QUERY message with a payload length that exceeds the user-defined maximum number of bytes. The default line length is 256.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-RTABLE-UPD	This protocol anomaly is a Gnutella ROUTE_TABLE_UPDATE message with a payload length of 0 bytes.	info	sos5.1.0

P2P:AUDIT:GNUTELLA-SEARCH	This protocol anomaly is a Gnutella message with a search criteria field that does not end with a NULL character.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-SVR-RESP	This protocol anomaly is a Gnutella server response that does not use the expected syntax. Correct syntax for Gnutella 0.4 is: GNUTELLA OK<CR><LF>; correct syntax for Gnutella 0.6 is: GNUTELLA/0.6 200 OK<CR><LF>.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-TTL	This protocol anomaly is a Gnutella message with a TTL that exceeds the user-defined maximum. The default TTL is 8. The Gnutella RFC recommends an 8 to 10 TTL maximum for Gnutella messages.	info	sos5.1.0
P2P:AUDIT:GNUTELLA-UNSUP-VER	This protocol anomaly is a Gnutella message with a connect string that does not conform to Gnutella RFC or the requesting Gnutella version is not 0.4 or 0.6.	info	sos5.1.0
P2P:BITTORRENT:TRACKER-QUERY	This signature detects requests to a BitTorrent tracker website. Users may be querying the tracker to look for files to download.	info	sos5.1.0
P2P:BITTORRENT:TRACKER-SCRAPE	This signature detects 'scrape' requests to a BitTorrent tracker website. Users may be querying the tracker to look for files to download.	info	sos5.1.0
P2P:DC:DC-PP-ACTIVE	This signature detects use of the Direct Connect Plus Plus (DC++) file sharing client.	info	sos5.1.0
P2P:EDONKEY:CLIENT-VER-CHECK	This signature detects version checks by eDonkey 2000, a peer-to-peer file sharing client. The eDonkey client occasionally checks its own version number to ensure that the client is current.	info	sos5.1.0
P2P:GNUTELLA:CONNECT	This signature detects Gnutella client connection requests. Because Gnutella does not use a fixed port number, this signature searches TCP connections to port 1024 and higher by default.	info	sos5.1.0
P2P:GNUTELLA:CONNECTION-OK	This signature detects GNUTella server responses to a connection request. Because GNUTella does not use a fixed port number, this signature searches TCP connections to port 1024 and higher by default.	info	sos5.1.0
P2P:GNUTELLA:CONNECTION-OK-V06	This signature detects Gnutella server responses to a connection request. Because Gnutella does not use a fixed port number, this signature searches TCP connections to port 1024 and higher by default.	info	sos5.1.0
P2P:MLDONKEY:CLIENT-ACTIVE	This signature detects activity by the peer-to-peer (P2P) file sharing client MLDonkey, a multi-protocol P2P file sharing application.	info	sos5.1.0
P2P:SKYPE:VERSION-CHECK	This signature detects a Skype client request (to a central server) that checks for the latest version of the client software.	info	sos5.1.0

P2P:WINMX:CLIENT-MATCHMAKE-DNS	This signature detects a WinMX client performing DNS lookups for matchmaking servers. WinMX is a peer-to-peer file sharing client that tests firewall rules and reverse-connectivity to determine the most effective way to share files. WinMX queries a matchmaking server to obtain Supernode lists, which enable the WinMX client to share files.	info	sos5.1.0
P2P:WINMX:CLIENT-NET-PRB-DNS	This signature detects a WinMX client performing DNS lookups for hosts that WinMX will probe for connectivity. WinMX is a peer-to-peer file sharing client that tests firewall rules and reverse-connectivity to determine the most effective way to share files.	info	sos5.1.0
P2P:WINMX:CLIENT-VER-CHK	This signature detects an initial connection by WinMX, a peer-to-peer file sharing client. WinMX queries a Web site for new versions of the WinMX client software.	info	sos5.1.0
P2P:WINMX:CLIENT-VER-CHK-DNS	This signature detects attempts to obtain the IP address of the host that tracks WinMX client versions. WinMX is a peer-to-peer file sharing client.	info	sos5.1.0
POP3:DOS:MDAEMON-POP-DOS	This signature detects denial-of-service attempts against the Mdaemon POP3 Server. Mdaemon v.6.0.7 and earlier versions are vulnerable. Attackers may send a maliciously crafted DELE or UIDL request to the POP3 daemon to crash the POP3, SMTP, and IMAP services.	medium	sos5.0.0, sos5.1.0
POP3:ERROR:BOUNDARY_MISSING	This protocol anomaly is a message with a multipart content type but no boundary.	high	sos5.1.0
POP3:EXT:DOT-386	This signature detects e-mail attachments that have the extension .386 and were received via POP3. Because .386s (Windows Enhanced Mode Driver) files contain executable code, this may indicate an incoming e-mail virus. Attackers may create malicious executables, tricking users into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-ADE	This signature detects e-mail attachments that have the extension .ade and were received via POP3. Because .ADEs (Microsoft Access Project Extension) files can contain macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the macros and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-ADP	This signature detects e-mail attachments that have the extension .adp and were received via POP3. Because .ADPs (Microsoft Access Project) files can contain macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the macros and infecting the system.	medium	sos5.1.0

POP3:EXT:DOT-BAS	This signature detects e-mail attachments that have the extension .bas and were received via POP3. Because .BASs (Microsoft Visual Basic Class Module) files contain executable code, this may indicate an incoming e-mail virus. Attackers may create malicious executables, tricking users into executing the file and infecting the system.	low	sos5.1.0
POP3:EXT:DOT-BAT	This signature detects e-mail attachments with the extension '.bat' received via POP3. This may indicate an incoming e-mail virus. .BATs (executable files) contain one or more scripts. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-CHM	This signature detects e-mail attachments that have the extension .chm and were received via POP3. Because .CHMs (Compiled HTML Help File) files can contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the files and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-CMD	This signature detects e-mail attachments with the extension '.cmd' sent via POP3. This may indicate an incoming e-mail virus. CMD files contain commands that when executed can cause significant damage to a windows system.	high	sos5.1.0
POP3:EXT:DOT-COM	This signature detects e-mail attachments with the extension '.com' received via POP3. This may indicate an incoming e-mail virus. .COMs (executable files) contain one or more scripts. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-CPL	This signature detects e-mail attachments with the extension '.cpl' received via POP3. This may indicate an incoming e-mail virus. CPLs (Control Panel elements) are standard Microsoft Windows files that contain Windows Control Panel settings. Attackers may hide malicious executables within a CPL file, tricking users into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-CRT	This signature detects e-mail attachments that have the extension .crt and were received via POP3. Because .CRTs (Security Certificate) files can contain executable code, this may indicate an incoming e-mail virus. Attackers may create malicious executable code, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-EXE	This signature detects e-mail attachments with the extension '.exe' sent via POP3. This may indicate an incoming e-mail virus. EXEs (Executable files) contain one or more scripts. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	high	sos5.1.0

POP3:EXT:DOT-GRP	This signature detects GRP files sent over POP3. GRP files can contain Windows Program Group information, and may be exploited by malicious users to deposit instructions or arbitrary code on a target's system. User involvement is required to activate GRP files; typically they are attached to a harmless-appearing e-mail message.	medium	sos5.1.0
POP3:EXT:DOT-HLP	This signature detects e-mail attachments that have the extension .hlp and were received via POP3. Because .HLPs (Help File) files can contain macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the macros and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-HT	This signature detects e-mail attachments with the extension '.ht' sent via POP3. This may indicate an incoming e-mail virus or other attack. HT files contain configuration information for the Hyperterm console program, shipped with every Windows operating system since Windows 95. It is the default handler program for .ht files. A recent vulnerability in Hyperterm could allow an attacker to take control of your computer via an infected .ht file. These files are not normally sent via e-mail.	medium	sos5.1.0
POP3:EXT:DOT-HTA	This signature detects e-mail attachments with the extension .hta received using POP3. This may indicate an incoming e-mail virus. HTA files are HTML application files that can be executed by a web browser. Generally, HTA files are not sent via e-mail. As a general network security precaution, ensure that all users are aware of the dangers of sending and receiving binary files in e-mail attachments.	medium	sos5.1.0
POP3:EXT:DOT-INF	This signature detects e-mail attachments that have the extension .inf and were received via POP3. Because .INFs (Setup Information) files contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-INS	This signature detects e-mail attachments that have the extension .ins and were received via POP3. Because .INSs (Internet Naming Service) files contain configuration parameters, this may indicate an incoming e-mail virus. Attackers may include malicious configurations, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-ISP	This signature detects e-mail attachments that have the extension .isp and were received via POP3. Because .ISPs (Internet Communication Settings) files contain configuration parameters, this may indicate an incoming e-mail virus. Attackers may include malicious configurations, tricking users into executing the file and infecting the system.	high	sos5.1.0

POP3:EXT:DOT-JS	This signature detects e-mail attachments that have the extension .js and were received via POP3. Because .JSs (JavaScript File) files contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-JSE	This signature detects e-mail attachments that have the extension .jse and were received via POP3. Because .JSEs (JavaScript Encoded) files contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-LNK	This signature detects e-mail attachments that have the extension .lnk and were received via POP3. Because .LNKs (Windows link) files can point to any program, this may indicate an incoming e-mail virus. Attackers may create a link pointing to a dangerous program, tricking users into executing the link and affecting the system.	medium	sos5.1.0
POP3:EXT:DOT-MDB	This signature detects e-mail attachments that have the extension .mdb and were received via POP3. Because .MDBs (MS Access Application) files can contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-MDE	This signature detects e-mail attachments that have the extension .mde and were received via POP3. Because .MDEs (Microsoft Access MDE database) files can contain scripts and macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-MSC	This signature detects e-mail attachments that have the extension .msc and were received via POP3.	high	sos5.1.0
POP3:EXT:DOT-MSI	This signature detects e-mail attachments with the extension .msi received via POP3. This may indicate an incoming e-mail virus. .MSIs (Microsoft Windows Installer Package) contain executable code. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-MSP	This signature detects e-mail attachments with the extension .msp received via POP3. This may indicate an incoming e-mail virus. .MSPs (Microsoft Windows Installer Patch) contain executable code. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0

POP3:EXT:DOT-OCX	This signature detects e-mail attachments that have the extension .ocx and were received via POP3. Because .OCXs (Object Control Extension) files can contain multiple scripts, this may indicate an incoming e-mail virus. Attackers may create malicious executables, tricking users into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-PCD	This signature detects e-mail attachments that have the extension .pcd and were received via POP3. Because .PCDs (Photo CD MS Compiled Script) files can contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-PIF	This signature detects e-mail attachments with the extension '.pif' sent via POP3. This may indicate an incoming e-mail virus. PIFs (Program Information Files) are standard Microsoft Windows files that contain start up properties for DOS applications. Attackers may hide malicious executables within a PIF file, tricking users into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-REG	This signature detects e-mail attachments that have the extension .reg and were received via POP3. Because .REGs (Registry Entries) files contain entries for the Registry, this may indicate an incoming e-mail virus. Attackers may create malicious entries, tricking users into executing the file and infecting the system.	medium	sos5.1.0
POP3:EXT:DOT-SCR	This signature detects e-mail attachments with the extension '.scr' sent via POP3. This may indicate an incoming e-mail virus. SCRs (ScreenSaver files) are renamed '.exe' files containing executable code. Attackers may disguise malicious executables to appear as harmless screensaver files, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-SCT	This signature detects e-mail attachments with the extension .sct received via POP3. This may indicate an incoming e-mail virus. .SCTs (Windows Script Component) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-URL	This signature detects e-mail attachments with the extension .url received via POP3. This may indicate an incoming e-mail virus. .URLs (Internet Shortcut) contain a link to a web location. Attackers may create a malicious shortcut, tricking the user into executing the file and send the user to a malicious website.	high	sos5.1.0
POP3:EXT:DOT-VB	This signature detects e-mail attachments with the extension .vb received via POP3. This may indicate an incoming e-mail virus. .VBs (VBScript File) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	high	sos5.1.0

POP3:EXT:DOT-VBE	This signature detects e-mail attachments with the extension .vbe received via POP3. This may indicate an incoming e-mail virus. .VBEs (VBScript Encoded Script File) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-VBS	This signature detects e-mail attachments with the extension '.vbs' sent via POP3. This may indicate an incoming e-mail virus. VBSs (Visual Basic files) contain one or more executable scripts. Attackers may create malicious VB files, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-WMF	This signature detects Metafiles files sent over POP. Windows Metafiles and Enhanced Metafiles files can exploit a Windows GDI vulnerability and may be exploited by malicious users to deposit instructions or arbitrary code on a target's system. User involvement is required to activate Metafiles; typically they are attached to a harmless-appearing e-mail message.	medium	sos5.1.0
POP3:EXT:DOT-WSC	This signature detects e-mail attachments with the extension .wsc received via POP3. This may indicate an incoming e-mail virus. .WSCs (Windows Script Component) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-WSF	This signature detects e-mail attachments with the extension .wsf received via POP3. This may indicate an incoming e-mail virus. .WSFs (Windows Script File) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-WSH	This signature detects e-mail attachments with the extension .wsh received via POP3. This may indicate an incoming e-mail virus. .WSHs (Windows Script Host Settings File) contain configuration parameters. Attackers may create malicious configurations, tricking the user into executing the file and infecting the system.	high	sos5.1.0
POP3:EXT:DOT-ZIP	This signature detects e-mail attachments with the extension .zip received using POP3. This may indicate an incoming e-mail virus. Zip files are compressed files that can contain one or more executables. Attackers may compress malicious executables within a .zip file, tricking unsuspecting users into executing the file and infecting the system. Because Zip files are frequently used for non-malicious purposes, this signature can generate false positives. As a general network security precaution, ensure that all users are aware of the dangers of sending and receiving binary files in e-mail attachments.	low	sos5.1.0

POP3:EXT:DOUBLE-DOT-DOT	This signature detects e-mail attachments that contain two file extensions. Attackers or viruses may send e-mail attachments that use two file extensions to disguise the actual file name and trick users into opening a malicious attachment.	high	sos5.1.0
POP3:FAILURE:BRUTE-FORCE	This protocol anomaly is multiple login failures within a short period of time between a unique pair of hosts. The default is 4.	high	sos5.1.0
POP3:OUTLOOK:TROUBLE-QUERY-OF	This signature detects buffer overflow attempts against an ActiveX control in Microsoft Outlook. The Local Troubleshooter ActiveX control has inadequate bounds for checking for its Query function, and this exploit bypasses normal Outlook/IE ActiveX security controls. Attackers may create a malicious Web site that contains a call to this ActiveX control; this call contains an overly long string that overflows the control buffer, enabling the attacker to gain control of the target system with the user privileges.	high	sos5.1.0
POP3:OVERFLOW:APOP	This protocol anomaly is a POP3 APOP command argument that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
POP3:OVERFLOW:BOUNDARY_OVERFLOW	This protocol anomaly is a message with more than 70 boundary characters.	high	sos5.1.0
POP3:OVERFLOW:COMMAND	This protocol anomaly is a POP3 command that exceeds 4 bytes, the standard length for a POP3 command. This may indicate a nonstandard POP3 client/server or an attacker has gained command-line access to the server.	high	sos5.0.0, sos5.1.0
POP3:OVERFLOW:CONTENT_NAME	This protocol anomaly is a mime header content-type with a name length that is longer than the defined value. The default value is 128.	high	sos5.1.0
POP3:OVERFLOW:FILENAME2LONG	This protocol anomaly is a message with a content_disposition header containing a 'name' attribute value that is too long.	high	sos5.1.0
POP3:OVERFLOW:LINE	This protocol anomaly is a text-line from a POP3 client to the server that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
POP3:OVERFLOW:PASS	This protocol anomaly is a POP3 PASS command argument that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
POP3:OVERFLOW:QPOP-OF1	This signature detects buffer overflow attempts against Qpopper, a POP3 server for Unix. Qpopper 3.0beta20 and earlier versions are vulnerable.	critical	sos5.0.0, sos5.1.0
POP3:OVERFLOW:QPOP-OF2	This signature detects a buffer overflow attempt to exploit a vulnerability in Qpopper. Version 3.0beta30 and many earlier versions are vulnerable.	critical	sos5.0.0, sos5.1.0

POP3:OVERFLOW:QPOP-OF3	This signature detects buffer overflow attempts to exploit a vulnerability in the Qpopper daemon. Some 3.0 beta versions are vulnerable.	critical	sos5.0.0, sos5.1.0
POP3:OVERFLOW:QPOP-OF4	This signature detects a buffer overflow attempt to exploit a vulnerability in Qpopper using custom shellcode. Version 3.0beta20 and many earlier versions are vulnerable.	critical	sos5.0.0, sos5.1.0
POP3:OVERFLOW:TXTLINE_2LONG	This protocol anomaly is a message data line that exceeds the defined maximum length (sc_mime_textline_length).	high	sos5.1.0
POP3:OVERFLOW:USER	This protocol anomaly is a POP3 USER command argument that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
POP3:REQERR:REQ-MESSAGE-NUMBER	This protocol anomaly is a POP3 message number that is unreasonably high. This may indicate a huge mailbox or an exploit attempt.	high	sos5.0.0, sos5.1.0
POP3:REQERR:REQ-SYNTAX-ERROR	This protocol anomaly is an unparsed POP command line or header line. This may indicate a nonstandard e-mail client or server or a backdoor/exploit attempt.	medium	sos5.0.0, sos5.1.0
SCAN:AMAP:FTP-ON-HTTP	This signature detects the scanner tool amap, made by the Hacker's Choice. THC-AMAP is used in initial reconnaissance for an attacker to determine services running on target hosts before launching other attacks.	low	sos5.1.0
SCAN:AMAP:SAP-R3-ON-HTTP	This signature detects the scanner tool AMAP, made by The Hacker's Choice (THC). Attackers may use THC-AMAP during their initial reconnaissance to determine services running on target hosts before launching other attacks.	low	sos5.1.0
SCAN:AMAP:SSL-ON-HTTP	This signature detects the scanner tool AMAP, made by The Hacker's Choice (THC). Attackers may use THC-AMAP during their initial reconnaissance to determine services running on target hosts before launching other attacks.	low	sos5.1.0
SCAN:AMAP:SSL-ON-POP3	This signature detects the scanner tool AMAP, made by The Hacker's Choice (THC). Attackers may use THC-AMAP during their initial reconnaissance to determine services running on target hosts before launching other attacks.	low	sos5.1.0
SCAN:METASPLOIT:SMB-ACTIVE	This signature detects traffic generated by the open-source exploiting tool Metasploit Framework. Other signatures may also trip. This indicates that someone is using this tool on your network. Follow-up investigation of source or target machines may be required.	high	sos5.1.0
SCAN:MISC:HTTP:HTR-OVERFLOW	"This signature detects denial-of-service (DoS) attacks against Microsoft IIS 4.0 and 5.0. Attackers may send maliciously crafted HTR requests (.htr) with long variable names to overflow the buffer in the ism.dll ISAPI extension that implements HTR scripting and create a denial of service or execute arbitrary commands.	medium	sos5.0.0 sos5.1.0

SHELLCODE:AIX:NOOP-PKT	This signature scans PACKETS for at least four in a row AIX NOOP instructions, which are very common in buffer overflow exploits. You may want to apply this signature to all non-TCP traffic to your AIX servers.	medium	sos5.0.0, sos5.1.0
SHELLCODE:BSDX86:GEN-1-PKT	This signature scans PACKETS for an x86 BSD (all flavors) instruction sequence, common in buffer overflow exploits. You may want to apply this signature to all non-TCP traffic to your BSD servers.	high	sos5.0.0, sos5.1.0
SHELLCODE:BSDX86:GEN-2-PKT	This signature scans PACKETS for an x86 BSD (all flavors) instruction sequence, common in buffer overflow exploits. You may want to apply this signature to all non-TCP traffic to your BSD servers.	high	sos5.0.0, sos5.1.0
SHELLCODE:DIGITAL:NOOP-PKT	This signature scans PACKETS for at least four in a row DEC ALPHA NOOP instructions, which are very common in buffer overflow exploits. You may want to apply this signature to all non-TCP traffic to your DEC ALPHA servers.	medium	sos5.0.0, sos5.1.0
SHELLCODE:HP-UX:HP-NOOP-1-PKT	This signature scans PACKETS for a HP-UX PA-RISC instruction sequence, common in buffer overflow exploits. You may want to apply this signature to all non-TCP traffic to your HP-UX servers.	medium	sos5.0.0, sos5.1.0
SHELLCODE:HP-UX:HP-NOOP-2-PKT	This signature scans PACKETS for a HP-UX PA-RISC instruction sequence, common in buffer overflow exploits. You may want to apply this signature to all non-TCP traffic to your HP-UX servers.	medium	sos5.0.0, sos5.1.0
SMB:AUDIT:INV-PROTOCOL	This protocol anomaly is an invalid SMB protocol. The first four bytes of valid SMB messages are 0xff, 'S', 'M', 'B'. This may be a misbehaving client or an attempt to tunnel through the NETBIOS port.	info	sos5.1.0
SMB:CONNECT-FROM-LOCALHOST	This signature detects attempts to remotely connect to SMB shares with the NetBIOS hostname of Localhost. Because Localhost logins are not typically performed over the network, this may indicate that an attacker is trying to bypass host-based access controls.	low	sos5.1.0
SMB:ENUM:NAME-LOOKUP	This protocol anomaly is the \pipe\lsarpc (Local Security Authority) named pipe transaction used to execute the LookupAccountName function. Programs such as user2sid and Hyena use this named pipe transaction to validate usernames on the target host.	medium	sos5.1.0
SMB:ERROR:GRIND	This protocol anomaly is multiple login/authentication failures between a unique pair of hosts within a short period of time. Vulnerability scanners and programs like enum that perform dictionary based or password-guessing attacks will likely trigger this attack.	high	sos5.1.0
SMB:ERROR:INV-MSG-LEN	This protocol anomaly is an invalid session message length in an SMB message.	high	sos5.1.0

SMB:ERROR:MAL-MSG	This protocol anomaly is a malformed SMB message in which the wcount field is larger than the message size.	high	sos5.1.0
SMB:EXPLOIT:ACCOUNT-NAME-OF	This signature detects attempts to overflow the SMB Account Name. ISS BlackICE, Proventia, and RealSecure products are vulnerable to this buffer overflow. A successful attack could give an attacker complete control of these systems.	critical	sos5.1.0
SMB:EXPLOIT:DOT-JOB	This signature detects a Microsoft Task Scheduler (.job) file being copied over an SMB network share. Microsoft Windows XP Service Pack 1 and Microsoft Windows 2000 Service Pack 2 and earlier are vulnerable. Attackers may open a malicious .job file in Task Scheduler to execute arbitrary code and compromise the system.	medium	sos5.1.0
SMB:EXPLOIT:LANMAN-NUKE	This protocol anomaly is a LANMAN request (NetServerEnum, NetServerEnum2, or NetShareEnum) over a named pipe transaction where the max-param-count and/or the max-data-count of the Transaction header is zero. Attackers can use this malformed request to crash an unpatched Microsoft NT, 2000, or XP server.	critical	sos5.1.0
SMB:EXPLOIT:LINUX-TRANS2-OF	This signature detects attempts to exploit a vulnerability in the Server Message Block File System (SMBFS) implemented in the Linux kernel. Kernels 2.4 and 2.6 are vulnerable. Attackers may gain root access on the target host.	high	sos5.1.0
SMB:EXPLOIT:NULL-FILENAME	This protocol anomaly is an empty Filename field in the Delete, Rename, Move or Copy SMBs.	medium	sos5.1.0
SMB:EXPLOIT:NULL-PATH	This protocol anomaly is an empty Path field in the Tree Connect SMB. This may be a misbehaving client or an attempt to exploit vulnerabilities in the SMB server.	medium	sos5.1.0
SMB:EXPLOIT:NULL-SERVICE	This protocol anomaly is an empty Service field in the Tree Connect SMB. This may be a misbehaving client or an attempt to exploit vulnerabilities in the SMB server.	medium	sos5.1.0
SMB:EXPLOIT:REGISTRY-DOS	DI has detected a suspiciously large registry key in the OpenKey function executed using a named-pipe transaction. Large key sizes in the OpenKey function can cause the winlogon.exe process in Window NT 4.0 to crash.	critical	sos5.1.0
SMB:EXPLOIT:SAMBA-DIR-TRAV	This signature detects SMB requests for pathnames that attempt to traverse the server root. Samba 3.0.5 and earlier versions are vulnerable. Malicious users can send "get", "put", and "dir" commands to a Samba server to access files outside the shared directories.	medium	sos5.1.0
SMB:EXPLOIT:WINBLAST-DOS	Microsoft Windows Samba File Sharing Resource Exhaustion Vulnerability	medium	sos5.1.0

SMB:NETBIOS:INV-CDNAME-ENC	This protocol anomaly is an invalid calling name encoding in the NETBIOS header that encapsulates an SMB. NETBIOS names are 16 bytes and may encode to a maximum of 34 bytes.	high	sos5.1.0
SMB:NETBIOS:INV-CDNAME-LEN	This protocol anomaly is an invalid called name length in the NETBIOS header that encapsulates an SMB. NETBIOS names are 16 bytes and may encode to a maximum of 34 bytes.	high	sos5.1.0
SMB:NETBIOS:INV-CGNAME-ENC	This protocol anomaly is an invalid calling name encoding in the NETBIOS header that encapsulates an SMB. NETBIOS names are 16 bytes and may encode to a maximum of 34 bytes.	high	sos5.1.0
SMB:NETBIOS:INV-CGNAME-LEN	This protocol anomaly is an invalid calling name length in the NETBIOS header that encapsulates an SMB. NETBIOS names are 16 bytes and may encode to a maximum of 34 bytes.	high	sos5.1.0
SMB:NETBIOS:INV-SHDR-LEN	This protocol anomaly is an invalid session header length in the NETBIOS header that encapsulates an SMB. The minimum length of an SMB message is 33 bytes.	high	sos5.1.0
SMB:NETBIOS:INV-SNAME-LEN	This protocol anomaly is an invalid session name length in the NETBIOS header that encapsulates an SMB. NETBIOS names are 16 bytes and may encode to a maximum of 34 bytes.	high	sos5.1.0
SMB:NETBIOS:RMT-REG-ACCESS	This signature detects attempts to remotely access the Windows registry. Attackers may use a malicious client to view or modify the contents of the Windows registry.	medium	sos5.1.0
SMB:NETBIOS:SHARE-LVL-SEC	This protocol anomaly is an SMB session with share-level security. A user may gain access to various resources on the server without username or password authentication.	medium	sos5.1.0
SMB:TOOLS:PSEXEC	This signature detects attempts to upload psexec.exe, an SMB tool for uploading and executing programs interactively. This signature also indicates that the psexec.exe has already logged into the system; Psexec.exe can upload itself to the host only after successful login. Worms often use psexec.exe to propagate.	high	sos5.1.0
SMTP:AUDIT:REQ-INVALID-CMD-SEQ	This protocol anomaly is an invalid sequence of SMTP commands, which would normally not be issued by an SMTP client or server. This may indicate an attacker manually trying to exploit an SMTP server	info	sos5.1.0
SMTP:AUDIT:TEXT-LINE	This protocol anomaly is a text line (in the data section) in an SMTP connection that is too long. This may indicate a buffer overflow attempt.	info	sos5.1.0

SMTP:COMMAND:EXPN	This protocol anomaly is an attempt to use the EXPN command. This command is not used by most standard clients and servers and may reveal information about e-mail accounts.	medium	sos5.0.0, sos5.1.0
SMTP:COMMAND:TURN	This protocol anomaly is an attempt to use the TURN command that exchanges the roles of the e-mail client and server. You may want to ban this command and drop the connection, or edit the SMTP attack objects and change their direction to 'BOTH.	'medium	sos5.1.0
SMTP:COMMAND:VRFY	This protocol anomaly is an attempt to use the SMTP VRFY command. This command is not used by most standard clients and servers and may reveal sensitive information about e-mail accounts.	medium	sos5.0.0, sos5.1.0
SMTP:COMMAND:WIZ	This signature detects attempts to determine if the SMTP server supports the "WIZ" command, which may provide anonymous root access.	critical	sos5.0.0, sos5.1.0
SMTP:EMAIL:EUDORA-SPOOF3	This signature detects attempts to spoof an e-mail attachment. Eudora Windows versions prior to up to 6.0.3 are vulnerable. Attackers may send a maliciously crafted e-mail with an illegal "Attachment Converted:" line in the message body to spoof attachments, which can lead to remote code execution.	medium	sos5.1.0
SMTP:EMAIL:EUDORA-SPOOF4	This signature detects attempts to spoof an e-mail attachment. Eudora Windows 6.2.0.7 and earlier versions are vulnerable. Attackers may send a maliciously crafted e-mail with an illegal "Attachment Converted:" line in the message body to spoof attachments, which can enable remote code execution.	medium	sos5.1.0
SMTP:EMAIL:HEADER-FROM-PIPE	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the pipe passthrough vulnerability. Attackers may use the invalid "from " as the return e-mail address to cause Sendmail to reroute data to another program.	medium	sos5.0.0, sos5.1.0
SMTP:EMAIL:HEADER-TO-PIPE	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the pipe passthrough vulnerability. Attackers may use the invalid 'to ' as the return e-mail address to cause Sendmail to reroute data to another program.	medium	sos5.0.0, sos5.1.0
SMTP:EMAIL:MAIL-FROM-PIPE	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the pipe passthrough vulnerability. Attackers may use the invalid "mail from " as the return e-mail address to cause Sendmail to reroute data to another program.	medium	sos5.0.0, sos5.1.0

SMTP:EMAIL:RCPT-TO-DECODE	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the "decode" e-mail alias vulnerability. Attackers may use the invalid 'rcpt to decode' as the "rcpt to" e-mail address to cause Sendmail to reroute data to the program uuencode. Attackers may then send uuencoded data to overwrite files or place an arbitrary .rhosts files onto the system.	high	sos5.0.0, sos5.1.0
SMTP:EMAIL:RCPT-TO-PIPE	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the pipe passthrough vulnerability. Attackers may use the invalid 'rcpt to ' as the "rcpt to" e-mail address to cause Sendmail to reroute data to another program. Some SMTP servers have been shown to use the ' ' character legitimately.	medium	sos5.0.0, sos5.1.0
SMTP:EMAIL:REPLY-TO-PIPE	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the pipe passthrough vulnerability. Attackers may use the invalid 'reply to ' as the "reply to" e-mail address to cause Sendmail to reroute data to another program. This may also be legitimate traffic from several types of SMTP servers.	medium	sos5.0.0, sos5.1.0
SMTP:EXCHANGE:DOS	This signature detects denial-of-service (DoS) attempts that exploit a MIME header vulnerability in Microsoft Exchange Server 5.5. Attackers may send an e-mail message with an empty charset value ("") in the MIME header to cause a denial-of-service (DoS).	high	sos5.0.0, sos5.1.0
SMTP:EXCHANGE:INV_BDAT_CMD	This protocol anomaly is a BDAT command that is not chunk-size.	high	sos5.1.0
SMTP:EXCHANGE:INV_BDAT_SEC_LEN	This protocol anomaly is a BDAT with a chunk-size larger than 0x7fffffff.	high	sos5.1.0
SMTP:EXCHANGE:MAL-VERB-XEXCH50	This signature detects attempts to exploit a vulnerability in Microsoft Exchange Server 5.5 and 2000. The command verb "Xexch50", which is valid only for communication between validated Exchange servers, is handled incorrectly. Attackers may send the command verb with a negative number or a very large positive number to crash the Exchange server, and, in extreme cases with Exchange Server 2000, may also take control of the server.	critical	sos5.1.0
SMTP:EXPLOIT:EUDORA-URL-SPOOF	This signature detects attempts to exploit a vulnerability in the Eudora mail client. By supplying a link containing character entities, an attacker can force Eudora to display a link as something other than what it really is.	low	sos5.1.0
SMTP:EXPLOIT:HCP-QUOTE-SCRIPT	This signature detects attempts to exploit a vulnerability in URL handling with the Microsoft Help and Support Center (HSC) when invoked with an hcp:// URL. By embedding a quote (") character in the URL, HSC can be instructed to load an arbitrary local file or remote web page, which can then be used to execute scripts in the local zone.	high	sos5.1.0

SMTP:EXPLOIT:MIME-TOOLS-EVADE	This signature detects attempts to evade antivirus tools such as MIME Tools, a Linux-based e-mail MIME scanner. The MIME RFC allows for an empty boundary, but most all mail clients use one, while many viruses will not.	medium	sos5.1.0
SMTP:EXT:DOT-386	This signature detects e-mail attachments that have the extension .386 and were sent via SMTP. Because .386s (Windows Enhanced Mode Driver) files can contain executable code, this may indicate an incoming e-mail virus. Attackers may create malicious executables, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-ADE	This signature detects e-mail attachments that have the extension .ade and were sent via SMTP. Because .ADEs (Microsoft Access Project Extension) files can contain macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the macros and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-ADP	This signature detects e-mail attachments that have the extension .adp and were sent via SMTP. Because .ADPs (Microsoft Access Project) files can contain macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the macros and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-BAS	This signature detects e-mail attachments that have the extension .bas and were sent via SMTP. Because .BASs (Microsoft Visual Basic Class Module) files contain executable code, this may indicate an incoming e-mail virus. Attackers may create malicious executables, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-BAT	This signature detects e-mail attachments with the extension '.bat' sent via SMTP. This may indicate an incoming e-mail virus. .BATs (executable files) contain one or more scripts. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-CHM	This signature detects e-mail attachments that have the extension .chm and were sent via SMTP. Because .CHMs (Compiled HTML Help File) files can contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the files and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-CMD	This signature detects e-mail attachments with the extension '.cmd' sent via SMTP. This may indicate an incoming e-mail virus. CMD files contain commands that when executed can cause significant damage to a windows system.	medium	sos5.1.0

SMTP:EXT:DOT-COM	This signature detects e-mail attachments with the extension '.com' sent via SMTP. This may indicate an incoming e-mail virus. .COMs (executable files) contain one or more scripts. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-CPL	This signature detects e-mail attachments with the extension '.cpl' sent via SMTP. This may indicate an incoming e-mail virus. CPLs (Control Panel elements) are standard Microsoft Windows files that contain Windows Control Panel settings. Attackers may hide malicious executables within a CPL file, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-CRT	This signature detects e-mail attachments that have the extension '.crt' and sent received via SMTP. Because .CRTs (Security Certificate) files can contain executable code, this may indicate an incoming e-mail virus. Attackers may create malicious executable code, tricking users into executing the file and infecting the system.	low	sos5.1.0
SMTP:EXT:DOT-EXE	This signature detects e-mail attachments with the extension '.exe' sent via SMTP. This may indicate an incoming e-mail virus. .EXEs (executable files) contain one or more scripts. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-GRP	This signature detects GRP files sent over SMTP. GRP files can contain Windows Program Group information, and may be exploited by malicious users to deposit instructions or arbitrary code on a target's system. User involvement is required to activate GRP files; typically they are attached to a harmless-appearing e-mail message.	medium	sos5.1.0
SMTP:EXT:DOT-HLP	This signature detects e-mail attachments that have the extension '.hlp' and sent received via SMTP. Because .HLPs (Help File) files can contain macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the macros and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-HT	This signature detects e-mail attachments with the extension '.ht' sent via SMTP. This may indicate an incoming e-mail virus or other attack. HT files contain configuration information for the Hyperterm console program, shipped with every Windows operating system since Windows 95. It is the default handler program for .ht files. A recent vulnerability in Hyperterm could allow an attacker to take control of your computer via an infected .ht file. These files are not normally sent via e-mail.	medium	sos5.1.0
SMTP:EXT:DOT-INF	This signature detects e-mail attachments that have the extension '.inf' and were sent via SMTP. Because .INFs (Setup Information) files contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	medium	sos5.1.0

SMTP:EXT:DOT-INS	This signature detects e-mail attachments that have the extension .ins and were sent via SMTP. Because .INSS (Internet Naming Service) files contain configuration parameters, this may indicate an incoming e-mail virus. Attackers may include malicious configurations, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-ISP	This signature detects e-mail attachments that have the extension .isp and were sent via SMTP. Because .ISPs (Internet Communication Settings) files contain configuration parameters, this may indicate an incoming e-mail virus. Attackers may include malicious configurations, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-JS	This signature detects e-mail attachments that have the extension .js and were sent via SMTP. Because .JSs (JavaScript File) files contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-JSE	This signature detects e-mail attachments that have the extension .jse and were sent via SMTP. Because .JSEs (JavaScript Encoded) files contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-LNK	This signature detects e-mail attachments that have the extension .lnk and were sent via SMTP. Because .LNKS (Windows link) files can point to any program, this may indicate an incoming e-mail virus. Attackers may create a link pointing to a dangerous program, tricking users into executing the link and affecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-MDB	This signature detects e-mail attachments that have the extension .mdb and were sent via SMTP. Because .MDBs (MS Access Application) files can contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	low	sos5.1.0
SMTP:EXT:DOT-MDE	This signature detects e-mail attachments that have the extension .mde and were sent via SMTP. Because .MDEs (Microsoft Access MDE database) files can contain scripts and macros, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-MSC	This signature detects e-mail attachments that have the extension .msc and were sent via SMTP. Because .MSCs (Microsoft Common Console Document) files can contain configuration information, this may indicate an incoming e-mail virus. Attackers may change the configuration to point to a dangerous command, tricking users into executing the files and infecting the system.	medium	sos5.1.0

SMTP:EXT:DOT-MSI	This signature detects e-mail attachments with the extension .msi sent via SMTP. This may indicate an incoming e-mail virus. .MSIs (Microsoft Windows Installer Package) contain executable code. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-MSP	This signature detects e-mail attachments with the extension .msp sent via SMTP. This may indicate an incoming e-mail virus. .MSPs (Microsoft Windows Installer Patch) contain executable code. Attackers may create malicious executables, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-OCX	This signature detects e-mail attachments that have the extension .ocx and were sent via SMTP. Because .OCXs (Object Control Extension) files can contain multiple scripts, this may indicate an incoming e-mail virus. Attackers may create malicious executables, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-PCD	This signature detects e-mail attachments that have the extension .pcd and were sent via SMTP. Because .PCDs (Photo CD MS Compiled Script) files can contain scripts, this may indicate an incoming e-mail virus. Attackers may create malicious scripts, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-PIF	This signature detects e-mail attachments with the extension '.pif' sent via SMTP. This may indicate an incoming e-mail virus. PIFs (Program Information Files) are standard Microsoft Windows files that contain start up properties for DOS applications. Attackers may hide malicious executables within a PIF file, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-REG	This signature detects e-mail attachments that have the extension .reg and were sent via SMTP. Because .REGs (Registry Entries) files contain entries for the Registry, this may indicate an incoming e-mail virus. Attackers may create malicious entries, tricking users into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-SCR	This signature detects e-mail attachments with the extension '.scr' sent via SMTP. This may indicate an incoming e-mail virus. SCR files (ScreenSaver files) are renamed '.exe' files containing executable code. Attackers may disguise malicious executables to appear as harmless screensaver files, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-SCT	This signature detects e-mail attachments with the extension .sct sent via SMTP. This may indicate an incoming e-mail virus. .SCTs (Windows Script Component) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	medium	sos5.1.0

SMTP:EXT:DOT-URL	This signature detects e-mail attachments with the extension .url sent via SMTP. This may indicate an incoming e-mail virus. .URLs (Internet Shortcut) contain a link to a web location. Attackers may create a malicious shortcut, tricking the user into executing the file and send the user to a malicious website.	low	sos5.1.0
SMTP:EXT:DOT-VB	This signature detects e-mail attachments with the extension .vb sent via SMTP. This may indicate an incoming e-mail virus. .VBs (VBScript File) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-VBS	This signature detects e-mail attachments with the extension '.vbs' sent via SMTP. This may indicate an incoming e-mail virus. VBSs (Visual Basic files) contain one or more executable scripts. Attackers may create malicious VB files, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-WMF	This signature detects metafiles (files with .emf or .wmf extensions) in an e-mail attachment. Some versions of Microsoft Windows produce boundary errors when processing metafiles, enabling attackers to create a denial of service (DoS) and execute arbitrary code.	info	sos5.1.0
SMTP:EXT:DOT-WSC	This signature detects e-mail attachments with the extension .wsc sent via SMTP. This may indicate an incoming e-mail virus. .WSCs (Windows Script Component) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-WSF	This signature detects e-mail attachments with the extension .wsf sent via SMTP. This may indicate an incoming e-mail virus. .WSFs (Windows Script File) contain scripts. Attackers may create malicious scripts, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:DOT-WSH	This signature detects e-mail attachments with the extension .wsh sent via SMTP. This may indicate an incoming e-mail virus. .WSHs (Windows Script Host Settings File) contain configuration parameters. Attackers may create malicious configurations, tricking the user into executing the file and infecting the system.	medium	sos5.1.0
SMTP:EXT:JOB	This signature detects an attached Microsoft Task Scheduler (.job) file. Opening a malicious .job file in Task Scheduler may allow for arbitrary code execution, leading to system compromise. This vulnerability is present in Microsoft Windows 2000 Service Pack 2 and later. It is also present in Microsoft Windows XP Service Pack 1.	medium	sos5.1.0

SMTP:IIS:IIS-ENCAPS-RELAY	This signature detects attempts to exploit a vulnerability in the Microsoft SMTP Service in Microsoft IIS. Versions 4.0 and 5.0 are vulnerable. A maliciously crafted 'rcpt to:' command can circumvent e-mail relaying rules. Attackers may impersonate trusted e-mails or send spam anonymously.	medium	sos5.0.0, sos5.1.0
SMTP:INVALID:2MANY-BOUNDARY	This protocol anomaly is an SMTP boundary depth that exceeds the user-defined maximum. The boundary depth indicates the number of nested attachments in a MIME multipart message. The default boundary depth is 4.	medium	sos5.1.0
SMTP:INVALID:BASE64-CHAR	This protocol anomaly is an SMTP message with base64 encoding that contains an invalid character.	high	sos5.1.0
SMTP:INVALID:BOUNDARY-MISS	This protocol anomaly is an SMTP message with a content-type multipart that has no boundary parameter. The boundary parameter specifies a text string that is used to delimit the parts of the multipart message.	medium	sos5.1.0
SMTP:INVALID:DUP_AUTH	This protocol anomaly is multiple AUTH commands within a single SMTP transaction.	medium	sos5.1.0
SMTP:INVALID:DUP-BOUNDARY	This protocol anomaly is an SMTP message with a MIME multipart content-type that uses duplicate boundaries.	high	sos5.1.0
SMTP:INVALID:UNFIN-MULTIPART	This protocol anomaly is an SMTP message with a MIME multipart boundary that exceeds actual multipart data (all data is processed but unfinished boundary delimiters exist).	high	sos5.1.0
SMTP:MAJORDOMO:COMMAND-EXEC	This signature detects attempts to send shell commands via an SMTP e-mail message by exploiting the back-tick (`) vulnerability in Great Circle Associates Majordomo, a perl-based Internet e-mail list server. When processing a list command, Majordomo compares the "reply to" e-mail address again the advertise/noadvertise lists (if configured). During this comparison, Majordomo may be tricked into executing commands when it expands the back-tick operator (used by UNIX to enclose executable commands in a shell command line). Attackers may use the back-tick operator in the "reply to" e-mail header to execute arbitrary commands on the server.	high	sos5.0.0, sos5.1.0
SMTP:MAL:ACROBAT-UUEXEC	This signature detects a maliciously crafted PDF file attached to an e-mail. Attackers may insert certain shell metacharacters at the beginning of a uuencoded PDF file to force Adobe Acrobat to execute arbitrary commands upon loading the file.	high	sos5.1.0
SMTP:MAL:EMAIL-URL-HIDING-ENC	This signature detects attempts to exploit a vulnerability in Microsoft Outlook Express. Attackers may embed binary control characters in a URL that is included in an e-mail; when the URL is viewed, these control characters prevent Outlook Express and Internet Explorer from displaying the complete URL, which may have malicious content.	high	sos5.1.0

SMTP:MAL:NOTES-BIGMAIL	This signature detects large e-mail messages (>12 MB) sent to Lotus Domino servers via a commonly published exploit. Attackers may cause Lotus Domino to exhaust all system memory and cause the service to stop responding.	medium	sos5.1.0
SMTP:MAL:OUTLOOK-MAILTO-QUOT	This signature detects attempts to exploit a vulnerability in the Outlook 2002 mail client. Attackers may use mailto: URLs that contain " strings to execute arbitrary script commands, enabling them to execute code remotely.	high	sos5.1.0
SMTP:MAL:SQM-CONTENT-XSS	This signature detects attempts to exploit a vulnerability in SquirrelMail, a PHP4 Webmail package. Attackers may send e-mail messages that contain Javascript in the Content-Type field; when SquirrelMail receives the message, it may interpret and execute the Javascript, enabling the attacker to compromise the target system.	medium	sos5.0.0, sos5.1.0
SMTP:MDAEMON:SEND-OF	This signature detects buffer overflow attempts against the MDAemon mail server. MDAemon 6.7.9 and older versions are vulnerable. Attackers may send an overly long SMTP SAML, SOML, or SEND command to overflow the buffer and crash the MDAemon service; attackers may also obtain complete control of the server with SYSTEM level access.	high	sos5.1.0
SMTP:MSSQL-WORM-EMAIL	This signature detects attempts to send an e-mail to ixltd@postone.com. This may indicate the presence of SQLsnake, a MSSQL worm. SQLsnake infects Microsoft SQL Servers that have SA (administrative) accounts without passwords. The worm sends a password list and other system information via e-mail to ixltd@postone.com, then begins scanning for vulnerable hosts listening on TCP/1433.	critical	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:BOUNDARY	This protocol anomaly is an SMTP message with a boundary length that exceeds 70 characters. The SMTP RFC specifies 70 as the maximum number of characters in a boundary.	medium	sos5.1.0
SMTP:OVERFLOW:COMMAND-LINE	This protocol anomaly is a text line (in the command section, before the DATA command) in an SMTP connection that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:CONTENT-NAME	This protocol anomaly is an SMTP content-type name that exceeds the user-defined maximum. The default number of bytes in a content-type name is 128.	high	sos5.1.0
SMTP:OVERFLOW:EMAIL-ADDRESS	This protocol anomaly is an e-mail address that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:EMAIL-DOMAIN	This protocol anomaly is a domain name within an e-mail address (for example, localhost.localdomain in root@localhost.localdomain) that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:EMAIL-USERNAME	This protocol anomaly is a user name within an e-mail address (for example, root in root@localhost.localdomain) that is too long. This may indicate a buffer overflow attempt.	high	sos5.0.0, sos5.1.0

SMTP:OVERFLOW:FILENAME	This protocol anomaly is an SMTP content-disposition filename that exceeds the user-defined maximum. The default number of bytes in a content-disposition filename is 128.	high	sos5.1.0
SMTP:OVERFLOW:METAMAIL-HDR-FS2	This signature detects SMTP messages with headers that contain format string errors. Metamail 2.7 and earlier versions are vulnerable. Because Metamail does not handle SMTP headers correctly, attackers may send maliciously crafted SMTP messages to execute arbitrary code at the same privilege level as the target (typically user). Note: Systems that typically carry non-English e-mail messages should not include this attack object in their security policy.	high	sos5.1.0
SMTP:OVERFLOW:METAMAIL-HDR-OF1	This signature detects SMTP messages with large headers that contain character set information. Metamail 2.7 and earlier versions are vulnerable. Because Metamail does not handle SMTP headers correctly, attackers may send maliciously crafted SMTP messages to execute arbitrary code at the same privilege level as the target (typically a user). Note: Systems that typically carry non-English e-mail messages should not include this attack object in their security policy.	high	sos5.1.0
SMTP:OVERFLOW:METAMAIL-HDR-OF2	This signature detects SMTP messages with large headers that contain character set information. Metamail 2.7 and earlier versions are vulnerable. Because Metamail does not handle SMTP headers correctly, attackers may send maliciously crafted SMTP messages to execute arbitrary code at the same privilege level as the target (typically a user). Note: Systems that typically carry non-English e-mail messages should not include this attack object in their security policy.	high	sos5.1.0
SMTP:OVERFLOW:OUTLOOK-CERT-OF	This signature detects buffer overflow attempts against Microsoft Outlook Express, which ships with Internet Explorer 5.5. Attackers may send a maliciously crafted e-mail to a host; if the host opens the e-mail in Outlook Express, attackers may execute arbitrary code on the host.	high	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:REPLY-LINE	This protocol anomaly is a server reply line in an SMTP connection that is too long. This may indicate a buffer overflow attempt by a compromised or malicious SMTP server.	high	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:SENDMAIL-CMT-OF1	This signature detects attempts to exploit a vulnerability in Sendmail. Sendmail versions 5.79 to 8.12.7 are vulnerable. Attackers may include multiple empty address containers in an SMTP header field to overflow the SMTP header buffer and force Sendmail to execute arbitrary code on the host; attackers may obtain root access.	critical	sos5.0.0, sos5.1.0
SMTP:OVERFLOW:SENDMAIL-CMT-OF2	This signature detects attempts to exploit a vulnerability in Sendmail. Sendmail versions 5.79 to 8.12.7 are vulnerable. Attackers may include multiple empty address containers in an SMTP header field to overflow the SMTP header buffer and force Sendmail to execute arbitrary code on the host.	critical	sos5.0.0, sos5.1.0

SMTP:OVERFLOW:SENDMAIL-MIME-OF	This signature detects buffer overflow attempts against Sendmail. Sendmail versions 8.8.0 and 8.8.1 are vulnerable. Attackers may embed a maliciously crafted MIME header in an e-mail to overflow a buffer in Sendmail and execute arbitrary commands as root.	critical	sos5.1.0
SMTP:OVERFLOW:SQRLMAIL-HDR-INJ	This signature detects SMTP messages with Base-64 encoded headers. SquirrelMail 1.4.3a and earlier versions do not correctly sanitize SMTP headers. Attackers may send maliciously crafted SMTP messages to execute arbitrary code at the same privilege level as the target (typically user). Note: Systems that typically carry non-English e-mail messages should not include this attack object in their security policy.	medium	sos5.1.0
SMTP:OVERFLOW:TOO-MANY-RCPT	This protocol anomaly is too many 'RCPT TO:' recipients in an SMTP connection. This may indicate a very popular e-mail message or a DoS/buffer overflow attempt.	medium	sos5.0.0, sos5.1.0
SMTP:REQERR:REQ-SYNTAX-ERROR	This protocol anomaly is an unparsed SMTP command line or header line due to a missing ':'. This may indicate a nonstandard e-mail client or server or a backdoor/exploit attempt.	medium	sos5.1.0
SMTP:RESPONSE:PIPE-FAILED	This signature detects SMTP server responses that are generated when an unsuccessful attempt is made to send shell commands via an SMTP e-mail message by exploiting the pipe () passthrough vulnerability in SendMail. If the ' ' operator was used within specified "mail to" and/or "rcpt to" e-mail addresses to cause Sendmail to reroute data to another program, attackers receive a '550' error message.	high	sos5.1.0
SMTP:SAGTUBE-DOS	This signature detects character strings within an e-mail message that are designed to exploit a vulnerability in SpamAssassin. SpamAssassin Project SpamAssassin 2.63 and earlier are vulnerable. SpamAssassin uses a weighting system to determine when an e-mail message is spam. Attackers may send a maliciously crafted e-mail with a spoofed address to cause SpamAssassin to consider all further e-mail from the spoofed address as spam, regardless of the target's whitelist settings. After the malicious e-mail has been received by the target, SpamAssassin blocks all e-mails from the spoofed address.	medium	sos5.1.0
SMTP:SENDMAIL:ADDR-PRESCAN-ATK	This signature detects attempts to exploit a vulnerability in Sendmail SMTP server versions prior to 8.12.9. Because the prescan() procedure that processes e-mail addresses in SMTP headers does not perform some char and int conversions correctly, attackers may send a maliciously crafted request to corrupt the Address Prescan Memory on a Sendmail SMTP server and execute arbitrary code.	high	sos5.0.0, sos5.1.0
SMTP:SENDMAIL:SENDMAIL-FF-OF	This signature detects attempts to exploit a vulnerability in Sendmail versions 8.12.8 and earlier. Under certain conditions, the Sendmail address parser does not perform sufficient bounds checking when converting char to int. Attackers may use this exploit to gain control of the server.	high	sos5.0.0, sos5.1.0

TROJAN:AUTOPROXY:INFECTED-HOST	This signature detects the AutoProxy trojan attempting to contact a master server and register the IP address and open ports of the infected host. AutoProxy is a trojan that installs a proxy server on Microsoft Windows hosts. Attackers may use an infected host to attack other targets while masking their actual IP address.	critical	sos5.0.0, sos5.1.0
TROJAN:MISC:MOONPIE3-FTP-RESP	This signature detects a banner from the FTP server embedded in the MoonPie backdoor version 3.0 (other versions may also be detected).	high	sos5.1.0
TROJAN:MISC:WANREMOTE-ADMIN	This signature detects access to the WanRemote administration interface using the HTTP protocol.	medium	sos5.1.0
TROJAN:MS-04-028:BACKDOOR-LOGIN	This signature detects login attempts from a client infected with a trojan installed as part of the Microsoft GDI+ Library JPEG overflow exploit.	critical	sos5.1.0
TROJAN:MS-04-028:TOOL-DOWNLOAD	This signature detects attempts by a specific trojan to download files. The trojan, installed as part of the Microsoft GDI+ Library JPEG Overflow exploit, is attempting to download updated files from a remote host.	high	sos5.1.0
TROJAN:PHATBOT:FTP-CONNECT	This signature detects Phatbot FTP connections. Phatbot, a trojan similar to Agobot but with more functionality, sends spam from an infected host machine.	high	sos5.1.0
TROJAN:QAZ:TCP25-CALLING-HOME	This signature detects the string 'nongmin_cn' within an SMTP header-from field sent from a remote system to local server port 25. This may indicate an attacker is attempting to access the Trojan/Worm QAZ. The QAZ Trojan/Worm, famous for infecting the Microsoft network October 2000, allows attackers to access data and gain control over some functions on remote Microsoft Windows systems.	high	sos5.0.0, sos5.1.0
VIRUS:POP3:BABYLONIA	This signature detects e-mail attachments with the file name 'x-mas.exe' sent via POP3. This may indicate the Babylonia e-mail virus is attempting to enter the system. The executed virus infects all files greater than 8kb, installs automatic virus updaters, and allows attackers to further compromise the system by uploading trojans, creating backdoors, etc.	high	sos5.1.0
VIRUS:POP3:BADASS	This signature detects e-mail attachments with the file name 'badass.exe' sent via POP3. This may indicate the BadAss e-mail virus is attempting to enter the system. The executed virus displays a message box with specified text, opens the Microsoft Outlook database, and sends infected messages containing a Dutch phrase to all addresses found.	high	sos5.1.0
VIRUS:POP3:EICAR-ATTACHMENT	This signature detects the EICAR antivirus test file sent as an e-mail attachment.	info	sos5.1.0

VIRUS:POP3:EUROCALCULATOR	This signature detects e-mail attachments with the file name 'Eurocalculator.exe' sent via POP3. This may indicate the Eurocalculator Trojan is attempting to enter the system. The executed file installs a remote administration Trojan similar to Back Orifice, allowing attackers to access data and gain control over some functions on remote Microsoft Windows systems.	critical	sos5.1.0
VIRUS:POP3:EXPLOREZIP-B	This signature detects e-mail attachments with the file name 'zippati.exe' sent via POP3. This may indicate the e-mail virus ExploreZip.B is attempting to enter the system. The executed.ZIP file (zippati in Italian) installs the program explore.exe, which edits the host and visible networked WIN.INI files to run explore.exe on startup. The virus also searches all local and visible networked drives for common file types (.ASN, .C, .CPP, .DOC, .H, .XLS, .PPT) and reduces them to zero bytes.	critical	sos5.1.0
VIRUS:POP3:FIX2001	This signature detects e-mail attachments with the file name 'fix2001.exe' sent via POP3. This may indicate the e-mail virus Fix2001 is attempting to enter the system. The executed file edits the Registry to run the virus on startup, obtains e-mail addresses from sent and received messages, and sends infected e-mail messages to all addresses found. If the virus is patched or corrupted, it also overwrites the C:COMMAND.COM file with a denial-of-service (DoS) (DoS) trojan that erases all drive data upon reboot.	critical	sos5.1.0
VIRUS:POP3:FREELINK	This signature detects e-mail attachments named 'Link.vbs' sent via POP3. This may indicate the VBS.Freelink e-mail virus is attempting to enter the system. The executed virus edits Microsoft Windows Registry entries, opens the Microsoft Outlook database, and sends infected messages to all addresses found.	high	sos5.1.0
VIRUS:POP3:HAPPY99	This signature detects e-mails with the header 'X-Spanska: Yes' and the UU-encoded attachment 'Happy99.exe' sent via POP3. This may indicate the e-mail virus/worm Happy99/Ska is attempting to enter the system. The executed file edits files (notably WSOCK32.DLL) in the system directory, obtains e-mail addresses from sent and received messages, and sends infected e-mail messages to all addresses found. Once WSOCK32.DLL is successfully modified, the virus/worm also exhibits a message box animation routine of a fireworks display.	high	sos5.1.0
VIRUS:POP3:IROK	This signature detects e-mail attachments named 'irok.exe' sent via POP3. This may indicate the e-mail virus Irok is attempting to enter the system. The executed file exhibits a message box animation routine of a starfield while copying itself to the Windows system directory and writing the file Irokrun.vbs to the Startup directory. Upon reboot, the VB script uses Windows Scripting Host (WSH) to open the Microsoft Outlook database and send infected files to up to 60 addresses found. This virus also install the file script.ini to the m IRC directory and use dcc to send irok.exe to IRC clients who join the channel.	high	sos5.1.0

VIRUS:POP3:MATRIX	This signature detects e-mails with the content 'Software provided by Matrix' sent via POP3. This may indicate the e-mail virus Matrix is attempting to enter the system. The executed file first checks for antivirus software running on the host and terminates if found. Otherwise, the virus copies itself to the Windows directory as le_pack.exe, runs, and renames to Win32.dll. Matrix also installs the downloader program Mtx_.exe (which downloads plug-ins for the virus upon reboot), and infects Win32 executables.	critical	sos5.1.0
VIRUS:POP3:MYPICS	This signature detects e-mail attachments named 'pics4you.exe' sent via POP3. This may indicate the e-mail virus MyPics is attempting to enter the system. The executed file installs as Pics4You.exe and writes itself to the Windows Startup directory, obtains e-mail addresses from the Microsoft Outlook database, and sends infected e-mail messages to 50 addresses at a time. MyPics was also designed to corrupt CMOS data and reformat hard drives on 1/1/2000.	high	sos5.1.0
VIRUS:POP3:MYROMEO-BLE-BLA	This signature detects e-mails with the subject 'ble bla' with the attachments myjuliet.chm and myromeo.exe sent via POP3. This may indicate the e-mail virus Verona is attempting to enter the system. Because CHM files are compressed HTML files, myjuliet.chm is activated when viewed in the Microsoft Outlook preview pane; once triggered, the CHM file runs myromeo.exe in the background. Myromeo.exe obtains e-mail addresses from the Microsoft Outlook database, sends infected e-mail messages to all addresses found, and edits the Window directory file hh.dat.	high	sos5.1.0
VIRUS:POP3:MYROMEO-EXE	This signature detects e-mail attachments with the name 'myromeo.exe' accompanied by myjuliet.chm and sent via POP3. This may indicate the e-mail virus Verona is attempting to enter the system. Because CHM files are compressed HTML files, myjuliet.chm is activated when viewed in the Microsoft Outlook preview pane; once triggered, the CHM file runs myromeo.exe in the background. Myromeo.exe obtains e-mail addresses from the Microsoft Outlook database, sends infected e-mail messages to all addresses found, and edits the Window directory file hh.dat.	high	sos5.1.0
VIRUS:POP3:MYROMEO-I-LOVE-YOU	This signature detects e-mails with the attachments myjuliet.chm and myromeo.exe sent via POP3. This may indicate the e-mail virus Verona is attempting to enter the system. Because CHM files are compressed HTML files, myjuliet.chm is activated when viewed in the Microsoft Outlook preview pane; once triggered, the CHM file runs myromeo.exe in the background. Myromeo.exe obtains e-mail addresses from the Microsoft Outlook database, sends infected e-mail messages to all addresses found, and edits the Window directory file hh.dat.	high	sos5.1.0

VIRUS:POP3:MYROMEO-MYJULIET	This signature detects e-mail attachments with the name 'myjuliet.chm' accompanied by myromeo.exe and sent via POP3. This may indicate the e-mail virus Verona is attempting to enter the system. Because CHM files are compressed HTML files, myjuliet.chm is activated when viewed in the Microsoft Outlook preview pane; once triggered, the CHM file runs myromeo.exe in the background. Myromeo.exe obtains e-mail addresses from the Microsoft Outlook database, sends infected e-mail messages to all addresses found, and edits the Window directory file hh.dat.	high	sos5.1.0
VIRUS:POP3:NAVIDAD	This signature detects e-mail attachments named 'navidad.exe' sent via POP3. This may indicate the e-mail virus Navidad is attempting to enter the system. The executed file copies itself as winsvrc.vxd to the Windows system directory and edits the Registry to run the virus on reboot, installs into the system tray, and displays a dialog box with the text 'Ul.' The virus also intercepts new incoming e-mail addresses and sends infected e-mail messages to all senders.	high	sos5.1.0
VIRUS:POP3:NIMDA	This signature detects e-mail attachments named 'readme.exe' sent via POP3. This may indicate the e-mail virus Nimda is attempting to enter the system. The executed file installs to the Windows directory, edits the Registry to run the virus on reboot, and infects Internet-related files. Nimda then obtains e-mail addresses and sends infected messages to all addresses found using its own SMTP server.	critical	sos5.1.0
VIRUS:POP3:PAPA	This signature detects e-mail attachments named 'xpass.xls' sent via POP3. This may indicate the e-mail virus Papa is attempting to enter the system. The executed Microsoft Excel file obtains e-mail addresses from Microsoft Outlook database and sends infected messages to the first 60 addresses found. Papa also attempts to create a denial-of-service (DoS) by pinging the all.net Web server.	critical	sos5.1.0
VIRUS:POP3:PASSION	This signature detects e-mail attachments named ICQ_Greeting.exe sent using POP3. This may indicate the e-mail virus Passion is attempting to enter the system. The executed file copies itself to local root drive, edits the registry to run the virus on reboot, and deletes files. Passion then obtains e-mail addresses from the Microsoft Outlook database and sends infected messages to the first 50 addresses found.	critical	sos5.1.0
VIRUS:POP3:PIKACHU-POKEMON	This signature detects e-mails with the subject 'Pikachu Pokemon' sent via POP3. This may indicate the e-mail virus Pikachu Pokemon is attempting to enter the system. The executed file displays a "friendly" message featuring Pikachu while it overwrites the Autoexec.Bat file to delete most Microsoft Windows 9x system files upon reboot. Pikachu then obtains e-mail addresses from Microsoft Outlook database and sends infected messages to all addresses found.	critical	sos5.1.0

VIRUS:POP3:PRETTY-PARK	This signature detects e-mails with the subject 'C:\CoolProgs\Pretty Park.exe' sent via POP3. This may indicate the e-mail virus Pretty Park is attempting to enter the system. The executed file copies itself to the Windows System directory as FILES32.VXD and edits the Registry to run the virus on reboot. Pretty Park then obtains e-mail addresses from Microsoft Outlook database and sends infected messages to all addresses found every 30 minutes. The virus also attempts to contact its author via IRC chat every 30 seconds; attackers may use the installed virus as a backdoor remote access tool to further compromise the system.	critical	sos5.1.0
VIRUS:POP3:SIMBIOSIS	This signature detects e-mail attachments named 'SETUP.EXE' sent via POP3. This may indicate the e-mail virus Simbiosis (Cholera worm executable containing a CTX virus) is attempting to enter the system. The executed Cholera worm copies itself to the Windows directory and edits either the WIN.INI file (Windows 9x) or the Registry (NT) to run the virus on reboot. Simbiosis then obtains e-mail addresses from Internet-related files and sends infected messages to all addresses found using its own SMTP server. The executed CTX virus appends and infects Microsoft Windows PE executables; the virus does not carry a payload and is apparent only through a video effect.	high	sos5.1.0
VIRUS:POP3:SUPPL	This signature detects e-mail attachments named 'Suppl.doc' sent via POP3. This may indicate the e-mail virus/trojan Suppl is attempting to enter the system. The executed file macros copy the active (virus) document to the Windows directory as Anthrax.ini and decompress the malicious Wsock32.dll file appended to Suppl.doc. On reboot, the virus file DLL.tmp replaces the malicious Wsock32.dll and the original Wsock32.dll is renamed to Wsock33.dll. Suppl then attaches to all outgoing SMTP e-mail messages, locates files with common extensions (DOC, .TXT, .ZIP, etc) on available hard drives, and truncates those files to zero bytes.	critical	sos5.1.0
VIRUS:POP3:THEFLY	This signature detects e-mail attachments named 'The_Fly.chm' sent via POP3. This may indicate the e-mail virus The Fly is attempting to enter the system. The executed file copies itself as THE_FLY.CHM to the Windows directory, as DXGFXB3D.DLL to Windows system directory, and opens a graphic with message 'If you ride a motorcycle, close your mouth'. The Fly then copies MSJSVM.JS to the Windows system directory and edits the Registry to run this JavaScript upon reboot. The virus also obtains e-mail addresses from the Microsoft Outlook database and sends infected messages to all addresses found.	high	sos5.1.0

VIRUS:POP3:TIMOFONICA	This signature detects e-mail attachments named 'Timofonica.txt.vbs' sent via POP3. This may indicate the e-mail virus Timofonica is attempting to enter the system. The executed file creates cmos.com and edits the Registry to run the virus on reboot. When cmos.com is run, it erases CMOS memory, MBRs from the first four physical hard disks, and MBRs and DOS Boot Records of extended partitions. Timofonica also obtains e-mail addresses from the Microsoft Outlook database and sends infected messages to all addresses found. Simultaneously, the virus e-mails the SMS gateway at Moviestar.net and send SMS messages to random cellular phone numbers.	critical	sos5.1.0
VIRUS:POP3:TOADIE	This signature detects e-mail attachments named 'Toadie.exe' sent via POP3. This may indicate the e-mail virus Toadie is attempting to enter the system. The executed file infects EXE files by relocating the initial 7800 bytes to the end of the file, encrypting those bytes, and writing 7800 bytes of its own DOS program to the beginning of the file, thus changing EXE files to DOS files. When run, the virus code first infects more EXE files before passing control. Toadie also replaces unsent e-mail messages in Pegasus Mail, and may send copies of itself via IRC.	high	sos5.1.0
VIRUS:POP3:TRIPLESIX	This signature detects e-mail attachments named '666test.vbs' sent via POP3. This may indicate the e-mail virus TripleSix is attempting to enter the system. The executed file displays three dialogue boxes leading the user through the game "Does your name add up to 666?". The virus then copies WINTMP.TXT to the Windows directory; this file creates WINTMP.EXE (a PkZip executable), which in turn creates 666TEST.ZIP (an archive). The archive is copied to the Windows system directory as WINSWAP.SWP. Triplesix also writes REGSVR.VBS to the Windows system directory and edits the Registry to run that script on reboot. When REGSVR.VBS is activated, it obtains e-mail addresses from the Microsoft Outlook database and sends infected messages to all addresses found, overwrites mIRC and Pirch setup files, and sends infected messages via IRC.	high	sos5.1.0
VIRUS:POP3:TUNE	This signature detects e-mail attachments named 'Tune.vbs' sent via POP3. This may indicate the e-mail virus Tune is attempting to enter the system. The executed file copies itself to the Windows, Windows system, and Temporary directories and edits the Registry to run the virus on reboot. When activated, it obtains e-mail addresses from the Microsoft Outlook database and sends infected messages to all addresses found, overwrites mIRC and Pirch setup files, and sends infected messages via IRC.	high	sos5.1.0

VIRUS:POP3:UUENCODED-DOT-VBS	This signature detects e-mail attachments containing the string 'begin' and the file extension 'vbs' sent via POP3. This may indicate the e-mail virus LoveLetter is attempting to enter the system. The executed file copies itself to the Windows system directory and edits the Registry to run the virus on reboot; when activated, it downloads a trojan from a specified web site that deletes security keys and sends stolen passwords to its owner. LoveLetter also obtains e-mail addresses from the Microsoft Outlook database and sends infected messages to all addresses found, overwrites mIRC and Pirch setup files, and sends infected messages via IRC.	high	sos5.1.0
VIRUS:POP3:WSCRIPT-KAK	This signature detects e-mails containing 'kak.hta' sent via POP3. This may indicate the e-mail virus Kak is attempting to enter the system. The virus arrives embedded within Microsoft Outlook message signature file as kak.htm, and activates when viewed in the Microsoft Outlook preview pane. Once triggered, the file copies itself as kak.hta to the Windows startup and system directories; on reboot, the virus overwrites the autoexec.bat file to delete the virus from the startup directory. Kak then replaces the Microsoft Outlook message signature with infected file kak.htm. The virus also displays an alert box after 6pm on the first day of the month and shows down Windows.	medium	sos5.1.0
VIRUS:POP3:Y2K-ZELU	This signature detects e-mail attachments named 'Y2k.exe' sent via POP3. This may indicate the e-mail virus Zelu is attempting to enter the system disguised as the utility ChipTec Y2K Freeware Version. The executed file scans available directories, corrupts writeable files, and inserts a message at the beginning of infected files. Zelu may reset the system, making the operating system unusable and erasing all data.	critical	sos5.1.0
VIRUS:POP3:ZIPPED	This signature detects e-mail attachments named 'ZippedFiles.exe' sent via POP3. This may indicate the e-mail virus Zipped_Files is attempting to enter the system. The executed.ZIP file installs the program explore.exe, which edits the host and visible networked WIN.INI files to run explore.exe on startup. The virus also searches all local and visible networked drives for common file types (.ASN, .C, .CPP, .DOC, .H, .XLS, .PPT) and reduces them to zero bytes.	critical	sos5.1.0
VIRUS:SMTP:BAGLE.Q-SMTP	This signature detects the Q through T variants of the Bagle SMTP virus. Bagle sends e-mails containing an attachment with a malicious payload. Viewing the e-mail message loads an external link using HTTP; this link is actually an executable program that infects the target. The virus then sends a copy of itself to e-mail addresses found on the target's hard drive using the target's e-mail address as the return address.	high	sos5.1.0
VIRUS:SMTP:DOUBLE-DOT-DOT	This signature detects e-mail attachments that contain two file extensions. Attackers or viruses may send e-mail attachments that use two file extensions to disguise the actual file name and trick users into opening a malicious attachment.	high	sos5.1.0

VIRUS:SMTP:DUMARUJ	This signature detects the J variant of the Dumaru SMTP virus. Dumaru sends e-mails with the subject line: "Important information for you. Read it immediately!"; the e-mail includes a .zip attachment that contains a malicious payload disguised as a picture. When the picture is viewed, the malicious executable program infects the target host. The virus then sends a copy of itself to e-mail addresses found in the target's address book, using the target's e-mail address as the return address.	high	sos5.1.0
VIRUS:SMTP:EICAR-ATTACHMENT	This signature detects the EICAR antivirus test file sent as an e-mail attachment.	info	sos5.1.0
VIRUS:SMTP:EXE-ATTACH-1	This signature detects Win32 executables sent as a MIME attachment. Many viruses, worms, and other malicious programs are transmitted through SMTP attachments. You might want to block all executable attachments and instead require your users to send executables in a compressed format.	medium	sos5.1.0
VIRUS:SMTP:EXE-IN-ZIP	This signature detects Win32 executables sent within a ZIP file as a MIME attachment. Many viruses, worms, and other malicious programs are transmitted through SMTP attachments. You might want to block all executable attachments.	medium	sos5.1.0
VIRUS:SMTP:NAIL	This signature detects attempts by the e-mail virus Nail to enter the system. When executed, the virus assigns the Microsoft Word auto.dot template to a template located on an attacker Web site, enabling the attacker to upload new virus code. Nail then starts a MAPI (Mail API) session, obtains e-mail addresses from the Microsoft Outlook database, and sends infected e-mail messages to all addresses found. Finally, the virus sends an e-mail message to chainnail@hotmail.com, assumed to be the e-mail address of the virus author.	high	sos5.0.0
VIRUS:SMTP:RESUME-EXPLORER-DOC	This signature detects e-mail attachments named 'EXPLORER.DOC' sent via SMTP. This may indicate the e-mail virus Resume is attempting to enter the system. The executed file obtains e-mail addresses from Microsoft Outlook database and sends infected messages to all addresses found. When the file is closed, Resume creates directory C:\Data, copies itself there as Normal.dot, and edits the Registry to run the virus on reboot. The virus then attempts to delete all files from several directories (including Windows) and all drives from A: to Z:.	low	sos5.1.0
VIRUS:SMTP:SOBIG-ATTACHMENTS	This signature detects e-mail attachments with one of the following file name sent via SMTP: approved.pif, application.pif, doc_details.pif, movie28.pif, password.pif, ref-39xxxx.pif, screen_doc.pif, screen_temp.pif, _approved.pif. This may indicate the SOBIG e-mail virus is attempting to enter the system.	medium	sos5.1.0

VOIP:MGMT:XPRESSA-HTTP-DOS	This signature detects attempts to exploit a vulnerability in Pingtel Xpressa phones. Attackers may supply an overly long request to the HTTP management server on the phone to execute arbitrary code or crash the phone (the phone must be rebooted).	medium	sos5.1.0
WORM:AGOBOT:HTTP-SHARE-ENUM	This signature detects attempts by the Agobot worm to enumerate SMB shares via HTTP.	medium	sos5.1.0
WORM:AGOBOT:PY-HTTP-PROP	This signature detects the PY variant of the Agobot worm as it attempts to infect another host. This signature could be prone to false positives.	high	sos5.1.0
WORM:BAGLE:AF-HTTP	This signature detects the AF variant of the Bagle SMTP virus. Bagle sends e-mails that contain an attachment with a malicious payload. When the attachment is viewed, the payload uses HTTP to load an external link, which is actually an executable program that infects the target host. The virus then sends a copy of itself to e-mail addresses found on the target's hard drive, using the target's e-mail address as the return address.	high	sos5.1.0
WORM:BAGLE:AF-SMTP	This signature detects the AF variant of the Bagle SMTP virus. Bagle sends e-mails to victims with an attachment with malicious payload. Attempting to view the attachment, which is actually an executable program, infects the user. The virus then sends a copy of itself to e-mails found searching the victim's hard drive for addresses, with the victim's e-mail address as the return address.	high	sos5.1.0
WORM:BERBEW:KEYLOGGER-UPLOAD	This signature detects the Berbew worm as it uploads keylogger information to a listening post. Berew monitors user keystrokes for financial data and reports that information to an attacker via HTTP to a listening post. Source IP addresses that trigger this signature are extremely likely to be infected with the Berbew worm.	high	sos5.0.0, sos5.1.0
WORM:BOBAX:C-PHONE-HOME-DNS	This signature detects Bobax worm activity. The C variant of the Bobax worm attempts to lookup the correct IP addresses for listening post servers set up by the Bobax virus authors. Because lookups for these addresses are extremely suspicious, you should investigate the source device for Bobax infection. However, this signature detects Bobax activity (not Bobax infection attempts), and cannot be used to prevent Bobax infection. To prevent Bobax infection, configure your security policy to drop traffic that matches the signatures "Windows RPC: LSASS Malicious OpCode" and "Windows RPC: LSASS DCE-RPC Oversized Fragment".	high	sos5.1.0
WORM:CODERED:INFECTION-ATTEMPT	The signature detects attempts to infect an Microsoft IIS server with the Code Red worm using a .ida buffer-overflow attack. The installed worm downloads code from the donor host, creates a backdoor on the victim, and sets up 100 threads of the worm that scan for other vulnerable hosts using random IP addresses. Code Red also checks the host system time; on the 20th of each month (GMT), all infected systems send 100k bytes of data to TCP/80 of www.whitehouse.gov, causing a denial-of-service (DoS).	medium	sos5.0.0, sos5.1.0

WORM:CODERED-2:CMD-BACKDOOR	This signature detects attempts to access a backdoor web script installed by the Code Red II worm. The Code Red II worm, like the original Code Red worm, allows attackers to remotely access the server.	medium	sos5.0.0, sos5.1.0
WORM:CODERED-2:INFECT-ATTEMPT	This signature detects attempts by the CodeRedII worm to infect a host. The CodeRedII worm, also known as CodeRed.F, exploits the same vulnerability as the original CodeRed worm.	high	sos5.0.0, sos5.1.0
WORM:CODERED-2:ROOT-BACKDOOR	This signature detects attempts to access a backdoor web script installed by the Code Red II worm. The Code Red II worm, like the original Code Red worm, allows attackers to remotely access the server.	medium	sos5.0.0, sos5.1.0
WORM:EMAIL:BAGLE-INFECTION	This signature detects the Bagle worm activity on a host. After infecting a host, the Bagle worm attempts to contact a Web server listening post. The Bagle worm, which affects Microsoft Windows, copies itself to the system directory, and edits the system registry. The worm uses an e-mail attachment to propagate itself to other hosts, and has a hard-coded expiration date (January 28). This signature could be prone to false positives.	medium	sos5.0.0, sos5.1.0
WORM:EMAIL:W32.SOBIG.E	This signature detects e-mail attachments containing the W32.Sobig.E worm sent via SMTP.	medium	sos5.1.0
WORM:MIMAIL:MIMAIL.A	This signature detects the Mimapil.A worm attachment in SMTP traffic. After infecting a Windows-based host, Mimapil sends itself as an attachment to another target using its own SMTP engine.	high	sos5.1.0
WORM:MIMAIL:MIMAIL.L	This signature detects the Mimapil.L worm attachment in SMTP traffic. After infecting a Windows-based host, Mimapil sends itself as an attachment to another target using its own SMTP engine.	high	sos5.1.0
WORM:MOFEI:MOFEI-B-PROPAGATION	This signature detects the MoFei worm attempting to propagate to another host. After infecting a host, the MoFei worm propagates by depositing a copy of itself in a vulnerable NetBIOS folder on another host. The MoFei worm is known by several aliases, including W32.Mofei-B and W32.Femot.D.	high	sos5.1.0
WORM:NACHI:B-C-D-INFECT-ATTEMPT	This signature detects infection attempts of the Windows RPC Locator Service by the B, C or D variants of the Nachi worm. This signature only triggers on a successful connect to an accessible victim. Follow up is strongly suggested.	critical	sos5.1.0
WORM:NACHI:D-WEBDAV-ATK	This signature detects WebDAV overflows, which can indicate an infection attempt by the Nachi worm (D variant). Nachi.D, a worm, typically attempts to infect the target host by exploiting several vulnerabilities.	high	sos5.0.0, sos5.1.0

WORM:NETSKY:V-SMTP-PROP	This signature detects the V variant of the NetSky worm. The V variant encodes a malicious HTML script in the body of an e-mail sent to the target host. Due to a known vulnerability, Microsoft Outlook and Outlook Express process the encoded script when the e-mail appears in the preview pane (the e-mail does not need to be opened). The script downloads the NetSky worm from known Internet sites and installs the worm on the target host.	high	sos5.1.0
WORM:NIMDA:BIN-255-CMD	This signature detects attempts to infect a Microsoft IIS Web server with the Nimda worm. Nimda may infect other Web servers by obtaining e-mail addresses and sending a copy of itself in infected messages using its own SMTP or POP3 server; adding files to a system configured to allow Windows file shares; or posting an infected HTML e-mail to the Web server where it can be accessed via HTTP.	medium	sos5.0.0, sos5.1.0
WORM:NIMDA:MSADC-ROOT	This signature detects attempts to infect a Microsoft IIS Web server with the Nimda worm. Nimda may infect other Web servers by obtaining e-mail addresses and sending a copy of itself in infected messages using its own SMTP or POP3 server; adding files to a system configured to allow Windows file shares; or posting an infected HTML e-mail to the Web server where it can be accessed via HTTP.	medium	sos5.0.0, sos5.1.0
WORM:NIMDA:NIMDA-EML	This signature detects attempts to create .EML files on the system, a common sign of the NIMDA worm. The worm browses remote directories and creates .EML files (the worm's multi-part messages containing a MIME-encoded worm) with the same names as existing documents or Web page files.	medium	sos5.1.0
WORM:NIMDA:NIMDA-NWS	This signature detects attempts to create a .NWS file on the system, a common sign of the NIMDA worm. The worm browses remote directories and creates .NWS files (the worm's multi-part messages containing a MIME-encoded worm) with the same names as existing documents or Web page files.	medium	sos5.1.0
WORM:NIMDA:NIMDA-RICHED20	This signature detects attempts to create the file RICHED20.DLL on the system, a common sign of the NIMDA worm. The worm may overwrite the original RICHED20.DLL in the Windows systems folder with a binary copy of itself, and place additional copies in all folders containing .DOC or .EML files.	high	sos5.1.0
WORM:NIMDA:SCRIPTS-C11C-CMD	This signature detects attempts to infect a Microsoft IIS Web server with the Nimda worm. Nimda may infect other Web servers by obtaining e-mail addresses and sending a copy of itself in infected messages using its own SMTP or POP3 server; adding files to a system configured to allow Windows file shares; or posting an infected HTML e-mail to the Web server where it can be accessed via HTTP.	medium	sos5.0.0, sos5.1.0

WORM:NIMDA:SCRIPTS-CMD	This signature detects attempts to infect a Microsoft IIS Web server with the Nimda worm. Nimda may infect other Web servers by obtaining e-mail addresses and sending a copy of itself in infected messages using its own SMTP or POP3 server; adding files to a system configured to allow Windows file shares; or posting an infected HTML e-mail to the Web server where it can be accessed via HTTP.	medium	sos5.0.0, sos5.1.0
WORM:NIMDA:SCRIPTS-ROOT	This signature detects attempts to infect a Microsoft IIS Web server with the Nimda worm. Nimda may infect other Web servers by obtaining e-mail addresses and sending a copy of itself in infected messages using its own SMTP or POP3 server; adding files to a system configured to allow Windows file shares; or posting an infected HTML e-mail to the Web server where it can be accessed via HTTP.	medium	sos5.0.0, sos5.1.0
WORM:PHPINCLUDE:SEARCH-REQ	This signature detects the Santy.C worm attempting to find targets by sending a search request to a Google or Yahoo search engine.	medium	sos5.1.0
WORM:SANTY:GOOGLE-SEARCH	This signature detects a machine infected with the Santy worm querying Google to locate new targets for infection. The source IP of this log is likely infected with a variant of Santy.	medium	sos5.1.0
WORM:SANTY:INFECT-ATTEMPT	This signature detects a machine infected with the Santy worm attempting to infect a new target host. The source IP of this log is likely infected with a variant of Santy.	high	sos5.1.0
WORM:SMB:DELODER	This signature detects attempts to upload the deloder worm. This signature also indicates that the worm has already logged into the system; the deloder worm can upload itself to the host only after successful login as Administrator (deloder uses one of 50 default passwords to login).	critical	sos5.1.0
WORM:SMB:W32-SLACKOR	This signature detects SMB transmissions of the W32/Slackor worm, which targets file shares. The worm scans the /16 of the infected host for systems listening on TCP/445; if a system is found, the worm uses pre-programmed usernames and passwords to connect to the \$IPC share on the system, copies itself to the C:\sp directory, and runs its payload.	high	sos5.1.0

Configuration Log Entries

The Configuration category contains the subcategories shown in [Table 127 on page 997](#):

Table 127: Configuration Log Entries

Configuration Log Entry Subcategories	ScreenOS Message ID
Address	Addresses > Notification > 00001
Admin	Admin > Notification > 00002

Table 127: Configuration Log Entries (continued)

Configuration Log Entry Subcategories	ScreenOS Message ID
Auth	Auth > Notification > 00015
Clock	System > Notification > 00014
CLS	Notification > 00043
CMS	Device > Notification > 00022
Console	Admin > Notification > 00003
DHCP CLI	DHCP > Notification > 00027
DHCP IP	DHCP > Notification > 00009
DHCP Opt	DHCP > Notification > 00024
DIP	DIP > Notification > 00021
DNS	DNS > Notification > 00004
DNS REP	DNS > Notification > 00029
Erase	System > Notification > 00023
Hostname	System > Notification > 00006
Interface	Interface > Notification > 00009
MIP	MIP > Notification > 00021
NSRP	High Availability > Notification > 00007
OSPF	OSPF > Notification > 00038
PKI	PKI > Notification > 00002
Policy	Policies > Notification > 00018
PPP	HDLC > Notification > 00042
PPPoE	PPPoE > Notification > 00034
RIP	RIP > Notification > 00045
Route	Route > Notification > 00011
Route Map	Route > Notification > 00048

Table 127: Configuration Log Entries (continued)

Configuration Log Entry Subcategories	ScreenOS Message ID
Schedule	Schedule > Notification > 00020
Service	Service > Notification > 00012
Set ARP	ARP > Notification > 00051
Shaper	Traffic Shaping > Notification > 00002
SIP ALG	Flow > Notification > 00047
SME	NSM > Notification > 00033
SNMP	SNMP > Notification > 00031
S/W Key	Entitlement > Notification > 00036
SSH	SSHv2 > Notification > 00026
SSL	SSL > Notification > 00035
Syslog	Syslog and WebTrends > Notification > 00019
Track IP	High Availability > Notification > 00050
URL	WEB Filtering > Notification > 00013
User	User > Notification > 00014
VPN	VPN > Notification > 00017
Vrouter	Virtual Router > Notification > 00049
Vsys	Vsys > Notification > 00032
Zone	Zone > Notification > 00037
Set ARP Always On Dest	ARP > Notification > 0005
Unset ARP Always On Dest	ARP > Notification > 00054

Information Log Entries

The Information category contains the subcategories shown in [Table 128 on page 1000](#):

Table 128: Information Log Entries

Information Log Entry Subcategories	ScreenOS Message ID
Auth Challenge	Auth > Information > 00546
Auth Failed	Auth > Warning > 00518
Auth Status Change	Auth > Information > 00525
Auth Passed	Auth > Warning > 00519
Auth Timeout	Auth > Warning > 00520
Anti Virus	AntiVirus Scanning (External) > Information > 00547
BGP	BGP > Information > 00542
Clock	NTP > Notification > 00531
Configuration Size	System > Notification > 00553
Device Connect	N/A
Device Disconnect	N/A
DHCP CLI	DHCP > Information > 00530
DHCP DNS	DNS > Information > 00004
Generic	System > Information > 00767
VIP Svr Up	VIP > Notification > 00533
Link Status	Interface > Notification > 00513
Log Cleared	Logging > Information > 00534
NSRD	NSRD > Information > 00551
NTP failure	NTP > Notification > 00531
NTP timeout	NTP > Notification > 00531
OSPF	OSPF > Information > 00541
Password Change	Admin > Information > 00002
PKI	PKI > Information > 00535
PPP	PPP > Notification > 00539

Table 128: Information Log Entries (continued)

Information Log Entry Subcategories	ScreenOS Message ID
PPPoE	PPPoE > Notification > 00034
RIP	RIP > Information > 00544
SME	NSM > Information > 00538
SNMP	SNMP > Information > 00524
SSH	SSHv1 > Information > 00528 SSHv2 > Information > 00026
SSL	SSL > Information > 00284
URL Blk	WEB Filtering > Notification > 00523
Username Change	Admin > Information > 00002
VPN	VPN > Information > 00536 L2TP > Information > 00536 IKE > Information > 00536
VIP Server Status	VIP > Notification > 00533
DHCP Server Status	DHCP > Information > 00527



NOTE: For security devices running ScreenOS 5.0.x or higher, Network and Security Manager does not generate information logs for device connect and disconnect events. The **Realtime Monitor** however, does display the correct up/down status of the device.

Self Log Entries

Self log entries appear in the **Log Viewer** under the category Self, which contains a single subcategory: Self Log.

Traffic Log Entries

Traffic log entries appear in the **Log Viewer** under the category Traffic, which contains a single subcategory: Traffic Log.

GTP Log Entries

When you enable logging in a GTP object, you can configure a security device to create log entries with Basic or Extended information. Additionally, when counting is also enabled the GTP object, the device also generates log entries for deleted GTP tunnels.

For log entries generated by GTP objects with Basic logging enabled, you can view the following information:

- Timestamp
- Source IP address
- Destination IP address
- TID (Tunnel Identifier) or TEID (Tunnel Endpoint Identifier)
- Message type
- Packet status: forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited
- Interface, vsys, or vrouter name (if applicable)

For log entries generated by GTP objects with Extended logging enabled, you can view the following information:

- IMSI
- MSISDN
- APN
- Selection Mode
- SGSN address for signaling
- SGSN address for user data
- GGSN address for signaling
- GGSN address for user data

For log entries generated by deleted GTP tunnels, you can view the following information:

- Timestamp
- Interface name (if applicable)
- SGSN IP address
- GGSN IP address
- TID
- Tunnel duration time in seconds
- Number of messages sent to the SGSN
- Number of messages sent to the GGSN